Making Responsible Decisions (When it Seems that You Can't)
Engineering Design and Strategic Planning Under Severe Uncertainty

What happens when the uncertainties facing a decision maker are so severe that the assumptions in conventional methods based on probabilistic decision analysis are untenable? Jim Hall and Yakov Ben-Haim describe how the challenges of really severe uncertainties in domains as diverse as climate change, protection against terrorism and financial markets are stimulating the development of quantified theories of robust decision making.

When random factors influence the future performance of design options or policy alternatives, then decision makers can readily use probabilistic decision analysis. Unfortunately, some of the most serious decisions confronting engineers, managers and policy-makers contain very severe uncertainties that, we will argue, do not succumb to the conventional methods of probabilistic decision analysis. Take protection against terrorism, where designers responsible for reducing risks to infrastructure systems are faced by adversaries whose behaviour is wilful and malicious. Game theory has something to say about analysis of this type of problem, but relies on being able to second guess the adversary's likely priorities and strategies. Consider also the fluctuations of stock or commodity markets, which exhibit random behaviour over some time scales, but also discontinuities that the best of analysts fail to predict. In the field of structural safety, David Blockley has argued that engineering failures are a result of the unforeseen consequences of human actions, about which conventional reliability theory has little or nothing to say.

Making decisions about how to deal with climate change presents a particular challenge because the future we plan for will be influenced by a host of human economic, technological and political choices that have yet to be made. This is an instance of a fundamental limit, sometimes referred to as Shackle-Popper indeterminacy, to what can be predicted about systems including human agents who will learn from new information and modify their behaviour accordingly.

Decision making under severe uncertainty

All of the situations discussed above have similar characteristics of severe uncertainty, by which we mean that evidence upon which to base a decision is scarce and only of limited relevance to predicting what may happen in the future. Situations of severe uncertainty are characterised by the distinct possibility of surprise, of situations that were unforeseen or unexpected. This is more than a question of analysing the extreme tails of a probability distribution – it has to do with fundamental and unexpected shifts in the functioning of the system. In the face of this type of uncertainty, analysts may attempt to construct a probability distribution over the future states of nature, but there will be precious little evidence upon which to base such a distribution. Experts may have intuitions about what is probable, but can be hard pressed to quantify the rare but possible. This raises fundamental questions about the representation of uncertainty, which have, especially since the 1970s, fuelled a proliferation of mathematizations of uncertainty that depart in one way or another from Kolmogorov's axioms of probability, including fuzzy set theory, evidence theory and the theory of imprecise probabilities. Connected, and more fundamental, is the question of how in practice agents should go about making decisions in the face of this type of uncertainty. In situations of severe uncertainty, should we seek to optimise the outcome or would it be wiser to make a decision that delivers a satisfactory outcome under a wide range of possible conditions? We refer to the latter of these two approaches as “robust satisficing”. To ‘satisfice’ means to ‘achieve an acceptable outcome’, as opposed to achieving the best possible outcome.

There is always a trade-off between optimising and robust satisficing. It is a matter of fact that optimal designs, which maximally exploit every feature and resource of the system, are minimally robust to conditions that depart from the assumptions under which the optimisation was conducted. Consider London’s Heathrow airport, which operates at 98.5% capacity making it the most highly optimised airport in Europe, but air passengers will know that Heathrow is susceptible to the knock-on effects of unplanned disturbances.

In the face of severe uncertainty it becomes necessary to trade-off some outcome quality against the confidence (robustness to uncertainty) in attaining that outcome. The decision maker invests some resources in managing uncertainty, at the expense of garnering tangible benefit directly from those resources.

Scenario analysis

Governments are waking up to the importance of robust solutions in some of the most important policy areas. For example the UK policy on flood risk management calls for “adaptability to climate change through robust and resilient solutions”. Yet much less is said about how these solutions might actually be identified and analysed.

