# Introduction

**Luis C. Dias[a], Alec Morton[b] and John Quigley[b]**

[a]University of Coimbra, Portugal

[b]University of Strathclyde, UK

> *I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely in your thoughts advanced to the state of Science, whatever the matter may be.*
>
> Lord Kelvin (1891, p72-72)

## 1. Conceptual background

A useful definition of a problem is that it is a situation where there is a current state, and a desired state, and they are not the same. Most people are familiar with this sort of situation and many day to day problems can be dealt with by largely subconscious or automatic processes (the coffee is too bitter, so I add sugar; the water is too cold, so I turn the tap). But some problems (I want to take up a new hobby, perhaps a new sport, a new language, or a new instrument) require reflection: I have to reflect what goals I want to achieve and whether the actions I have at my disposal will help me achieve them. In such cases I have to build a mental model of my problem to organize my thoughts and help me choose wisely. Other problems, even more complex, involve the significant others in my life (where should we go on holiday?; should we move to a new city, or new country, to take that new job?): in these cases, the model I build should be a shared one, so as to ensure that all those involved in the problem understand what they are getting into. At a higher level still, society has to take important decisions about responses to threats to our environmental and economic wellbeing and security: in a democracy these decisions should take account of the views of the public in some organised fashion.

Tackling problems at multiple levels therefore, requires models of value and models of (our knowledge of) the world around us which may be to a greater or lesser extent implicit or explicit depending on the background nature of the problem. These models can also be more or less complex. For example, in the event of an uncontrolled emission of radioactive materials from a nuclear plant into the atmosphere, the core decision may look as shown in Table 1, where the rows are the choice of actions, the columns are scenarios which may be realized and the consequences in the cells are the outcomes experienced by humanity. This is a very simple model. At the other end of the scale, there are much more complex models (e.g. Geldermann et al, 2009). Such systems may allow (probabilistic) forecasting of wind direction, and model the dispersion of radionuclides and the consequent damage to human health. These more complex models may require drawing on extensive amounts of data and cutting-edge science.

To whom should we go when we wish to deliberate on these models? In the case of models of the world around us, it seems reasonable to privilege experts – those who have relevant knowledge about the subject matter - above lay members of the public. However, identifying these experts may not be straightforward. Senior professors may have long since ceased to keep up with the research literature and be primarily expert at obtaining research funding, and managing grants. Industry experts may be blinkered by social norms and conflicts of interest, especially if they depend for employment or consulting income on other powerful stakeholders. Moreover, if we want to make a genuinely informed decision, we want experts who are able, not just to offer an opinion, but to give us an assessment of how much confidence we can have in their assessment. This requires a cognitive ability which is entirely distinct from actual knowledge.

We may seem to be on safer ground when it comes to models of value. In these cases, surely the person to engage with on whether the value model is the decision maker? Yet this is not a particularly helpful observation. In many situations there is no single unitary decision maker. Even if one person has to sign on the dotted line, the agreement of many people is required if the decision is to be real – is actually to result in action and change. What is more, research strongly suggests that even when making consequential decisions people do not know the goals that they have – even if asked to spontaneously list their goals, there are many other not less important goals which they also recognize as being relevant to their choices (Bond, Carlsson, and Keeney, 2008). Therefore, it is wise to engage in reflective dialogue with their friends and partners about significant choices – even if the decision falls to you alone.

Table 1. A simple model of a decision in nuclear emergency management

| | | Scenarios | |
|---|---|---|---|
| | | Wind blow seawards are radioactive material is dispersed over the ocean | Wind blow landwards and radionuclides are dispersed over land |
| Actions | Evacuate nearby town | Unnecessary evacuation with result cost and hardship | Population are moved out of the path of harm's way |
| | Do not evacuate but encourage people to shelter indoors | Damage and inconvenience both minimised | Population are exposed to potentially hazardous levels of radiation |

This book is a book about elicitation, which may be defined as the *facilitation of the quantitative expression of subjective judgement*, whether about matters of fact or matters of value. Why should anyone want to express their judgements quantitatively, or to help others to do so? So far, we have stressed the role of models in underpinning decision making. But these models are often – and always in the case of models which are exclusively mental models – qualitative in nature: the human mind has not evolved to do floating point operations.

