

Management science for managing risk

An introduction to useful techniques

By Matthew Leitch, 27 March 2012 and expanded 2 May 2012.

Everything most people and organizations need to know to manage 'risk' better is known already and part of management science¹. The alternative to the time-wasting torture of making risk registers is long established, freely available, and widely taught in schools and universities. It is management science, and management science has powerful techniques to help us wherever limitations of our knowledge (and consequent uncertainty) are important to us.

Many people already know at least some of these techniques from school, university, or professional studies. We recognize their characteristic use of clear diagrams, explicit mental models, mathematics, and calculations. Most children in the UK start learning about the mathematics of probability before they are 10 years old.

This introduction to the management science most relevant to 'managing risk' pulls together relevant techniques and resources, some of which you may not have come across before, and organizes them according to where in management they are useful for dealing with uncertainty.

The areas of management thinking covered build on each other, with new aspects being introduced along the way. The areas are:

- Understanding what is going on now, and why
- Making predictions about what might happen
- Making decisions
- Developing designs and plans
- Evaluating progress
- Communicating

¹ Management science is not a defined body of knowledge, but here I just mean the sort of things covered in management science textbooks, usually just because they are topics in management that have been studied in a scientific way. Management scientists disagree with each other on many things, but tend to agree that clear thinking, objectivity, mathematics, models, and research are good things. Probability is the leading approach to uncertainty within this tradition.

The purpose is to help people identify opportunities to improve the way they manage, and find resources that give more understanding and guidance on the specific techniques they want to use.

In each area I point out where uncertainty is important², offer principles for a management science approach, illustrate how they apply in everyday conversations, then give brief overviews of some more advanced techniques that might be useful in some situations. At the end of the document there are suggestions for further reading about the techniques mentioned.

Understanding what is going on

A lot of attention in 'risk management' is devoted to thinking about what might happen in future, but before we can do that we almost always need to understand what is happening now, what has happened in the past, and why. Making sense of our situation is a crucial starting point and often our knowledge is surprisingly limited.

Here are some examples of situations where uncertainty about what is happening can be crucial:

- A company boss asks "Why have sales of our new product been so high? How much does it really cost to produce it?"
- A government minister asks "Why is my government in so much debt? Why are factories closing in my country?"
- A parent asks "Where is my 14 year old daughter, who is she with, and what is she doing?"
- A doctor asks "What is causing my patient's pain?"

In these examples we are familiar with the actions people take to find out more, such as gathering and analysing data about a company or economy, calling the daughter, and performing medical tests to reach a diagnosis. We also know that these often do not eliminate uncertainty. It can be extraordinarily difficult to learn anything useful about causality from sales reports. Economic statistics, while available in great numbers, still give only a rough idea of what might be happening to the millions of very different people in a country. Meanwhile, the daughter may give false assurances and medical tests are often unreliable even in the relatively simple case where there is only one disease to be diagnosed.

The underlying principles to apply to deal with this sort of uncertainty include these:

- Keep a very open mind about what the truth might be and what data might turn out to be most relevant.
- Pay attention to the data you have already.
- Get good data efficiently and continue trying to find out which data matter most.
- Make inferences from your data that point towards the true explanation and away from false explanations.

² A survey carried out in April 2012 showed that most people see value in being able to manage uncertainty around several aspects of decision making, and these go beyond the usual focus on uncertainty around predictions of the future. See 'Results of a survey on the locations of uncertainty' on the Working In Uncertainty website, available at http://www.workinginuncertainty.co.uk/study_uncloc_report.shtml.

- Try not to choose an explanation unless you are certain it is *the* explanation. Instead, hold on to multiple possibilities but work with them efficiently.

If you are a manager and want to encourage people to apply these principles effectively then you can do so by the things you say in conversations. For example:

- To encourage open mindedness: "What else could be going on here?", "I don't think we can be entirely sure that's true.", and "How do we know they want to do this at all?"
- To use data you have: "Look at this. The sales of this product started to increase two months before our advertising campaign and fell again three months after the advertising stopped. The greatest increases were in Europe. What do those things suggest about who is buying the product and why?"
- To uncover more relevant data: "I know we spent a lot on that advertising but I'm not sure it is relevant here. Can we get some more information about why major customers made those purchases then?"
- To get more data: "Could you please call Tom at Big Co and ask why he hasn't bought anything from us this month?"
- To make inferences: "Our biggest customer said he hasn't bought from us this month because he built up stocks ahead of the tax change. The timing of the tax change corresponds to the rise in our sales, but the advertising campaign was too late. It looks like the tax change is the main reason our sales rose then fell, and our advertising had little or no effect."
- To avoid premature conclusions: "OK, so the evidence suggests that the tax change was probably more important than our advertising for our biggest customer and perhaps others as well, but don't forget that the advertising might still have had an effect and there could be other explanations too."

If numbers are necessarily involved, or if you are comfortable with numbers and calculations and want to use their power and convenience, then there are some well-established techniques that are worth knowing. Understanding what is going on often involves making measurements and trying to make sense of them, so this list starts with measurement uncertainty.

Assessing measurement uncertainty: If you weigh yourself with bathroom scales ten times in rapid succession you may be surprised to see that your scales do not give the same weight each time. Clearly, none of the measurements is entirely reliable. This phenomenon has been a problem for scientists for hundreds of years and some of the most important ideas about how to deal with it were worked out a long time ago by people forced to use unreliable astronomical measurements. There are many specific techniques, but by looking at the distribution of multiple measurements it is possible to make a better best estimate of the true value and to say how likely it is that the true value lies within a range.

But measurement uncertainty isn't just the result of unreliable instruments. In business, many of the numbers given in management reports are wrong or just estimates. Some, such as customer satisfaction numbers, rely on taking samples. Measurement uncertainty is usually greater than most people realise.

Quantifying rounding errors: It's not usually important, but occasionally rounding errors matter, especially if you do calculations with a number that could magnify the size of the errors. There are simple mathematical formulae that make it possible to calculate how errors will propagate through a calculation.

Information graphics: The right graphs for your data can help with understanding what is going on, including getting a feel for past variability. A 5%

drop in revenue is alarming if revenue rarely fluctuates by more than 1%, but is boring if it usually rises and falls like a rollercoaster.

Information graphics can also be used to show measurement uncertainty.

Characterizing patterns in data: A surprisingly large proportion of scientific research and theorising does not provide explanations for observations. Instead, it just characterises those observations, revealing and describing regularities in them rather than describing how a familiar underlying mechanism is producing them. The first step is usually just to plot graphs. The next is to fit simple shapes to those graphs. Tukey's Exploratory Data Analysis method is quite explicit in selecting simple line shapes that appear to fit the data.

This kind of characterization, despite its lack of explanation, is valuable. Sometimes we can just extrapolate from the past to predict the future. Also, having a good summary of past observations makes it easier to compare them with explanations we might come up with in future.

