Appearing in *Risk Analysis,* 32(10): 1638–1646.

Perspective

Why Risk Analysis is Difficult, and Some Thoughts on How to Proceed

Yakov Ben-Haim[1]

Contents

1 Introduction 1

2 Why Risk Analysis is Difficult 2
2.1 A Bit of History 2
2.2 Shackle-Popper Indeterminism 2
2.3 Hume and the Problem of Induction 3

3 Some Thoughts on How to Proceed 4
3.1 Epistemic Paralysis 4
3.2 Models and Robustness 4
3.3 Robustness and the Innovation Dilemma 5
3.3.1 Problem Formulation 5
3.3.2 Two Solution Strategies 5
3.3.3 Robustness and Probability 7
3.3.4 Summing Up: Robust-Satisficing and the Innovation Dilemma 7
3.4 Robustness and Worst Cases: Two Approaches 8
3.5 The Other Side of the Coin: Opportuneness 9

4 Conclusion 10

5 References 10
Keywords: robustness, uncertainty, info-gap theory, satisficing, optimizing

Abstract

Risk analysis is challenged in three ways by uncertainty. Our understanding of the world and its uncertainties is evolving; indeterminism is an inherent part of the open universe in which we live; and learning from experience involves untestable assumptions. We discuss several concepts of robustness as tools for responding to these epistemological challenges. The use of models is justified, even though they are known to err. A concept of robustness is illustrated in choosing between a conventional technology and an innovative, promising, but more uncertain technology. We explain that non-probabilistic robust decisions are sometimes good probabilistic bets. Info-gap and worst-case concepts of robustness are compared. Finally, we examine the exploitation of favorable but uncertain opportunities and its relation to robust decision making.

1 Introduction

This essay discusses two ideas: uncertainty and robustness. Uncertainty is a central challenge in the analysis of risk. In its severe forms, uncertainty is a lack of information, a deficiency of understanding, or the potential for surprise. In section 2 we use historical and philosophical

[1] Yitzhak Moda'i Chair in Technology and Economics, Faculty of Mechanical Engineering, Technion—Israel Institute of Technology, Haifa 32000 Israel, yakov@technion.ac.il

approaches in attempting to understanding the nature and origin of severe uncertainty. Section 3 presents some responses to the challenges of uncertainty, based mainly on concepts of robustness.

2 Why Risk Analysis is Difficult

2.1 A Bit of History

The ancient Greeks invented the axiomatic method and used it in the study of mathematics.[1,2] Some medieval thinkers explored the mathematics of uncertainty,[3] but it wasn't until around 1600 that serious thought was directed to the systematic study of uncertainty.[4] In the mid 18th century Thomas Bayes launched the modern idea of conditional probability and its application to inference.[5] Statistics as a separate and mature discipline emerged only in the 19th century,[6,7] and Kolmogorov provided an axiomatization of probability in the 1930s.[8] The 20th century saw a florescence of uncertainty models. Lukaczewicz discovered 3-valued logic in 1917,[9] and in 1965 Zadeh introduced his work on fuzzy logic.[10] In between, Wald formulated a modern version of min-max in 1945.[11] A plethora of other theories, including P-boxes,[12] lower previsions,[13,14] Dempster-Shafer theory,[15,16] generalized information theory[17] and info-gap theory[18] all suggest that the study of uncertainty will continue to grow and diversify. This suggests that innovations, discoveries, inventions, surprises, errors, and misunderstandings are to be expected.

Furthermore, the diversity of models of uncertainty reflects the variety of uncertainty itself. This indicates the value of methodological pluralism. The risk analyst needs a diverse set of tools for modeling and managing uncertainty. Uncertainty is not monolithic, so its treatment must be adapted to each circumstance.

2.2 Shackle-Popper Indeterminism

Risk analysis is difficult, in part, because there are many things that we do not yet know. This has an important implication that was developed separately and in different ways by the economist Shackle[19, pp.3–4, 156, 239, 401–402] and the philosopher Popper,[20, pp.80–81, 109] and that I will refer to as Shackle-Popper indeterminism.[21]

The basic idea is that the behavior of intelligent learning systems displays an element of unstructured and unpredictable indeterminism. By *intelligence* I mean: behavior is influenced by knowledge. This is surely characteristic of humans individually, of organizations and of society at large: what we know influences how we behave. Risk analysts confront this when people alter their behavior as a result of new facts, theories, products or processes. By *learning* I mean a process of discovery: finding out today what was unknown yesterday. Shackle-Popper *indeterminism* arises as follows:[19,20] because tomorrow's discovery is by definition unknown today, tomorrow's behavior is not entirely predictable today.