The UK Government’s Foresight Flood and Coastal Defence project used socio-economic scenarios to look 30 to 100 years into the future at flood risk in the UK in an attempt to find robust strategies for flood risk management. Whilst there are many versions of scenario analysis, they all tend to be based around construction of a small number of contrasting yet internally consistent narratives about the future. As well as the UK Foresight programme (Figure 1) a narrative scenario approach underpins the Intergovernmental Panel on Climate Change’s Special Report on Emissions Scenarios. The restriction to a small number of (typically four) possible futures is motivated by the need to keep the analysis manageable and intelligible for human participants, but overlooks the capacity to use computer simulations to explore much larger spaces of possible futures.

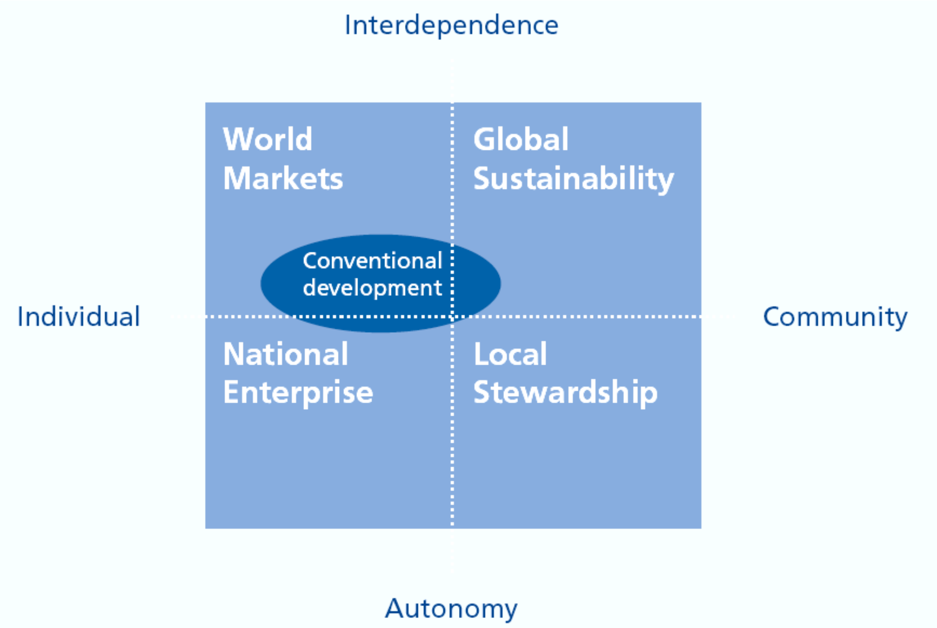

Figure 1: The Foresight Futures scenario grid of four scenarios (Office of Science and Technology, 2002)

In recent years Rob Lempert and Steve Bankes at the RAND Corporation have developed computer-intensive simulation models for analysing the possible outcomes of policy decisions over large spaces of possible futures. The approach recognises the deficiencies in any model of a complex system and does not attempt to represent the uncertainties in probabilistic terms. Rather, the approach, which the team at RAND refer to as ‘Robust Decision Making’, is based upon identifying options that perform acceptably well over the widest subset of the space of possible futures. The problem still remains, however, of specifying the range of that space of possibilities. Actually, if decision makers and their analysts know anything, it is about the central tendencies, not the bounds of variation.

A similar problem of precisely specifying the bounds of possibility besets the use of the theory of Imprecise Probability, which offers a coherent alternative to conventional probabilistic representation of uncertainties when probability distributions cannot be uniquely specified. Imprecise Probability is attracting enormous theoretical interest, but in practice the results are usually determined by the location of the extreme outer points of an uncertainty representation – the points about which a decision maker may be least confident.