While superficially similar, the equations for these simple shapes are not the same, in principle, as equations derived from an explanation. For example, early 'laws' about gasses were just mathematical summaries of how pressure, temperature, and volume were related. In contrast, models derived from statistical mechanics started with the idea that gasses were tiny particles moving around at random and then developed a mathematical model of how such particles would behave that then proved to fit actual gas behaviour, as summarised in the original laws.

Using correlation to detect causality: One typical goal of business reporting is to focus on the 'key' performance indicators. But how do we know what is 'key'? Often we don't. Some indicators turn out to be completely irrelevant, while others may be just slightly misleading. Perhaps 'customer satisfaction' is believed to be 'key' or 'staff turnover', but more relevant indicators turn out to be 'willingness to recommend' and 'staff turnover among supervisors'.

The strongest, most direct way to discover what really matters would be to do experiments. This involves deliberately creating differences in indicators (e.g. treating some customers well and others badly for three months) and measuring what happens to see if there are differences.

However, experiments can be difficult to organize. (For example, who would agree to treat customers badly for three months, or pay staff more for a period just to push staff turnover down for an experiment?) An efficient alternative – or at least a first step – is to filter available data from ordinary operations for interesting correlations.

Obviously this method has its limitations. As everyone should know, finding that A correlates with B does not prove that A *causes* B. For example, if a study finds that companies with a Chief Risk Officer (CRO) tend to provide better returns than others then that does not prove that CROs are worth their pay. It could be that only companies in a commanding position in growing markets can afford a CRO, or that the region of the world that companies operate in drives both their hiring and returns.

However, although correlation does not identify causation, it is a strong clue that causation is at work *somewhere*. Since experimental manipulations are often difficult in business, it can be efficient to filter available data for interesting correlations then investigate further using experiments or by studying the details more closely.

Direct observation of causality: One of the most difficult problems is to try to understand causal mechanisms from data alone. It can be done, eventually, with

enough data and skill. However, an easier way to understand causal mechanisms may be available: direct observation.

For example, imagine you are a farmer who keeps free-range chickens. From time to time a chicken has gone missing but nothing in the pattern of losses tells you why, though you can see that losses are getting more frequent. Suspecting a fox, you decide to observe your chickens. You install infrared CCTV and spend three sleepless nights keeping watch over your chickens until you spot a dark figure enter the coup and grab one of your chickens. Not a fox, but a man in a hoody and trainers. In the case of that particular chicken, causality is not in question. You saw the man enter, grab a chicken, and leave with it. The only question is whether this chicken thief is the sole reason for your losses, or whether there are other reasons, such as a fox, or chickens escaping somehow.

In a typical office work situation a correlation or trend might be easier to understand if you ask people for details of some of the sales, purchases, and other specific acts that contributed to the numbers. Get them to recall details of what they saw happening.

Choosing between explanations: Without direct observation, finding explanations can be tricky. In choosing between alternative interpretations of a situation, or of events, it helps to keep an open mind and look for information that sidelines the explanations that are not true while pointing towards the one that is. The Bayesian approach works just like this and Thomas Bayes's formula provides a mathematical rule for revising our opinions in response to each new piece of evidence. It even works when evidence is suggestive rather than conclusive.

Bayes's formula makes it easier to combine evidence, which is hard to do by judgement alone in many situations. If evidence says that one explanation is less likely then that should boost the others, but with unaided judgement we quickly lose track of the effects of each piece of evidence.

The Bayesian approach uses the fact that it is often easier to think about how likely it is that each piece of evidence might be true if you assume that an explanation is true. This is the opposite of the obvious direction, but Bayes's formula is the mathematical rule that reverses the logic and combines all the information for you.

The starting point is to set out a comprehensive set of possible explanations (though they don't have to be detailed) that are mutually exclusive. In other words, no two explanations can be true at the same time. This is always possible, though your explanations might not be very useful or relevant if you don't know much about what is going on, so you might revise them later.

(Doing this is a helpful practice even if you do not then use mathematics to combine evidence.)

The next step is to decide how likely you think it is that each explanation is the true one without using any new evidence. Again, this is always possible. If you think you have no idea at all then there are techniques for choosing a distribution of probabilities that make the weakest assumptions possible. In practice, as long as your initial views do not give a probability of 1 or 0 to any of the explanations then data will quickly change those views. (If you assign 1 or 0 then your views will never be changed by evidence.)

Now you are ready to consider new evidence. For each piece of evidence (e.g. another opinion, results from testing a sample) consider how likely it is that the evidence would have been seen given each explanation in turn, then apply Bayes's formula to see how likely you now think it is that each explanation is the true one.

Bayes's formula, used this way, is a way of reasoning logically with uncertainty. It is not a statement of facts about the world, such as the frequency of particular events.

One of the great advantages of the Bayesian approach is that it does not require us to select one best explanation (even though it helps us do that if necessary). The Bayesian approach lets us state a probability distribution over all the possible explanations and that distribution can be carried forward into other thinking, such as planning, predictions, and decision-making.

For example, imagine you are working in counter-terrorism and suspect that a terrorist cell exists in a major city. Perhaps you have several alternative hypotheses about how large, skilled, and connected the cell is. Depending on the type of cell you would make different predictions about its likely behaviour. Simply assume each hypothesis is true, in turn, and write out your predictions in the form of possible actions and their probability. Then account for the probability of each hypothesis being true by multiplying the probability of each action, given that the hypothesis is true, by the probability that the hypothesis is true. Finally, sum the probabilities across each action to get your probabilities for each action being taken by the cell.

Making predictions

Predictions about what might happen in future are important for recognizing when it is time to think again, and for choosing between alternative courses of action. You might forecast what will happen in the world around you, regardless of what you do. You might forecast what will happen if you continue as you are now. If these predictions worry you then it may be time to think again and do something different. With each new plan that you consider you will want to think about what might happen if you used it. Our predictions range from quick mental simulations to elaborate economic forecasts made using surveys and computerised models. Predictions are important.

Obviously, some predictions are easier to make accurately than others and most people understand that predictions are usually uncertain. However, we tend to think we are better at making predictions than we really are. There seem to be many reasons for this. Tests by psychologists have established (1) our tendency towards thinking things will turn out better than, on average, is the case and (2) our tendency to be too confident that our predictions are accurate.

The main thing to remember about predictions is not to accept or use a single-point, best-guess forecast of any kind, ever. Always think about the various different things that could happen and how likely they are (but do not jump into unhelpful detail).

In conversation, you might say things like:

- "What else could happen?"
- "Let's consider the broad alternative outcomes before we consider details."
- "What if this apparent trend is really just a bubble?"
- "What could happen that's outside our control but would be important to us?"
- "Let's think about what it would mean for us if things turned out according each of the scenarios we have thought of."