We believe that people should be encouraged to express their judgements quantitatively as a way of making their thoughts precise, and ensuring that they are testable against the evidence from the real world. Statements like "This year there will probably be a lot of rain in Glasgow" or "Artistic self-realisation is more important to me than money" are hopelessly vague: "This year there is a 50% chance of more than 1100mm of rain in Glasgow" and "I would be prepared to take a pay cut of up to £7K per annum to free up a day a week for my theatre workshop" can be tested against the actual realized weather and my actual choices respectively.

For important decisions, this clarity is critical, we believe, if we are to have high quality, transparent engagement of experts and stakeholders; if we seriously care about having high quality deliberative dialogue. It is not that words and qualitative reasoning are not important. However, significant decisions inevitably involve weighing competing risks and values and questions of relative magnitude inevitably arise. The only way to communicate clearly about relative magnitude is through the use of number. For such decisions, words and numbers are jointly necessary and indeed, complementary.

It is true that people are not (yet) accustomed to use numbers to express their judgements of fact or value. For some people this is difficult or uncomfortable; others have an ideological objection to it, as they view quantification as having a technocratic flavour. Yet we believe that the difficulties are overstated. As the chapters in this book show, there are many ways to enable people to express their quantitative judgements, which can be customised to quite different cognitive styles and tastes. Many of the elicitation methods we review involve asking respondents purely qualitative questions: the numbers are, so to speak, “backed out” from their answers.

Our purpose in this section has been to present our motivating philosophy and the conceptual underpinnings of the current volume. In the remainder of this chapter, we discuss in more detail the need for, and barriers to, using elicitation of probabilities and preferences to support decision making, outline the chapters of the book, and in conclusion, present some common themes and ways forward.

## 2. The need for and barriers to elicitation

> *For the purposes of sharpening assumptions and distinguishing them, nothing beats an exercise in probability*. (Neustad and Fineberg, 1983, p.118)

> *Values are what we care about. As such values should be the driving force for our decision making… But that is not the way it is. It is not even close to the way it is*. (Keeney. 1992, p.3)

### *2.1 The need for elicitation of judgement*

In this subsection, our aim is to reflect on the need for elicitation. We do so by considering cases where elicitation was or could have been profitably used.

These in-depth cases will give a sense of the breadth of potential application across time and across domains. Specifically we deal with four areas of applications, which are depicted in Table 2 below: they cover human health (swine flu); provision of public services (airport location); natural hazards (assessment of the risk of earthquakes) and environmental protection (in the case of radioactive waste). Although all cases involve both uncertainties about matters of fact (probabilities) and conflicts about values (preferences), two case studies are better used to highlight the former, and the other two, the latter.

Table 2. Four case studies which illustrate the potential for structured elicitation

| | Assessment of probabilities | Assessment of preferences |
|---|---|---|
| Historic (1960s/ 70s) | Case 1. Swine flu | Case 2. Airport location |
| Recent | Case 3. Assessment of risk of earthquake | Case 4. Radioactive waste |

*Case 1. Swine flu.*

In early February 1977, then US Secretary of Health, Education and Welfare Joseph A Califano Jr was confronted with the decision to release stocks of influenza vaccine; he had been in post for two weeks. The vaccine had been used in autumn 1976 to begin immunizing the nation against swine flu, a strain of the H1N1 influenza, and possibly prevent an epidemic on the scale of the Spanish flu which caused the death of 3% to 6% of the world's population in 1919. The vaccine had been withheld due to possible but not certain links with Guillain-Barre Syndrome, which is an often paralyzing and sometimes killing side effect. This unenviable time pressured task of decision making under uncertainty concerned trading between risks, where traditional "scientific" evidence from controlled lab based experiments did not exist and as such must rely on expertise. Today, the outbreak is most remembered for an unnecessary mass immunization that cost $135 million (Harrell 2009). The virus resulted in one fatality while side-effects from the vaccine are thought to have caused 25 deaths due to Guillain-Barre syndrome (Roan 2009). There is no guarantee that decision making under such circumstances will result in the best outcome post-hoc, however better processes for working with expert judgement seem to have been needed.