Info-gap theory

An alternative approach is to start an uncertainty analysis at our best estimate and then examine how decision options perform as conditions depart increasingly from expectations, without identifying a worst case. If an option continues to be preferred, even at a very large horizon of uncertainty, then it is thought to be robust to uncertainty. In order to do this type of analysis it is necessary to have some understanding of the ways in which the future may depart from expectations. If there are many uncertain variables that together determine the performance of the decision options, then insights into how they co-vary can be used to construct the uncertainty model. However, it is not necessary to distribute a probability measure or for that matter any other measure over the space of possibilities. Like RAND’s robust decision making method, uncertainty is represented in terms of sets of possibilities. However, in info-gap theory, the boundaries of those sets are unknown and vary up to a horizon that is measured by an unknown positive scalar α.

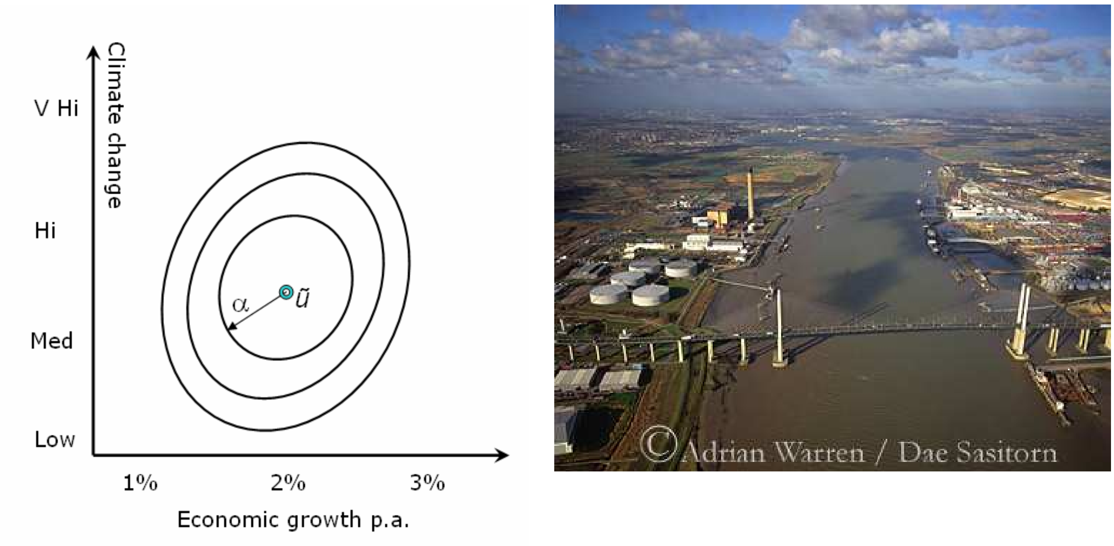

Figure 2: Illustration of an info-gap analysis dealing with two severe long term uncertainties in the Thames Estuary: the rate of economic growth (and hence vulnerability and exposure to flooding) and the rate of climate change.

Figure 2 illustrates an instance of this approach applied to long term uncertainties in climate (including sea level, storm surge frequency and rainfall in the Thames catchment) and economic growth (including sectoral, demographic and technological changes) in the Thames Estuary where London is situated. Of course we do not know for sure what the future holds on either of these axes, though we can obtain evidence from models and analysis of trends to make some estimates. In the face of this uncertainty, an info-gap analysis allows us to examine how each of a set of options for long term flood risk management performs as the horizon of uncertainty (scaled by α) expands from our best estimate $\tilde{u}$.

Evaluating strategies for protection against bioterrorism

To understand how info-gap analysis can help to inform strategic decision making, consider the management of an epidemic. SARS, AIDS and avian flu have caused serious injury, economic loss, and death in recent decades. Civilian populations face serious threat of attack with biological agents by rogue states or terror groups. The impact of a potential epidemic is evaluated with data and models from epidemiology and other fields. One simple epidemic model assumes constant population (no deaths, or death occurs much more slowly than infection), and infected individuals continue to infect the remaining susceptible population. The model contains four central parameters. β determines the rate of infection of susceptible by infected individuals and is very hard to predict. t_c is the duration from detection of the epidemic until medical treatment has been dispensed, which depends, amongst other factors, on the size of the population to be treated. y_0 is the number of initially infected people. N is the size of the population within which the disease can spread, which can be limited by quarantine.