Some useful techniques for when the stakes are higher include the following:

Multiple scenarios: One of the simplest ways to respond to uncertainty in predictions is to consider the implications of a number of different possibilities. For example, an accountant with a cash forecasting spreadsheet that produces best

guess forecasts can easily recalculate for different interest rates, sales growth assumptions, or rent increases. Even without calculations, it can be helpful to imagine the impact of major events. For example, what might happen if you lost your current job, or got a promotion, or a promotion that involved moving house? What might happen if interest rates rose by 2% and then you lost your job? Answering these questions does not overcome the fact that you don't know what will happen in future, but it does help you think through the implications and build more flexibility into your plans.

Classic management science techniques typically involve choosing between plans by looking at a table of outcomes in different scenarios and making choices using different rules. Often, a probability is assigned to each scenario.

Bayesian Model Averaging: People sometimes think that they have to choose one prediction model from all the possibilities. For example, they may have a lot of data about past returns from projects and want to use it to predict future returns on a new project. An infinite number of mathematical models could fit those data, but some are more likely to be true than others. One approach to this is to pick just one model, the one that seems to fit best. However, it's often a close thing, and setting aside all the other possibilities means ignoring some of your uncertainty – which is usually a bad thing to do.

Bayesian Model Averaging is the idea of combining the predictions from all the models, weighting each according to the probability that its model is true. The combination might be an average or, for something like the chance of a bridge collapsing, the sum of the probabilities from each model.

Assessing and improving forecasting skill: In this context the word 'skill' has a very specific and fascinating meaning. Research inspired by weather forecasting has led to ways to measure the 'skill' of a forecaster (or forecasting method) given knowledge of their past probabilistic forecasts and what actually happened.

A perfect forecaster would always assign a probability of 1 (i.e. certainty) to the outcome that then actually happens, and a probability of 0 to all other outcomes. Few forecasters are as good as this so their skill scores will be less than the perfect score.

One skill score is called Ignorance (IGN for short) and it shows how much less information a forecaster provides than a perfect forecaster. In real business situations, when people give probability numbers based on their subjective certainty these can be surprisingly good numbers compared to probabilities derived from statistics alone. We tend to take many factors into consideration and this compensates for our unfortunate tendency to be too confident.

Furthermore, with practice and feedback we can improve an element of our forecasting skill known as calibration, and this reduces our Ignorance. Calibration is a measure of how well our probabilities match up to experience. For example, if someone gives a probability of 0.8 for many similar forecasts and in fact the event does happen 80% of the time on average, then the person is well calibrated. Note carefully that calibration is only one part of forecasting skill and it is possible to produce perfectly calibrated probabilities that are next to useless.

Using measures of skill, such as Ignorance, it is possible to assess alternative sources of forecasts and even begin to track long-term improvements in forecasting skill.

Probability elicitation methods: When you ask someone to give a probability for something based on their expertise and judgement they will tend to give a

number that is biased in various ways. However, these biases are at least partly understood thanks to many experiments by scientists and a number of good methods have been developed for asking for probabilities that tend to counter the biases. These usually include some preliminary conversation to alert the expert to the types of bias that may occur, ways to word the questions you ask, and ways to use the answers from more than one person. (Remember that even biased probabilities can be more informative, and hence more useful, than historical statistics.)

Forecast markets: One relatively new idea for getting probabilistic forecasts from lots of people is to set up a forecast market. This usually involves a computerised market and giving people some money or tokens with some value that they can use to bet on outcomes. Most of these focus on getting a lot of people to weigh in on a small number of predictions, but designs also exist to coax many predictions out of a smaller number of participants.

Forecast markets do not work very well if participants simply have no relevant information or everyone is relying on exactly the same information (e.g. a news broadcast). In contrast, for predicting whether a movie will be a hit or not most people can use their taste and understanding of the tastes of others to take a punt.

Prediction intervals: Prediction intervals are relevant when you have to predict a number e.g. sales next month or the cost of a project. A prediction interval is a range between two numbers such that you have a specific level of confidence that the truth will lie in the range. For example, you might say you are 80% confident (i.e. probability is 0.8) that the cost of a project will be between £3m and £3.5m. People tend to take prediction intervals quite seriously because they sound scientific, so it is important to avoid giving a range that is too narrow for the stated confidence.

Empirical prediction intervals: If you have past experience as a guide it is possible to use it to calculate prediction intervals. For example, if you have given best estimates in the past (perhaps along with prediction intervals) then the distribution of the differences between your best estimates and the truth can be used to calculate prediction intervals for new forecasts.

A trap for the unwary is that sometimes prediction models are mistakenly tested using the same data that were used for building the model in the first place. This makes the model seem better at prediction than it really is. Test it against different data and you usually find that it does not perform as well. There are well established methods for doing this.

Prediction intervals from propagating uncertainty: If you don't have past experience as a guide then you can either set a prediction interval by sheer judgement or use a model to help. Forecasting models typically do calculations with various inputs, many of which are themselves predictions, and produce an output, such as a forecast cost for a project. If you find the model helps you make a prediction then that is usually because it is easier to think about each of the inputs and about the logic of the model than to do all that by a massive leap of judgement and just go straight to the final prediction. This is typical for situations where we lack directly comparable past experience of the number to be predicted.

If that is the situation, then it is also probably easier for you to express your uncertainty about the inputs and the model, then use software to calculate automatically what that implies for your uncertainty about the output to be predicted.

The best known technique for doing this calculation is Monte Carlo simulation, which is a simple idea made possible by the amazing number crunching power of modern computers. It involves a blizzard of random numbers but since software does

it for you it does not involve a lot of effort. With very simple models you can do it using a table in a spreadsheet and for more demanding situations you can use Excel add-ins ranging from good free macros to amazing packages costing several hundred pounds, or even more.

Making decisions

We make a staggering number of decisions in our lives. These include decisions forced on us, such as whether to accept a voluntary redundancy offer, and decisions we create because we have a new idea, such as whether to buy an attractive holiday or change a business policy. Most things about organizations are the result of one or more decisions taken at some point, ranging from where the organization is located to what colours appear in its logo. Many of those decisions are taken in bundles, which is a situation covered in a later section on design and planning.

Knowing what is going on and being able to predict the future are helpful when it comes to making decisions, but making decisions adds more opportunities to be uncertain.

- We are often uncertain about what would be in our best interests. Put another way, we are often uncertain about what our objectives should be. No amount of navel gazing can overcome this because our uncertainty stems from having limited knowledge about how things we can plan to achieve would translate, ultimately, into a better life.
- If the alternatives we must choose from are givens then there may be details of them that are uncertain, and it is possible that more alternatives will be arise if we wait a bit longer.
- If we have to invent the alternatives then there is uncertainty about what alternatives to consider initially, what each alternative might involve and, often, great uncertainty about when to stop looking for better alternatives and just make a choice. (In the extreme this becomes a design task, see Developing Designs and Plans, which is the next section.)
- There is usually uncertainty about when to stop looking for more *information* and just make a choice, particularly if we are also inventing new alternatives.