Much has been written on this outbreak with the most in-depth critique of the decision making process "The Epidemic That Never Was: Policy Making & The Swine Flu Affair" (Neustad and Fineberg, 1983), published after the event with the aim of learning lessons for the future. While there are a number of confounding issues that led to the decision to attempt to vaccinate the entire US population a key shortcoming identified in the process was the lack of probability assessments, explicitly identifying the need for experts to quantify their uncertainty in terms of probability, exposing their judgment for comparison with one another.

*Case 2. Airport location*

A perennial issue in UK politics over the last several decades has been airport capacity planning in the crowded South-East of the country around London. An instructive episode in this history is the Roskill Commission (Hall, 1980) appointed by the UK government in 1968, and which reported in 1971. The centerpiece of the Roskill Commission's Report was a highly detailed economic cost-benefit analysis ("without doubt the largest and most complex of its kind attempted anywhere" – Hall, 1980, p 32) which involved calculations and monetisation not only of capital investment and passenger time, but also noise impacts, agricultural impacts and the like. The Commission's calculations pointed to a site – Cublington – between London and Birmingham as the best choice. However the publication of the report and the substantive recommendation of Cublington generated a storm of controversy. One commission member wrote an impassioned note of dissent suggesting that the Commission's entire methodology had been misguided as it completely ignored the overriding importance of preserving open countryside. Academic commentators such as Mishan (1970) ("What is wrong with Roskill?") and Self (1970) ("Nonsense on stilts") piled into the discussion with trenchantly expressed take-downs of the study methodology. An important theme of the Mishan and Self critiques is that the Roskill Commission calculations embed disputable and critically important assumptions social values, such as equity. In large part because of a disconnect between the values embedded in the cost-benefit analysis and political and popular perceptions, the Roskill recommendation of Cublington was ultimately rejected and the government chose to explore the option of building an airport at Foulness.

The experience of the Roskill Commission is a reminder that complex decisions are "wicked problems" (Rittel and Webber, 1973) and feature conflicting stakeholders, with multiple, competing, objectives. Effective analyses have to grapple with these features of the problem context rather than wish them away. It is interesting to contrast the mode of analysis of the Roskill Commission with the Multicriteria Decision Analysis (MCDA) described in Keeney and Raiffa (1976) for the location of the new Mexico City airport. This very early decision analysis (originally reported in 1972) nevertheless features the use of computerised sensitivity analysis to explore and communicate the model, in order to assist decision makers to reflect on their value judgements.

*Case 3. Assessment of the risk of earthquake*

In early April 2009, an earthquake struck L'Aquila Italy killing 309 people. Six scientists and one government official who participated in Italy's National Commission for the Forecast and Prevention of Major Risks six days prior to the earthquake were sentenced to six years in prison in October 2012 for manslaughter. The prosecution argued that the expert advice from the Commission resulted in 30 people deciding to stay indoors which resulted in their death. The case led to outrage from many in the scientific community who argued that earthquakes cannot be predicted with certainty, so the trial was seen by some as an attack on science. The prosecuting attorney Fabio Picuti was not criticizing the experts on these grounds, rather on a lack of evaluation of the degree of risk present in L'Aquila; the presiding judge Marco Billi ruled the analysis was superficial. An appeal in November 2014 resulted in all six scientists being acquitted and the government official having his jail sentence reduced to two years, on the grounds that only the government official was responsible for the communication of the risk assessment that led to the death of the 30 individuals. For details see Nature (2011) and Science (2012 and 2014)

A further criticism of the L'Aquila risk assessment identified by Alessandro Martelli and Lalliana Mualchin who were respectively the President and General Secretary of the International Seismic Safety Organisation (ISSO) concerned the dangers of the lack of independence amongst expert judgments (Martelli and Mualchin 2012). This tragedy highlights a need for transparent, rigorous and widely accepted processes for assessing uncertain events.