Consider the choice of t_c (delay until treatment) and N (quarantine size). For any given plan, (N, t_c), we must ask: by how much can the models err, while acceptable outcomes are still guaranteed? The answer to this question, for various levels of morbidity and three different plans, are the robustness curves shown in Figure 3. The horizonal axis is the morbidity, y_c, and the vertical axis is the robustness: the greatest fractional error in the estimated mixing rate, β, up to which the corresponding morbidity, y_c, will not be exceeded.

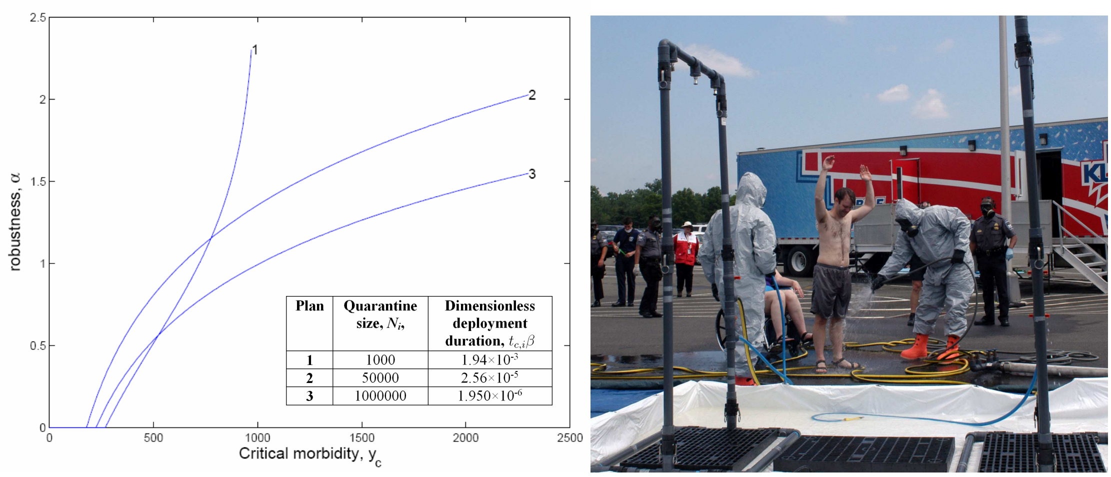

Plan	Quarantine size, N_i,	Dimensionless deployment duration, $t_{c,i}\beta$
1	1000	1.94×10^{-3}
2	50000	2.56×10^{-5}
3	1000000	1.950×10^{-6}

Figure 3: Robustness vs. critical morbidity, for three different quarantine plans. The curves illustrate the maximum number of infections at a given horizon of uncertainty α

We see two important features of the robustness curves in Figure 3. First, the slopes are positive, indicating trade-off between robustness and morbidity: morbidity, y_c, can be reduced only by also reducing the robustness against error in the models and data. Second, the robustness becomes zero at the value of morbidity predicted by the best available models and data. The best estimate of the mixing rate, β, indicates that morbidity in plans 1, 2 and 3 will not exceed 268, 182 and 227 infections, respectively. However, even tiny errors could result in greater morbidity. Best-model predictions are a poor basis for evaluating a plan.

Plan 2 has lower estimated morbidity than the other two plans displayed in Figure 3, and continues to guarantee lowest morbidity up to robustness of about 120% (1.2 fractional error). However, in light of the many unknown factors which can influence the infectious-mixing rate, β, the decision maker may well want far greater robustness. In this case plan 1 – severe quarantine – is indicated due to its greater immunity to uncertainty. Nonetheless it must be recognized that this greater robustness is obtained only at the expense of accepting the possibility of greater morbidity. The large robustness premium for plan 1 is particularly striking since this plan has the largest best-estimated morbidity. Making a strategic decision will not be easy, but with this type of analysis the implications of inevitable modelling inadequacies are made explicit.