To deal with risk/uncertainty well when making decisions:

- Keep an open mind throughout.
- Recognize all the areas of uncertainty, including uncertainty about what would be in the best interests of those involved.
- Expect to think iteratively, learning and modifying your ideas as you explore the alternatives and what they might lead to.
- Take the time and care that each decision deserves and allows, getting more information and using calculations where you can.

In particular, although it is a good idea to think very carefully about what would be good to achieve, there is usually no need to set specific target levels for performance³ and it is a big mistake to insist on specific targets before doing anything else.

³ Control theory has explored the idea of negative feedback control, in which a control system uses deviation from a 'target' to drive actions designed to reduce the deviation. This applies well to electronic circuits, steam engines, and other systems where there are known actions that can reduce deviations. However, in a general management situation there often is no action that can be taken to get 'back on track', let alone one that is known.

In conversation, you can promote good decision making under uncertainty by saying things like this:

- To explore uncertainty around our interests: "What do we know about what would be good outcomes for us?", "Who could be affected by this and what might they be hoping for or worried about?", "How could we measure the degree of our success?", "Just how valuable would a 10% improvement be to us? What about a 20% improvement? Are we guessing about this?", and "How much extra should we be willing to pay, as a maximum, in return for the better ergonomics of the second product?"
- To get detail of alternatives: "When you say you are offering us a choice between two contracts, can you give more detail about exactly what the terms would be?"
- To open up new alternatives: "So far we've considered competing directly with their product, or just withdrawing from that market. What else could we do?" and "I'm not convinced we've considered all the options. We've got time to explore alternative ways to make the purchase so let's look into them a bit further and see what we could do."
- To gain more information: "One of our difficulties with this decision is that we don't know much about what it will cost to buy the packaging. Let's do a bit of research to find out more."
- To test if it is time to make a final choice: "Are there any further alternatives that justify more exploration, or have we reached a point where we're unlikely to have any better ideas?" and "Is there anything more we can usefully find out before we make our choice?"

This is not rocket science; it's just common sense. However, it is also the basis for a lot of the decision theory at the core of management science. Here are some relevant techniques. Although they are usually used only when the stakes are high, it is useful to understand the principles underlying them because these principles are valid and useful even when the techniques are not used.

Objective functions: Thinking carefully about what is in our interests is often a good thing to do and one of the best known ideas for representing this in management science is the 'objective function'. The word 'objective' means different things to different people, but 'objective function' means something specific. It is a mathematical formula designed to translate different futures into single numerical values. Typically, it maps outcomes to how much you value them. For example, the objective function might be total profit from a venture over its first year, or the number of votes gained in an election. Usually, the idea is to maximise, or minimise, that objective function, subject to meeting some constraints. For example, if you were trying to pack a suitcase you might have an objective function reflecting the value of the stuff you have packed, with the constraint being that you have only the capacity of the suitcase.

Often, we have multiple considerations, so one thing the objective function has to do is combine those into one value. Even crude formulae can do better on average than unaided human judgement and the sobering conclusion from a lot of research is that we're just not good judges when there are lots of considerations.

However, the robustness of simple formulae should not be taken as evidence that they are free from obvious and important flaws. While multiplying each consideration by a weighting factor and adding up the results is a common technique and works better than unaided judgement, usually, its flaws are obvious in situations where we have time to go into detail. The relationship between scores on a

consideration and value is rarely linear. For example, in choosing a car, top speed may be a consideration but there is little to be gained from being able to drive at three times the legal speed limit rather than just twice the legal limit.

Ralph Keeney has used the term 'value model' for objective functions that reflect our valuation of alternative outcomes. This value is, in effect, a prediction about what each outcome means to our lives. For example, the value model for an aircraft design might be a function that takes basic performance measures for an aeroplane and forecasts the commercial value of the machine over its lifetime. This will reflect factors like how many seats it has and how efficiently it uses fuel. Obviously, this forecast is uncertain and the uncertainty can be represented, as usual, with probability distributions.

Utility curve: This is a specific type of objective function intended to represent the value a person puts on different quantities of something desirable, such as money, time, or food. The value is expressed in 'utility', which is an imaginary scale of value. Utility does not have defined units, so if you want to construct a utility curve you have to use techniques that, in effect, establish units.

Utility curves don't just reflect personality. They reflect circumstances too, and can shift from moment to moment as circumstances shift. Typically, the more of something desirable that we have the better, but the extra utility of each extra unit of the item tends to fall as we get more and more. A £10,000 lottery win usually means more to a poor person than to someone who already has many millions. However, if that multi-millionaire happened to need £10,000 in cash right now to avert some disaster, and anything less would not be enough, then the relevant utility curve would look very different.

Conjoint analysis: This is a technique for eliciting an objective function/value model from a person, and it can deal with non-linear relationships between performance and value, and with multiple criteria. The person whose views you want to understand simply makes a series of choices between realistic alternatives and a computer program then works out an objective function that captures their views as best it can. The more choices they make the more accurately their views are captured. It's so easy and nuanced that if you use conjoint analysis to study your own views you will learn something about yourself. If you suspect that people you work with have different priorities you could use it to find out *exactly* how they differ.

Once you've experienced conjoint analysis you will be much less likely to ask feeble questions like "How would you rank the following considerations for importance: price, performance, looks, brand?"

Willingness To Pay (WTP): This is yet another idea for finding out what people would value. The technique is simply to ask people what is the most they would be willing to pay for something. WTP has been used many, many times, usually to research how much people value healthcare treatments.

Direct choice between distributions of outcomes: For all the help that conjoint analysis and a well-defined objective function can provide, there are many situations that are difficult to understand in that way. The work involved may be so great that it is easier to just predict the potential results of different decision options and present that to decision makers for them to respond to. In effect, they just need to think about how much they value those particular outcomes, rather than all conceivable outcomes. That's why it is easier for decision makers even though it is less helpful to people trying to develop courses of action. It also misses the opportunity to give decision makers feedback on their values that might help them to think in a more coherent way.

Another reason for just presenting decision makers with probability distributions of forecast outcomes for alternatives under consideration is that the shape of the distributions is itself important. The value of a distribution of value is not necessarily the same as the probability weighted average of that distribution for a variety of simple, practical reasons. If we can act on foreknowledge then certainty makes a difference in itself.

Mean-variance approach: This is an approach to comparing probability distributions over one criterion in a decision that involves calculating the mean and variance of each distribution and then comparing these statistics. Unfortunately, although considering mean and variance can allow you to eliminate some alternatives from consideration it does not give you enough information to make all decisions. To do that you have to get people to identify combinations of mean and variance numbers that they find equally attractive. (How people are supposed to have a view about exactly how much they care about different variance numbers I don't know.)