*Case 4. Radioactive waste management*

In 2003 the UK government set up a Committee on Radioactive Waste Management (CoRWM) to address the problem of what to do with the UK's inventory of radioactive waste. This problem was not new, but acquired new urgency as the current fleet of nuclear reactors was coming to the end of its life and government wanted to commission new nuclear power plants in order to ensure continuity of generating capacity. However, previous efforts to arrive at a solution – involving "deep disposal" of waste stocks in a deep underground repository had left a legacy of popular distrust of the nuclear industry and of the government. CoRWM was asked to take a new alternative approach – open and participative, and capable of inspiring public confidence.

Early on, CoRWM decided that they would systematically involve a broad range of stakeholders and conduct as much as possible of their business in public. However, a challenge was how to reconcile this with the need to actually reach a decision which all members of the committee (who brought a diverse range of views) could actually sign up to. One of the strengths of the CoRWM process was their use of a systematic MCDA as a core (though not the only) component of their deliberative strategy (Morton et al, 2009). The MCDA model provided a transparent basis through which different concerns – for example about safety, or about the need to avoid a burden on future generations – and stakeholder perspectives could be discussed and weighed up against each other. The MCDA also played a key role in communicating the rationale for the decision in the final report (CoRWM, 2006). Thus, CoRWM provides a good example of how explicit elicitation and modelling of value tradeoffs can play an important role in supporting complex societal decisions.

### *2.2 Why do people resist expressing their uncertainty and values quantitatively?*

It is sometimes argued that attempting to employ analytic methods in situations which are characterised by uncertainty and conflict over objectives reflects a technocratic arrogance in the face of a fundamentally uncertain, unpredictable world and/ or a profane disregard for the role of human values in decision making. As examples of the former, *Black Swans* (Taleb, 2007) and *Perfect Storms* (Junger, 1997) are two metaphors used to describe rare events about which there is "deep uncertainty" which is impossible to quantify. As an example of the latter, consider Tetlock's (2003) discussion of "sacred values": "A sacred value can be defined as any value that a moral community implicitly or explicitly treats as possessing infinite or transcendental significance that precludes comparisons, trade-offs, or indeed any other mingling with bounded or secular values" (Tetlock et al, 2000, p 853).

We believe that appeals to "deep uncertainty" or "sacred values" often reflect lazy, superficial thinking about both possible future events and human objectives. Again, we frame our discussion through four case studies: Deepwater Horizon and the Fukushima nuclear disaster for deep uncertainty; the approval of new drugs and the concept of capability in military planning.

*Deep uncertainty Case 1: Deepwater Horizon.*

In April 2010, a geyser of seawater erupted onto the BP Deepwater Horizon rig located in the Gulf of Mexico resulting in the largest offshore oil spill in US history; eleven platform workers were killed and seventeen injured. The National Academy of Engineering and National Research Council (2010) argued early indications of the problem existed from several repeated tests of well integrity. Bea (2010) attributes the cause stemming from the failure of multiple processes, systems and equipment. While this event may appear as a Black Swan as we have never experienced such an event before, it was not beyond the boundaries of reasoned imagination (Paté-Cornell 2012), as early warning signals were present. To model this is possible: we would require an assessment not only of each event but the dependency between events, where all the events which precipitated the disaster are made more likely through a certain management style.

*Deep uncertainty Case 2: the Fukushima disaster*

In March 2011, an earthquake in Japan resulted in the release of seismic energy into a place of convergent boundaries of tectonic plates, i.e. a subduction zone, causing a tsunami that reached 14m. The Fukushima Daiichi nuclear reactors which were designed for a maximum wave height of 5.7m, were affected by the tsunami, resulting in nuclear meltdowns and release of radioactive material. The plants design was deemed safe as the likelihood of a wave in excess of 6m was less than 0.01 in the next 50 years, although historical evidence of such extreme waves existed albeit from the $9^{th}$ and $17^{th}$ Century (Paté -Cornell 2012). Moreover, while the buildings were designed to withstand a tsunami, the plants backup generators were not (Masys 2012). This event illustrates how analogous events for which data exists could inform the identification of events and the assessment of the associated uncertainty on events.