Planning flood defence systems for an uncertain future

The floods New Orleans illustrated the potential for surprisingly damaging flooding. It is important to ensure that new flood defence systems are robust to future uncertainties. For illustrative purposes two idealised design concepts are considered for protecting an urban area from flooding, either building levees or widening the existing river channel. Deciding which option is preferred involves analysing the costs of each option and also the benefits in terms of reduction in flood risk. Calculating flood risk involves analysis of the probability of flooding reaching certain depths, which is a probabilistic risk assessment. Further insight can be gained by 'wrapping' an info-gap analysis around this probabilistic assessment.

The best models predict that the Net Present Value (NPV) of the levee is \$13 million, while the NPV of channel-widening is \$10 million. Based on these best-model estimates, one would prefer the levee. However, the models upon which these estimates are based are uncertain. Uncertainties in both the hydraulic modelling of flood depths and the flood frequency analysis have been analysed, but Figure 4 shows only the uncertainty in the flood frequency analysis. The negative slope expresses the trade-off between NPV and robustness. We see that the best-estimated NPV values have zero robustness, and thus are not a good basis for choosing between the designs. Furthermore, the two designs have essentially the same positive robustness over a substantial range of NPV values. In other words, the robust satisficing analysis demonstrates that the channel-widening option is still a serious contender if uncertainty is taken into account.

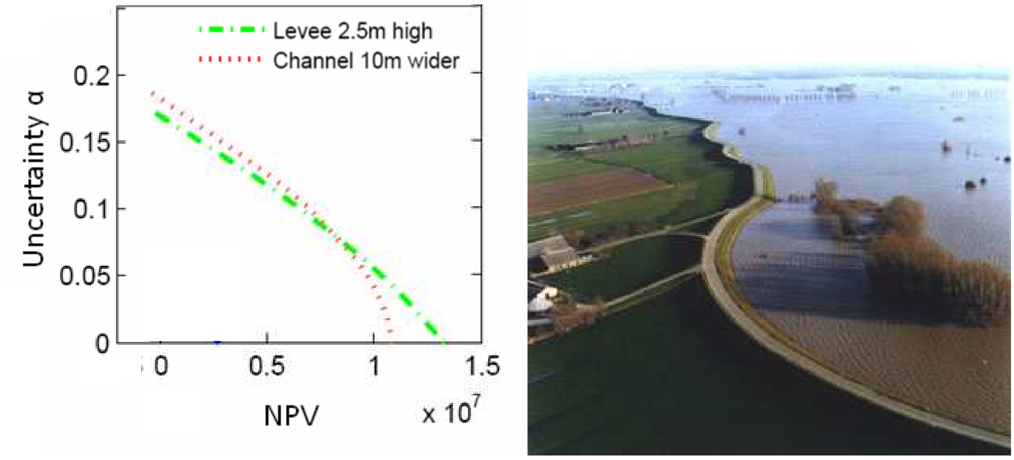

Figure 4: Robustness curves for alternative flood risk management options (courtesy Dan Hine). The curves illustrate the minimum level of NPV that can be guaranteed at a given horizon of uncertainty α

System identification and robust control

Our discussion so far has dwelt upon strategic and design decisions, but robustness may be attractive in a host of technical settings. A commonplace problem for numerical modellers is the identification of model parameters from observed data. This applies to parameter identification in complex dynamical numerical models, for example finite element models (Figure 5) or identification of black box models such as Artificial Neural Networks. In each of these settings there are parameter, and often hyper-parameter (e.g. model grid resolution), choices to be made. Each of these choices is susceptible to inadequacies in available data and models, and have been the object of successful applications of info-gap theory.