The thinking behind the mean variance approach is that if you have a typical flattening utility curve for a decision and the distributions are all symmetrical then the more spread a probability distribution the worse it is, other things being equal. In addition, the less predictable the results the harder it is to make worthwhile further plans. (In some real situations greater spread is desirable so be careful about using this approach without thinking.)

Decision trees: Decision trees are diagrams that show your alternatives and also different things that might happen, with probabilities assigned to the different outcomes. There are methods for working out more than one decision within a single tree. The main limitation of decision trees is that the number of alternatives has to be reasonably small. If you are trying to find, say, the best retail price for a product, then the alternatives might be small in number, e.g. £2.49, £2.59, or £2.69, due to the psychological preference for certain prices and a narrow range of competitive prices on a low margin product. A decision tree might be useful here. In contrast, if you were trying to choose the price per kilo for some industrial commodity then your price might easily be specified to fractions of a penny, giving hundreds of possibilities even within a small range of prices. To deal with that sort of problem you need a different approach.

Optimisation: Where there are very many alternatives it is better to avoid decision trees and instead set up a model that allows the best alternative to be found by algebra or by numerical methods of searching. There are lots of alternative methods and tools, but a familiar tool is the Solver feature of Excel, which uses numerical methods to search for the best answer.

Expected values: A common problem in decision making under uncertainty is to have some outcome that could be anywhere in a range of values. The usual way to express a view about the likely value is with a probability distribution. Now suppose you have two alternative courses of action, each with a predicted probability distribution. Which distribution do you prefer? In some cases the difference may be so much that it is obvious which is better. In other cases it depends on how you translate different values of the outcome into utility. This is a surprisingly deep question.

One way to deal with this is to just calculate the expected value of the outcome, which is its probability-weighted average. It's crude but still useful provided the outcomes in question are not extreme and provided you are nowhere near running out of resources.

For example, if you are a wealthy person making a short series of small bets then using expected values as the basis for your choices is a good approach. However, if you are a poor person making a series of rather large bets then the fact that money and utility are not the same becomes really important and losses will weigh much more heavily than gains of the same size. Also, there is a real possibility that you will run out of money and not be able to participate in all the bets.

Proportional betting: A strategy to consider if you could run out of money in a series of bets is to bet just a proportion of your remaining funds. That way you never quite run out. This means that the expected value of each individual bet is not maximised. Kelly betting is a strategy based on using a fixed proportion of your funds in every bet and the Kelly Criterion tells you what that proportion should be (though only if certain assumptions are true). In practice, people tend to bet less than the Kelly Criterion advises to reduce the size of wins and losses, and there is no need to keep the proportion the same at all times, as the Kelly Criterion does.

Discounting rates: One well known approach to valuing alternatives is applicable if all you care about is money and if you can estimate something called the beta of the alternative. The theory goes like this. If you are just interested in the money then predict the cash flows for each alternative and evaluate those in comparison to the cash you would get by just lending your money to a typical bank and receiving interest. The way to do that calculation is by 'discounting' future cash flows by a rate that is the interest rate you would get in a typical bank. If the net value of the cash flows discounted in this way (known as the Net Present Value) is positive then the alternative is better than just putting money in the bank. The higher the Net Present Value the better the alternative from this point of view.

In these cash flow models the cash flows are all best estimates but 'risk' in some sense is brought into consideration by varying the discount rate. The theory is that a bank or other investor needs to receive a higher rate of interest on investments that are 'risky' in the sense of having less predictable returns. Therefore, to be fair you should discount the alternative's cash flows using a rate that reflects the riskiness of the alternative. The higher the riskiness the higher the rate should be and so the less likely it is that the alternative will pay better than investing at the bank.

A further refinement in this theory is that the riskiness is not just any riskiness, but just the riskiness that an investor cannot avoid by diversification, known as beta. In practice that means that the only riskiness of interest here is the riskiness that is common to all alternatives of a broadly similar type. This kind of riskiness can be estimated for listed companies as a whole by doing statistics on their past share prices but is not an easy thing to assess in most decisions.

Iteration: Decisions involving uncertainty can be tricky. Each attempt to predict the outcomes that might arise from choosing each of the decision alternatives provides feedback about the prediction approach and the alternatives. The predictions and alternatives might be revised and perhaps improved as a result.

Iteration is a strategy to be pursued, not a burden to resent. Do not expect to go through a single, linear process of deliberation and then make a good decision in the face of uncertainty. Get set up to iterate efficiently.

Automated calculations: The mental work involved in predicting the outcomes if each decision alternative were to be chosen, and comparing those outcomes, can be tough and tedious. Searching your feelings to reach a gut feeling about which is best can be agonising, exhausting, and just plain boring. (Imagine making 20,000 credit decisions by judgement!) Uncertainty makes this more challenging still

because it means considering a wider range of alternative outcomes and because it tends to involve repeated attempts at making the decision, as ideas are improved.

To cut down the effort involved it is often worth investing a bit of time and thought in a numerical model that can be automated (usually a spreadsheet of course). An accountant today can recalculate a cash flow forecast with a click. Thirty years ago it might have taken hours. That's the power of automation.

A further bonus of automated calculations is that very often they are more accurate and reliable than unaided judgement. Judgement is particularly poor at combining multiple considerations and at tasks that require accurate quantification.

Also, the mental processes needed to set out the calculations clearly are valuable in themselves, promoting logical, coherent thought.

Value of Information: Is it worth doing some more research? There's a way to work it out. There are well known calculations for the value of perfect information (perfect information about something in your decision, not everything), and the value of imperfect information. These involve working out what your choice would be given current information and comparing it with the probability weighted average of your choices under alternative possible futures, assuming you knew them in advance (or at least had better information about them).

Automated calculation, of course, makes this easier.

The alternative of thinking or waiting some more: Thinking about the value of information is one aspect of a more general approach, which is to remember that one alternative in the decision that is almost always available is to do some more work on the decision before making a choice. That might be more work on the decision model, more work developing a better course of action, more work finding more information, or might just involve waiting for better opportunities to come along. There's nothing to stop you listing that alternative explicitly and evaluating its prospects along with other alternatives.

There are many everyday situations at work where waiting might be the best choice for now, such as evaluating project proposals and job candidates. Theoretical analysis of these problems shows that, typically, the best strategy is to build up an understanding of what is available to you by evaluating and rejecting a number of possibilities at first, then pick the first new alternative that is better than anything you have seen so far.

Developing designs and plans

Design is not just about choosing a nice logo or an attractive package for a product. It refers to a complex thought process by which many connected decisions are taken. We engage in design when we:

- develop a set of related policies;
- design an object or system, such as a bridge, house, ship, computer system, or work procedure; or
- plan a project.

Planning is a form of design, one in which we design a complex bundle of actions. Sometimes the whole job might be to design an object and the plan by which it will be constructed.