*Sacred values Case 1: the approval of new drugs*

Some of the hardest values to think about systematically are values which relate to one's own quality of life and, ultimately one's own mortality. However, since most of our healthcare is provided by third parties, either governments or insurance funds, there is a need to make tradeoffs since not all medical technologies, which influence one's health and survival are affordable. One tool for structuring such tradeoffs is the Quality Adjusted Life Year or QALY (Pliskin et al, 1980). QALYs provide a numerical assessment of health benefit which integrates quality of life and survival. Roughly speaking, QALYs are calculated via a quality of life score, which reflects different dimensions of quality of life such as level of pain, mental distress or mobility, multiplied by length of life. Over the last two or three decades QALYs (and their variants) have becomes widely used and accepted in many jurisdictions (Drummond et al, 2015), with the precise parameters used to calculate the QALYs being elicited from local populations to reflect local preferences. The success of the QALY in ensuring that public spending on medicines is in line with social values shows the potential of a simple, yet theoretically robust concept in making previously taboo tradeoffs discussable in the public sphere.

*Sacred values Case 2: the concept of "capability" in military planning*

A common way in which values become sacred in organizational management is where values are specified at an insufficiently strategic level. Protection of existing programmes becomes identified with loyalty to one's division of the organisation and accepting reallocations becomes identified with surrender. Addressing such issues requires creating an overarching conceptual framework in which the contribution of individual programmes to the common good can be traced and articulated. In businesses, profitability often provides this framework but in other sorts of organisations, the path to constructing such a framework might be less obvious. A good example of such a framework in a non-business setting is the idea of "military capability" which has been recently popular in countries such as the UK, US, Australia and Finland (Anteroinen, 2012). The idea in such frameworks is to substitute arguments between individual services about how many ships, tanks, or planes with arguments about how to deliver particular capabilities: for example a monitoring capability may be delivered by human reconnaissance, UAVs or satellites. Once this substitution has been made, it is possible to reframe decisions away from being about which branch of the service suffers and towards what constitutes the best way of delivering the ability to meet national military needs.

Reviewing the above cases, we freely admit that eliciting probabilistic or tradeoff information may be difficult: but we argue that the proper response is not to declare that the problems are somehow too profound for quantitative thinking to be useful, but rather to think carefully and creatively about what the difficulties are and how to tackle them. With this motivation, the rest of the book represents a sourcebook of methods and concepts for doing this.

## 3. Overview of the book

The idea for this book originated in the COST Action "Expert Judgment Network: Bridging the Gap Between Scientific Uncertainty and Evidence-Based Decision Making",[1] noting the importance of using sound elicitation processes when building models to inform decision making. Elicitation may be needed to populate models of uncertainty, interacting with subject experts, but they may also be needed to set up models of preferences, interacting with experts, decision makers, and other stakeholders. In both cases, it is important that analysts and experts follow process that allows them to think clearly about numbers, whether they concern probabilities or they concern the importance of attributes, for instance. Hence, this book covers elicitation processes having in mind both probabilities and preferences.

A first major group of chapters in this book (Chapters 2 to 9) focusses on processes to elicit uncertainty from experts. Chapter 2, by Quigley, Colson, Aspinall and Cooke, presents the Classical Method for aggregating judgements from multiple experts concerning a probability distribution; the method uses mathematical aggregation based on the performance of experts. In Chapter 3, Cooke discusses the issue of validation: what constitutes good uncertainty assessment and how can this be measured. This chapter addresses in particular the Classical Method, for which many studies have already been carried out. Chapter 4, by Gosling, presents the Sheffield elicitation framework, also to elicit probability distributions, covering its foundations, its extensions, and its applications. In contrast to the Classical method, a behavioral aggregation method is proposed in

---

[1] http://www.expertsinuncertainty.net/

this chapter. Chapter 5, by Hanea, Burgman and Hemming, outlines a protocol named IDEA, which is a mixed approach combining behavioral and mathematical aggregation techniques that can be used instead of, for example, the well-known Delphi protocol. The two ensuing chapters present approaches based on the principles of Bayesian updating of probability distributions. Chapter 6, by Hartley and French, discusses how one might use a full Bayesian model to combine the judgements of multiple experts into a posterior distribution, considering prior experts' judgements as data. In Chapter 7, Quigley and Walls present an approach to represent expert uncertainty through analogies with existing empirical data so reducing the burden of quantification on experts.