Given the richness of future discovery, (or conversely, the richness of our current ignorance), future behavior is incompletely determined by the past. The patterns and laws of behavior will grow or evolve in time as agents make discoveries. These laws cannot be known ahead of time. Indeed, they don't exist at all until they emerge, because by definition discoveries cannot be predicted and the laws of behavior depend in part on discoveries that will be made.

Risk analysts use mathematical models to describe the properties and behavior of the systems they analyze. Shackle-Popper indeterminism has important practical consequences for the use of such models. Complexity and high dimensionality are severe challenges in themselves. However, here we are dealing with the limited ability of laws or theories, derived from *past* behavior, to describe *future* behavior. Shackle-Popper indeterminism explains why such laws and theories are fallible and why models—based on laws and theories—are sometimes inadequate. Intelligent learning behavior, as we have defined it, entails an element of spontaneous innovation resulting from discovery. Consequently, models and theories cannot describe or predict all future behavior. In this way Shackle-Popper indeterminism explains the inability of models, even probabilistic ones, to predict structural changes that dominate the history of economic and social behavior. Shackle-Popper indeterminism accounts for the partially un-lawlike and surprising nature of the systems for which risk analysts are responsible. This indeterminism is a central cause of severe uncertainty. We will discuss some practical implications in section 3.

2.3 Hume and the Problem of Induction

Our understanding of the world, and the models and scientific theories based on that understanding, are obtained in part by induction: using evidence to draw new conclusions and to make generalizations and predictions. Why are we justified in accepting or believing the laws of physics or other sciences? The difficulty in conclusively answering that question gives insight into severe uncertainty.

As David Hume explained long ago, one can never demonstrate by deductive reasoning that past patterns will recur in the future. "[W]e cannot give a satisfactory reason why we believe, after a thousand experiments, that a stone will fall or fire burn".[22, p.160]

> For all inferences from experience suppose, as their foundation, that the future will resemble the past and that similar powers will be conjoined with similar sensible qualities. ... It is impossible, therefore, that any arguments from experience can prove this resemblance of the past to the future, since all these arguments are founded on the supposition of that resemblance.[22, p.57]

Furthermore, one cannot prove empirically that past experience is a guide to the future. By the time one tests the regularity of the future, that future has become the past. The future can never be tested, just as one can never step on the rolled up part of an endless rug unfurling always in front of you.[23]

Of course, Hume had no doubt that stones will continue to fall in the future and that fires will continue to burn. He ascribed this belief to the formation of habit, because neither experiment nor reasoning can provide a basis for this belief. Hume did not deal with the problem of how we decide which habitual predictions are warranted and which are not.

That problem has been studied by many philosophers and is very difficult and of great practical importance to risk analysts. Nelson Goodman's emerald example will make the point.[24, p.74]

Suppose that we have examined many emeralds up to time t and found them all to be green. This supports the inductive generalization "All emeralds are green." Now, says Goodman, consider the property "grue", which is "green up to time t and blue thereafter." Our evidence from inspecting emeralds up to time t is consistent with the statement "All emeralds are grue."

Goodman surely does not doubt that emeralds are, and will remain, green. The question is: what rule do we have for concluding from observation that "green" holds and "grue" does not. Both are equally warranted based on noting the color of many emeralds up to time t. It is no help to reply that we know other things, such as stability of color, or chemical composition. We can always construct grue-like hypotheses that are consistent with our knowledge, because the past does not uniquely constrain the future. As Hume wrote "Whatever *is* may *not be*."[22, p.161]

This is not the place to review the many attempts to establish criteria for warrant of inductive predictions. However, such criteria are of great importance to risk analysts, who use past data of many sorts to make predictions. Risk analysis is difficult, in part, because it is difficult to establish the degree of warrant of predictions and generalizations, especially when our contextual knowledge is limited or when the system is in flux.

3 Some Thoughts on How to Proceed

Risk analysis is difficult because we've only just begun to understand our world and its uncertainties (section 2.1), and because indeterminism is an inherent part of the open universe in which we live (section 2.2), and because learning from experience involves untestable assumptions and because warrant for predictions is problematic (section 2.3). Focussing on these challenges of severe uncertainty, and ignoring others (such as institutional and psychological issues), how can we proceed?

3.1 Epistemic Paralysis

We first of all identify the risk of epistemic paralysis, in order to avoid it. As John Locke wrote:[25, I.i.5]

> If we will disbelieve everything, because we cannot certainly know all things; we shall do muchwhat as wisely as he, who would not use his legs, but sit still and perish, because he had no wings to fly.

The absurdity of Locke's wingless gentleman starving in his chair leads us to belief and action, despite our doubts. The moral imperative of the risk analyst's professional responsibility sweeps aside the paralysis of uncertainty.