Decision-making (including understanding the situation and making predictions) is at the heart of developing designs and making plans, but uncertainty becomes important in new ways in these more elaborate tasks.

It's easier to understand why if you know something about the fundamental nature of design. What makes it special is that there are, usually, too many potentially alternative designs to count or even identify in detail, let alone individually evaluate in full.

For example, suppose the task is to design a road bridge across a river. There are unimaginably many alternative bridges that might be designed, with features ranging from the width of the central span (if there is one) to the exact routing of wires used to supply electricity to whatever arrangement of lighting is used.

As Herbert Simon described it⁴, design is a selective search through a large problem space (i.e. set of alternative solutions), rather than a matter of evaluating all alternatives and choosing the best. That search cannot consider every possible solution. It is a selective search capable of finding good solutions, but not of finding the best solution and knowing it to be the best solution. We cannot point to all the alternative designs that were set aside in favour of the final design proposed because they are innumerable and we did not describe them in full. Very often in practice we cannot even state a rule for generating the entire problem space because so much remains uncertain and undiscovered.

What we do when we design is to take a series of tentative decisions about which subsets of solutions to explore further. For example, in designing that road bridge the first decision might be between some kind of suspension bridge and a more traditional design. If a decision is taken to explore suspension bridge designs first then perhaps the next decision might be about how many pillars there will be, or perhaps where pillars can be placed. This is a difficult one because geological considerations may limit the places where a pillar can be located, or at least may make some locations easier and cheaper to tackle. But this decision could interact with decisions about how to route the road that leads up to and away from the bridge. Those routing decisions could be influenced by the length of road resulting, the amount of land that needs to be acquired, the opportunities for siting toll booths, and a host of other considerations. If pillars are tentatively located first those locations may well have to be revised once road routing is considered.

Some decisions are better ones to start with than others. Some decisions are relatively independent of all others, so work can proceed on them in parallel with other design work. A good decision to start with is often one whose best alternative is very clearly the best, making revision of the decision very unlikely. A skilled designer can make deductions about the solution that limit the search for solutions.

With all this in mind, the extra complexities related to uncertainty that design and planning introduce include:

- uncertainties about how to frame design decisions;
- about which decision to tackle next;
- about which decision alternative to explore first in each decision; and
- about what performance is likely to be achieved from each decision alternative given that many decisions have yet to be made and may be clouded in yet more uncertainty.

The uncertainty about which sequence of decisions to follow is driven in part by the difficulty of making predictions about future performance of the design/plan, which is harder during design/planning because often so much is still undecided.

⁴ For example, in Simon, H. A., 1996. *The Sciences of the Artificial*, 3rd edition. The MIT Press, Cambridge, Massachusetts.

Although difficult, these predictions are valuable. For example, Barnes Wallis⁵ designed some of the most effective bombs of World War II and his starting point was to make calculations about the force necessary to destroy large industrial targets, from which he deduced the size of bomb required, the size of aircraft needed, and so on. Doubtless his calculations were not particularly accurate, but they helped him realize that the design task was far more challenging than he might have thought. He realized that, even with the best explosive available at the time, no aircraft was available that could carry the big bomb needed. (When the Lancaster bomber arrived his bombs became feasible.)

Despite taking the trouble to make some powerful predictions, Barnes Wallis still relied heavily on testing his designs in practical ways. In developing his famous bouncing bomb he tested explosive charges on model dams and he tested bouncing balls off water in his back garden, in a ship model tank at the National Physical Laboratory, and by dropping real bombs from an aeroplane over the sea and at a disused dam in Wales.

Here are some principles for dealing with uncertainty during design/planning:

- Keep an open mind about what decisions to take, the order of them, what alternatives to consider, and what can be achieved by each route.
- Look at the design/planning problem from every angle, deducing things about the solution from observations you can make. Narrow down the search (which can lead to the most innovative solutions).
- Go down the most promising routes first but be prepared to backtrack.
- Apply relevant design skill; there is no substitute for years spent developing a powerful design/planning ability.
- Look for ways to learn from experience, for example by trialling ideas with a prototype or simulation.

My personal experience is that most people put too little effort into design thinking. They are more comfortable thinking of reasons why something cannot be done or why an idea will not work. When they do have a positive suggestion it is often obvious, vague, and not based on research or reasoning. They prefer not to explore ideas in depth or refine details.

If you want to get more from people then in conversation you might say things like this:

- To encourage effort: "Would anyone like to take on the task of researching this thoroughly and coming up with some properly thought through design proposals for this music festival?", "Let's take that idea and think about how it would work in practice."
- To encourage inferences about good solutions: "Let's begin by looking at this from every angle to see what we can deduce about how to design and plan our music festival.", "Our music festival is to be located on a rural island just off the coast of England. What does that imply about our plan for the festival?", "This is the first time there has been a music festival on the island. What does that imply for what we have to do to make it a success?"
- To encourage learning from experience: "Hopefully this is just the first of what will become an established annual festival. What can we do to learn more from the first festival and improve our next design?", "What are the ideas we've chosen this first time that we feel are a bit of a gamble?"

⁵ Morpurgo, J. E. (1972) *Barnes Wallis, A Biography*. Longman, London.

Management science has surprisingly little to say about the special wrinkles of design and planning that go beyond decision making. However, by looking closely at examples of it in action one can see some important ideas:

Model refinement: Models are very commonly used in decision-making and where uncertainty is represented explicitly in models it can be propagated from estimated inputs to the predicted performance of a design or plan. What happens in design is simply that the model(s) get revised and refined as work progresses, gradually incorporating more and more of the decisions that have been tentatively taken.

For example, in project planning a first cut plan (modelled on a computerised scheduling application) may have only rather high level tasks in it, with little work done on interdependencies and shared resources. Later versions will have more detail, especially for work that is to be done early in the project. As the project progresses, more detail goes into the later parts of the plan.

Structural heuristics: Some structures (e.g. physical structures, plan structures) have inherently superior characteristics in the face of uncertainty. For example, nuclear power station control systems are designed with *redundancy* in mind. This means that if one button stops working there will be another that does the same job that can be used instead. In planning, incremental delivery is generally superior if it can be done and the advantage increases with greater uncertainty.

Evaluating progress

We evaluate progress for a number of reasons. We do it to:

- learn and revise our expectations for the future;
- decide if it is time to think again and perhaps devise a revised approach; and to
- assess the effort and ability others are giving, with implications for job promotion and performance pay.

If we are evaluating progress to revise our expectations for the future then everything we know about uncertainty in establishing the current situation and making predictions is relevant.

If we are trying to decide if it is time to think again then what we know about establishing the current situation and about making decisions is relevant. In addition, we need to learn what usually justifies a rethink. That will include discovering that the situation is not what we thought, or that expected results are now significantly different from our previous expectations.

If we are trying to evaluate effort and ability then, again, what we know about establishing the current situation is relevant. In addition, there is a very interesting new uncertainty about what to use for comparison.