Chapters 8 and 9 address important issues that are of relevance for different probability elicitation approaches. Chapter 8, by Werner, Hanea and Morales-Nápoles, discusses the main elements of structured expert judgement processes for dependence elicitation, when eliciting multivariate distributions using either pooling or Bayesian approaches. Chapter 9, by Wilson and Farrow, discuss how mathematical methods for expert judgement aggregation, whether opinion pooling or Bayesian methods, can incorporate correlations between experts; they also consider behavioral approaches and the potential effects of correlated experts in this context.

A second major group of chapters in this book (Chapters 10 to 14) focusses on processes to elicit preferences from stakeholders or decision makers. The elicitation processes covered here consider situations in which a decision maker (or a group) needs to make a decision. The purpose of the elicitation is then to build a model of the preferences of the decision maker (and often, preferences of other stakeholders) that helps this decision maker in making sound and informed decisions. The first two chapters on preference elicitation deal with problems under uncertainty. Chapter 10, by Gonzalez-Ortega, Radovic, and Ríos Insua, presents the classical decision analysis paradigm of utility theory, to elicit models of preferences that can be combined with models of uncertainty. They cover the case of preferences concerning one attribute, multiple attributes, and the preferences of adversaries. In Chapter 11, Bordley presents a different perspective on utility elicitation based on targets, which allows using probability elicitation methods to elicit utility functions.

Chapters 12 to 14 address the elicitation of preferences independently of, or in the absence of, any uncertainty elicitation. These chapters concern preferences over multiple attributes or evaluation criteria. In Chapter 12, Morton presents the multiattribute value theory, from its foundations to process aspects important in practice. Chapter 13, by Matsatsinis, Grigoroudis and Siskos, is based on the same theoretical grounds, but introduces a new outlook on a very different elicitation paradigm, which involves eliciting preferences indirectly, "disaggregating" comparisons that a decision maker is able to make at an holistic level into the components of a multiattribute value function. In Chapter 14, Dias and Mousseau discuss the elicitation of an outranking-based preference model, considering in particular ELECTRE methods, which are based on principles that are different from the value measurement framework of the preceding chapters.

Chapters 15 and 16 are about cross-cutting issues that are relevant for elicitation of uncertainties as well as for elicitation of preferences. In both cases, the experts or stakeholders involved can incur into biases leading to answers that, upon reflection, they would wish to revise. In Chapter 15, Montibeller and von Winterfeldt overview the biases that individuals and groups are subject to, and also what might be done to reduce the occurrence of such biases. Chapter 16, by Bolger, addresses another issue present in any elicitation process involving expert judgement, which is the selection of the experts. This chapter presents a structured process having in mind mainly probability elicitation, but it is also relevant for preference elicitation.

The last group of chapters illustrates how some of the approaches presented in this book can be, and are being, applied in practice. In Chapter 17, Barons, Wright and Smith describe an integrating decision support system for probabilistic judgement elicitation (under a Bayesian approach), and they illustrate its potential on a food security case in the UK. Chapter 18, by Soares and Bojke, illustrates expert elicitation in health care decision making, discussing two examples of formal elicitation to inform Health Technology Assessments in the UK. Chapter 19, by Merrick and McLay, demonstrates an expert judgment based method using pairwise comparisons and parameter estimation to elicit nuclear threat risks concerning the security of US ports. In Chapter 20, Porthin, Rosqvist and Kunttu present a risk assessment concerning a new logistics system for a pulp and paper manufacturer in the Nordic countries, using a computer system designed to support decision-making (i.e. that can also be used to elicit preferences). Chapter 21 differs from the previous applications In that it deals with elicitation of preferences rather than uncertainties, but it also considers a situation involving a group of individuals. In this chapter, Delias, Grigoroudis and Matsatsinis present a case study of applying a multicriteria disaggregation approach to elicit a model of preferences for crop planning in a Greek region.