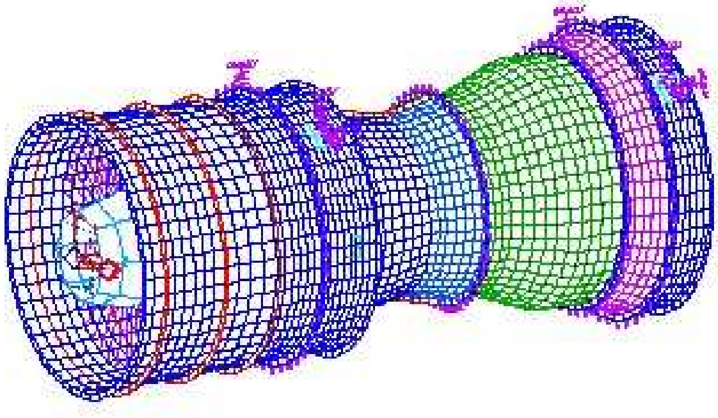

Figure 5: In analysis and control problems, engineers will wish to identify methods that are robust to data limitations and model deficiencies (courtesy of Scott Cogan, Applied Mechanics Lab., University of Franche-Comte, Besancon, France)

Control engineers face similar problems of maintaining system stability even if the system to be controlled differs from the model employed. In this setting, like those discussed previously, robustness analysis can help to identify controllers that are tolerant to substantial changes in the system that are not reflected in the historical data and to imperfections in the mathematical model of the system.

Robust decision making

Info-gap theory does not make the decision makers' decisions for them. It would be both naïve and arrogant to suppose that automated decision making would be possible for the types of decisions we are interested in. Indeed, responsible decision making requires designers and policy makers to incorporate objectives, constraints and intuitions that are not included in the formal decision analysis. This is particularly important when dealing with severe uncertainty, of which we have discussed a few characteristic examples. It is also particularly valuable for highly contested group decisions, where the arrival of an analyst with a nominally 'optimal' solution can be counter-productive if it does not admit alternative perspectives or concerns.

In any decision making situation it is the decision maker's responsibility to reflect upon the level of performance that is needed, and the level of uncertainty to guard against. Info-gap theory provides new insights for the decision maker by demonstrating how performance may deteriorate if the future does not turn out precisely as expected. These insights are not directly available from conventional probabilistic decision analysis. Info-gap analysis allows the decision maker to identify solutions that perform satisfactorily well under the widest possible range of conditions. This is a departure from the conventional approach to decision making, that seeks to maximise performance, but under conditions of severe uncertainty, a guaranteed level of performance (up to some horizon of uncertainty) may well be more attractive than optimal performance that is vulnerable to the unexpected.

Further References

Ben-Haim, Y. 2006. *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edition. Academic Press, London.

Evans, E.P. and Hall, J.W. 2004. A new climate for flood planning. *Ingenia*, No.21 pp.51-56.

Lempert, R., Popper, S.W. and Bankes, S.C. 2003. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. RAND, Santa Monica CA.

Office of Science and Technology, 2002. *Foresight Futures 2020: Revised Scenarios and Guidance*. Department of Trade and Industry, London.

Author biography

Jim Hall PhD CEng FICE holds the Chair in Earth Systems Engineering in the School of Civil Engineering and Geosciences at Newcastle University. He is Deputy Director for Engineering in the Tyndall Centre for Climate Change Research. His research is underpinning 21st Century flood risk management and other adaptations to the threats of climate change through the development of new risk and uncertainty analysis methodologies.

Yakov Ben-Haim PhD is a professor in the Faculty of Mechanical Engineering at the Technion - Israel Institute of Technology, Haifa, Israel, and holds the Yitzhak Moda'i Chair in Technology and Economics. The author of 4 books and numerous articles, he has developed info-gap decision theory which has found application in engineering, biological conservation, management, economics, homeland security, and medicine. He has lectured at universities, medical and technological research institutions, and central banks around the world.