Comparisons with initial expectations tend to be useful primarily for learning about our forecasting ability, and for detecting when things have changed. However, they are less useful for identifying whether someone is performing well or poorly. If conditions are different from those initially expected (which they usually are) then a person's results will have been affected by that as well as by their effort and ability. If there are other people doing the same work then it is often better to compare people with each other. Another approach is to compare results with what you would have expected if you had known then what you know now about conditions.

One long established area of management science directly related to evaluating effort and ability is this one:

Agency theory: This is theorising about what happens when one person (the agent) works for another (the principal). This is usually imagined as a board of executive directors working for shareholders, but it could be any principal-agent relationship. One of the questions that has been studied in detail is how the principal can judge the performance of the agent. The usual theoretical assumption is that the principal does not know how the agent has behaved but does know what results were achieved. A problem arises if those results are only loosely related to effort and ability, as is usually the case in many real situations. Part of the solution to that problem is to get information about behaviour and include that as part of the evaluation. If you can see that people have done a lot of good things but still the results seem disappointing then it is more likely that they've just been unlucky and a less competent, lazy person probably would have achieved even less.

Communicating

Communicating with others is important most of the time and there are some obvious and important uncertainties involved:

- If you say something, will the other person understand correctly?
- If someone says something to you, did you understand correctly what they were trying to say?

How much does anyone understand of anything? It's not easy to know but a good rule of thumb is that it will be less than you expect. For example, some interesting studies have been done to find out how well members of juries understand the law when it is explained to them, and how to explain the law better. Depressingly, a lot of jurors don't understand the law correctly when it is explained to them.

Some principles for dealing with uncertainty in communication include these:

- Expect misunderstandings.
- If misunderstandings could be important then use whatever means are available to avoid, detect, and correct them.
- Cut out opportunities for misunderstandings.

Here are some ideas from science that are relevant.

Numbers: Quantitatively vague phrases are a particularly big problem. These include phrases about probability such as 'quite likely', 'probably', and 'unlikely' as in "It's quite likely that we will win." These phrases do not have agreed on, precise meanings (even ranges) and their meaning tends to be relative to expectations rather than the same in all situations.

Using numbers reduces this problem dramatically, even if you use ranges rather than single points. The remaining problem is that people can make mistakes over what exactly a number represents (which they also do with the quantitatively vague phrases).

Information Theory: The amount of information in a message can be quantified and the theory of quantifying information makes use of probabilities. Improbable messages (the surprising ones) convey more information than messages telling you what you expected to hear. Unfortunately if you are not sure which of two messages you have just received the natural tendency is to assume it was the more likely message.

This means that if you want someone to understand something that will surprise them (e.g. a new insight contrary to common belief, evidence that the person was wrong about something) you will need to take special care to communicate clearly.

Finally – keep it simple

A lot has been written about uncertainty in management from a scientific perspective and browsing some of this material it would be easy to get the impression that it must always involve very complicated mathematics requiring enormous skill and effort. That is simply because, to get published these days, you have to do something very clever and impressive.

In reality, most everyday applications of this kind of thinking are simple, requiring little more than a quick calculation or spreadsheet, or perhaps just a diagram for clarity, at the most.

When applying these ideas to your work look for the things people do often and that are important. Seek out the various management tasks within them and pick some opportunities to improve, being very selective. Keep things simple and just make the tweaks that are most likely to be worthwhile.



About the author

Matthew Leitch is an author, educator, consultant, and researcher specialising in the way we deal with uncertainty at work.

About Working In Uncertainty

The idea is that, by making incremental improvements to the way we work that help us deal with uncertainty, we can achieve more with less stress. This applies to individuals and organizations. The starting point is usually to learn more about techniques that work well and try them. For more information please visit www.WorkingInUncertainty.co.uk.

Further reading

Here are some suggestions for further reading.

Technique	Reading suggestions
Assessing measurement uncertainty	<p>Bell, S., 1999. <i>A Beginner's Guide to Uncertainty of Measurement</i>. National Physical Laboratory. Available at: http://www.wmo.int/pages/prog/gcos/documents/gruanmanuals/UK_NPL/mgpg11.pdf</p> <p>NIST page on <i>Uncertainty of measurement results</i> available at: http://physics.nist.gov/cuu/Uncertainty/</p> <p>Securities Exchange Commission (SEC). Measurement Uncertainty in Financial Reporting: How Much to Recognize and How Best to Communicate It. Available at: http://www.sec.gov/about/offices/oca/ocafseries-briefing-measurement.htm</p>
Quantifying rounding errors	<p>Errors: Their origins and how to quantify them. Available at: http://www.iph.ufrgs.br/corpodocente/marques/cd/rd/errors.htm</p> <p>Significant figures and rounding off. In the <i>General Chemistry Virtual Textbook</i>, available at: http://www.chem1.com/acad/webtext/pre/mm3.html</p> <p>See also Wikipedia's page on "Quantization Error" at http://en.wikipedia.org/wiki/Quantization_error</p>
Information graphics	<p>Tufte, E., 2001. <i>The visual display of quantitative information (2nd edition)</i>. Graphics Press USA.</p> <p>Leitch, M., 2003. Design ideas for Beyond Budgeting management information reports. Available at: http://www.internalcontrolsdesign.co.uk/reportdesign/</p>
Characterising patterns in data	<p>"Exploratory Data Analysis" <i>Engineering Statistics Handbook</i>, NIST. Available at: http://www.itl.nist.gov/div898/handbook/eda/eda.htm</p> <p>Leitch, M., 2010. <i>A pocket guide to risk mathematics</i>, John Wiley & Sons, Chichester.</p>
Using correlation to detect causality	<p>Leitch, M., 2006. Better management of large scale financial and business processes using predictive statistics. Available at: http://www.internalcontrolsdesign.co.uk/procpredict/index.shtml</p>
Direct observation of causality	<p>Leitch, M., 2006. Better management of large scale financial and business processes using predictive statistics. Available at: http://www.internalcontrolsdesign.co.uk/procpredict/index.shtml</p> <p>Michotte demonstrations here http://cogweb.ucla.edu/Discourse/Narrative/michotte-demo.swf</p>