## 4. Conclusions and Future Directions

As we have emphasised in our initial section, we see the aim of elicitation as being to facilitate the quantitative expression of subject judgement, not as an end in itself, but to facilitate high quality dialogue and reflection about important decisions. As we stress in the historical survey in our overview of the chapters themselves, we think there is now a vibrant applied discipline which draws on a rigorous and well-developed theoretic base, and which has provided us with a toolbox of techniques, each custom-developed to meet the needs of particular sorts of problems and the preferred cognitive styles of different sorts of people.

The philosophy of this book is that quantification, through elicitation, is a way to refine and clarify the mental models which people inevitably use in thinking about complex problems. Quantification enables clearer communication about these models between people – all sorts of people – but also sharpens the predictions which these models make about the world and enable them to be tested empirically. A book on this subject is (in our view) particularly timely because of the following trends in the world.

- **Increasing range of choice**. Our experience of the world is increasingly mediated through digital technologies which are global in reach. This means that we are routinely confronted with choices broader than ever before. A trivial example is that we are now able to download virtually any book in print from Amazon onto my Kindle reader: our choices about our education, career, potential political or religious beliefs, choice of life partner, etc, have been similarly broadened. We need aids which will help us organise and make sense of these complex choices, and which will enable us to weigh, select and aggregate, and ultimately make better and more life-enhancing decisions.
- **Increasing availability of data.** As more transactions are conducted online and as the cost of data storage and processing drops, businesses and governments have been increasingly able to collect, process, and make available large volumes of data relating to their activity. Unfortunately, gathering data does not itself bring insight, and in the absence of a strong research design, making inferences about what caused what, and the generalising from *then* and *there* to *now* and *here* can be extremely difficult. To build meaning from this data requires somehow infusing the data with expert judgement.

- **Increasing demands for accountability.** During the debate prior to the 2016 Brexit referendum, the UK politician Michael Gove remarked that the "people in this country have had enough of experts" (Financial Times, 2016). Whatever one's view on the substantive issue, the outcome of that referendum clearly validated Gove's claim: the UK voting public did not trust the experts who predicted that leaving the EU would be a disaster, or the elites who purported to take decisions in their best interests. A technological society cannot survive without experts, or without political officeholders, but lay people may reasonably demand confidence that expert judgements - and claims to expertise - are as open to scrutiny and testing as possible, and that the values which inform public decisions are subject to open and transparent debate.

The tools in this book have a vital role meeting all of these challenges: empowering purposeful decisions in the face of these overwhelming choices; making sense of vast, complex and ill-structured datasets; and building bridges between experts and elites on the one hand, and (perhaps rightly) suspicious lay people on the other.

Elicitation is a young technology. Other quantitative technologies – counting, measurement of physical dimensions – have been around for millennia. Yet other quantitative technologies – cost accounting for example – have become well established in the space of a few decades when it became clear that they met a need of a modern complex industrial society. There is the potential for elicitation of value and uncertainty to have a no less central role in the society of the future. We hope that this book will give the reader some ideas as to how that might come about.

## References

Anteroinen J (2012) Integration of existing military capability models into the Comprehensive Capability Meta-model. Paper presented at the SysCon, Vancouver,

Bea R (2010) Failures of the deepwater horizon semi-submersible drilling unit. Deepwater Horizon Study Group, University of California, Berkeley. https://ccrm.berkeley.edu/pdfs_papers/bea_pdfs/DeepWaterBobBeaPrelimAnalyses-rev5-2.pdf. Accessed 8/6/2017.

Bond SD, Carlson KA, Keeney RL (2008) Generating objectives: Can decision makers articulate what they want? Manage Sci 54 (1):56-70

CoRWM (2006). Managing our radioactive waste safely. https://www.gov.uk/government/publications/managing-our-radioactive-waste-safely-corwm-doc-700. Accessed 8/6/2017.

Drummond M, Sculpher MJ, K C, GL S, Stoddart GL (2015) Methods for the Economic Evaluation of Health Care Programmes. Fourth edition. Oxford University Press, Oxford

Harrell E (2009) "How to Deal with Swine Flu: Heeding the Mistakes of 1976". Time Magazine http://content.time.com/time/health/article/0,8599,1894129,00.html Accessed 12/6/2017

Financial Times (2016) Britain has had enough of experts, says Gove https://www.ft.com/content/3be49734-29cb-11e6-83e4-abc22d5d108c?mhq5j=e1. Accessed 8/6/2017.

Geldermann J, Bertsch V, Treitz M, French S, Papamichail KN, Hamalainen RP (2009) Multi-criteria decision support and evaluation of strategies for nuclear remediation management. Omega-Int J Manage S 37 (1):238-251

Hall P (1980) Great Planning Disasters. University of California Press, Berkley.

Junger S (1997) The Perfect Storm. Norton, New York

Keeney RL, Raiffa H (1976) Decisions with multiple objectives: preferences and value tradeoffs. Wiley, Chichester

Kelvin, Lord (1891) Popular Lectures and Addresses. MacMillan, London and New York

Martelli A and Mualchin L (2012) Indictment and conviction of members of the Italian "Commissione Grandi Rischi" (open letter to the President of Italy). www.cngeologi.it/wp-content/uploads/2012/10/CoverletterandStatementISSO1.pdf. Accessed 8/6/2017.

Masys, AJ (2012) Black swans to grey swans: revealing the uncertainty. Disaster Prevention and Management: An International Journal 21 (3): 320-335.

Mishan EJ (1970) What Is Wrong with Roskill. J Transp Econ Policy 4 (3):221-234

Morton A, Airoldi M, Phillips LD (2009) Nuclear Risk Management on Stage: A Decision Analysis Perspective on the UK's Committee on Radioactive Waste Management. Risk Anal 29:764-779

National Academy of Engineering and National Research Council (2010) Interim report on causes of the Deepwater Horizon oil rig blowout and ways to prevent such events. www.nap.edu/catalog/13047.html (accessed 8/6/2017).

Nature (2011) Scientists on trial: At fault? www.nature.com/news/2011/110914/full/477264a.html. Accessed 8/6/2017.

Neustadt RE, Fineberg HV (1983) The epidemic that never was. Vintage, New York.

Paté-Cornell E (2012) On “Black Swans” and “Perfect Storms”: risk analysis and management when statistics are not enough. Risk analysis 32(11): 1823-1833.

Pliskin JS, Shepard DS, Weinstein MC (1980) Utility functions for life years and health status. Oper Res 28 (1):206-224

Rittel HWJ (1973) Dilemmas in a General Theory of Planning. Policy Sciences 4 (2):155-169

Roan, S (2009). "Swine flu debacle of 1976 is recalled". Los Angeles Times. http://articles.latimes.com/2009/apr/27/science/sci-swine-history27?pg=1. Accessed 12/6/2017

Science (2012) Earthquake Experts Convicted of Manslaughter.

www.sciencemag.org/news/2012/10/earthquake-experts-convicted-manslaughter. Accessed 8/6/2017.

Science (2014) Updated: Appeals court overturns manslaughter convictions of six earthquake scientists www.sciencemag.org/news/2014/11/updated-appeals-court-overturns-manslaughter-convictions-six-earthquake-scientists. Accessed 8/6/2017.

Self P (1970) Nonsense on Stilts - Cost-Benefit Analysis and Roskill-Commission. Polit Quart 41 (3):249-260

Taleb, NH (2007) The black swan: The impact of the highly improbable. Random house, New York.

Tetlock PE (2003) Thinking the unthinkable: sacred values and taboo cognitions. Trends Cogn Sci 7 (7):320-324

Tetlock PE, Kristel OV, Elson SB, Green MC, Lerner JS (2000) The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. J Pers Soc Psychol 78 (5):853-870