Technique	Reading suggestions
Choosing between explanations	<p>Chamberlain, T.C., 1931. The Method of Multiple Working Hypotheses. Available at: http://www.geology.und.edu/gerla/gge487_488_494/chamberlin1890science.pdf</p> <p>There are several demonstrations using Bayes' Theorem Calculators:</p> <p>http://statpages.org/bayes.html</p> <p>http://psych.fullerton.edu/mbirnbaum/bayes/BayesCalc.htm</p> <p>http://psych.fullerton.edu/mbirnbaum/bayes/BayesCalc3.htm</p>
Multiple scenarios	-
Bayesian Model Averaging	<p>Hoeting, J.A., Madigan, D., Raftery, A.E., and Volinsky, C.T., 1999. Bayesian model averaging: a tutorial. <i>Statistical Science</i>, 14(4), p.382–417. Available at: http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.ss/1009212519</p>
Assessing and improving forecasting skill	<p>Brocker, J. and Smith, L.A., 2006. Scoring Probabilistic Forecasts: the importance of being proper. <i>Weather and Forecasting</i>, Vol 22. Available at: http://journals.ametsoc.org/doi/pdf/10.1175/WAF966.1</p> <p>Hubbard, D.W., 2009. <i>The failure of risk management: why it's broken and how to fix it</i>, John Wiley & Sons.</p>
Probability elicitation methods	Part II Probability Elicitation. Available at: http://igitur-archive.library.uu.nl/dissertations/1952513/c5.pdf
Forecast markets	Surowiecki, J., 2005. <i>The Wisdom of Crowds</i> , Random House.
Prediction intervals	<p>Armstrong, J.S. and Green, K.C., 2012. Demand forecasting: evidence-based methods. Available at: http://marketing.wharton.upenn.edu/documents/research/DemandForecasting.pdf</p>
Empirical prediction intervals	<p>Armstrong, J.S. and Green, K.C., 2012. Demand forecasting: evidence-based methods. Available at: http://marketing.wharton.upenn.edu/documents/research/DemandForecasting.pdf</p>
Prediction intervals from propagating uncertainty	<p>Leitch, M., 2010. <i>A pocket guide to risk mathematics</i>. John Wiley & Sons, Chichester.</p> <p>Morgan, M.G. and Henrion, M., 1992. <i>Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis (second edition)</i>. Cambridge University Press.</p>
Objective functions	-

Technique	Reading suggestions
Utility curves	<p>Bernoulli, D., 1738. Exposition of a new theory on the measurement of risk. Available translated from the original Latin at: http://www.econ.ucsb.edu/~tedb/Courses/GraduateTheoryUCSB/Bernoulli.pdf</p> <p>Keeney, R. and von Winterfeldt, D., 2007. Practical Value Models. In W. Edwards, R.F. Miles, & D. von Winterfeldt (Eds.) <i>Advances in Decision Analysis: From Foundations to Applications</i>. New York: Cambridge University Press, pp. 232-252. Available at: http://www-bcf.usc.edu/~winterfe/practical%20value%20models.pdf</p>
Conjoint analysis	<p>Saul Dobney has provided two nice demonstrations of conjoint analysis: http://www.dobney.com/Conjoint/CnjtDemo.htm http://www.dobney.com/Conjoint/ModelDemo.htm</p>
Willingness To Pay	-
Direct choice between distributions of outcomes	<p>Chapman, C.B. and Ward, S.C., 2003. <i>Project Risk Management: Processes, Techniques and Insights (second edition)</i>. John Wiley & Sons.</p>
Mean-variance approach	<p>Ruefli, T. W., 1990. Mean-variance approaches to risk-return relationships in strategy: paradox lost. <i>Management Science</i>, 36(3), p.368-380. Available at: http://www.wiggo.com/mgmt8510/Readings/Readings8B/ruefli1990mgtsoci.pdf</p>
Decision trees	<p>Decision tree primer, by Arizona State University, available here http://www.public.asu.edu/~kirkwood/DASTuff/decisiontrees/index.html</p>
Optimisation	<p>Wikipedia has a lot of good material on mathematical optimisation.</p> <p>"Introduction to optimization with the Excel Solver tool" tutorial by Microsoft, available here http://office.microsoft.com/en-us/excel-help/introduction-to-optimization-with-the-excel-solver-tool-HA001124595.aspx</p>
Expected values	-
Proportional betting	-
Discounting rates	<p>Bowman, E. H., and Moskowitz, G. T., 2001. Real Options Analysis and Strategic Decision Making. <i>Organizational Science</i>, 12(6), p.772-777. Available at: http://www.bus.emory.edu/rcoff/630readings/rbowmanmoskowitz.pdf</p> <p>Luehrman, T. A., 1997. What's it worth? A general manager's guide to valuation. <i>Harvard Business Review</i>, May-June. Available at: http://karlin.sdsmt.edu/640/Luehrman%20-%20Valuation.pdf</p>

Technique	Reading suggestions
Iteration	Chapman, C. B. and Ward, S. C., 2002. <i>Managing project risk and uncertainty: A constructively simple approach to decision making</i> . Chichester, UK, John Wiley & Sons.
Automated calculations	Armstrong, J. S., 2001. Judgmental Bootstrapping: Inferring Experts' Rules for Forecasting. In <i>Principles of Forecasting: A Handbook for Researchers and Practitioners</i> (Ed. J. Scott Armstrong). Kluwer, 2001. Available at: http://repository.upenn.edu/cgi/viewcontent.cgi?article=1178&context=marketing_papers
Value of information	Hubbard, D. W., 2009. <i>The failure of risk management: why it's broken and how to fix it</i> . John Wiley & Sons.
The alternative of thinking or waiting some more	Hill, T., 2009. Knowing when to stop. <i>American Scientist</i> , March-April 2009. Available at: http://www.americanscientist.org/issues/num2/knowing-when-to-stop/1
Model refinement	Rekuc, S. J. and Paredis, C. J. J., 2005. Considering shared epistemic uncertainty in set-based design. Available at: http://www.srl.gatech.edu/Members/srekuc/SJR_EpUn2005.pdf
Structural heuristics	Leitch, M., 2003. Why is Evolutionary Project Management so effective? Available at: http://www.internalcontrolsdesign.co.uk/epmfactors/index.shtml See also Wikipedia's page on Redundancy (engineering), available at: http://en.wikipedia.org/wiki/Redundancy_(engineering)
Agency theory	ICAEW, 2005. Agency theory and the role of audit. Available at: http://www.icaew.com/~media/Files/Technical/Audit-and-assurance/audit-quality/audit-quality-forum/agency-theory-and-the-role-of-audit.ashx Sloof, R. and van Praag, C. M., 2008. The Effect of Noise in a Performance Measure on Work Motivation: A Real Effort Laboratory Experiment. A Tinbergen Institute Discussion Paper available at: http://www.tinbergen.nl/discussionpapers/08074.pdf
Numbers	McGlone, M. S. and Reed, A. B., 1998. Anchoring in the interpretation of probability expressions. <i>Journal of Pragmatics</i> , 30, p.723 – 733. Available at: https://webpace.utexas.edu/mm4994/www/anchoring.pdf Leitch, M., 2007. Favourite ways to characterise risks. Available at: http://www.internalcontrolsdesign.co.uk/pimquizresults/index.shtml
Information Theory	Gray, R. M., 2009. <i>Entropy and Information Theory</i> . New York: Springer Verlag. Available at: http://ee.stanford.edu/~gray/it.pdf