This has three practical implications. First, one should acquire information and understanding—which we will collectively refer to as "models"—subject to resource constraints. Second, it entails acknowledgement that the models may be of diverse sorts (linguistic, deterministic, probabilistic, etc.), that they may contain conflicting information, and that better models will probably become available in the future or given more resources. These two implications can be combined by saying that the analyst must balance between skepticism and the need to decide and act. Third, it implies the need for tools to manage this balancing act.

3.2 Models and Robustness

Many methods are relevant to managing our ignorance without falling prey to epistemic paralysis. Here we will focus on a generic concept of robustness.

'Robust' means[26] 'Strong and hardy; sturdy; healthy'. By implication, something that is robust is 'not easily damaged or broken, resilient'. A statistical test is robust if it yields 'approximately correct results despite the falsity of certain of the assumptions underlying it' or despite errors in the data.

A decision is robust if its outcome is satisfactory despite large error in the information and understanding that justified or motivated the decision. A decision is robust to the extent that it is resilient to surprise, immune to ignorance. This meaning of robustness has been operationalized in many ways. We have robust statistics,[27] robust control,[28] robust decision making,[29] robust flexibility,[30] robust economics,[31] info-gap robustness,[18] and many more tools.

The precise meaning of robustness differs among these tools. Sometimes it is probabilistic and sometimes not. Some methods are derived from axiomatic formulations and theorems of optimality. Other methods are derived by plausible reasoning from a given state of knowledge. Others are pragmatic and ad hoc. But all robust methods for analysis of risks and for prioritizing decisions attempt to balance between what we know—our models in the broad sense—and what we don't know. They all attempt, in different ways, to achieve acceptable outcomes based on our knowledge and despite our sometimes severe uncertainties. They all attempt to responsibly avoid epistemic paralysis.

Is a robust strategy a good probabilistic bet for achieving acceptable or desired outcomes? The question is important because likelihood of success is desirable. The question is difficult because most concepts of robustness are non-probabilistic. The surprising answer is that even non-probabilistic concepts of robustness often provide the best probabilistic bet. We touch on the explanation here and return to it in section 3.3.

Statistical or probabilistic concepts of robustness enable the evaluation of likelihood of success in one sense of another. Other concepts of robustness are inherently non-probabilistic, usually because the relevant probability distributions are not known. As Wald remarks,(11, p.267) "in most of the applications not even the existence of ... an a priori probability distribution [on a class of distribution functions] ... can be postulated, and in those few cases where the existence of an a priori probability distribution ... may be assumed this distribution is usually unknown."

Nonetheless, robustness to ignorance and surprise is often a successful strategy in competition. There is evidence that info-gap robust strategies are widespread in animal foraging(32) and in financial markets(18 sec.11.5, 33, 34), both of which are situations in which bad bettors would be weeded out by competitors who make better bets. The suggestion is that a robust strategy is a better bet than other strategies, as can be shown formally for a wide range of applications.(35, 36)

Sometimes the analyst knows that other strategies are, in principle, more likely to succeed. Nonetheless, when those strategies are not implementable due to lack of data or knowledge, a robust approach of one kind or another is a defensible and responsible avoidance of epistemic paralysis.

3.3 Robustness and the Innovation Dilemma

We will briefly and qualitatively discuss a robustness analysis of the innovation dilemma.(37)

3.3.1 Problem Formulation

Innovations are attempts to achieve a better future. However, innovations often present a challenging dilemma to risk analysts and decision makers. Many decisions require choosing between options, one of which is both potentially better in the outcome but markedly more uncertain. In these situations the decision maker faces an **"innovation dilemma."**

Innovation dilemmas arise in many contexts of concern to risk analysts. New and innovative technologies are often advocated because of their purported improvements on existing products or methods. However, what is new is usually less well-known and less widely tested than what is old. The range of possible adverse surprises of an innovative technology may exceed the range of surprise for a tried-and-true technology. In public health, for instance, new immunization programs may present policy officials with worries about uncertain side effects. New agricultural technologies promise improved production efficiency or new consumer choices, but with uncertain benefits and costs and potential unanticipated adverse effects resulting from use of manufactured inputs such as fertilizers, pesticides, and machinery, and, more recently, genetically engineered seed varieties.[2] And so on.

Innovation dilemmas are decision problems with three traits: critical needs must be met; the current situation may or may not be adequate; the innovation looks significantly better than current practice but is much more uncertain. To change, or not to change? What strategy to use in making a decision? What choice is the best bet?

3.3.2 Two Solution Strategies

The decision is easy in either of two extreme situations, and their analysis will reveal a general conclusion.

[2] I am indebted to L. Joe Moffitt and Craig Osteen for the agricultural example.

One extreme is that the status quo is clearly insufficient or unsatisfactory. Whether the realm is public health, or agricultural productivity, or homeland security, the current situation is unacceptable and must change.

The other extreme is that the status quo is just fine. Public health is steadily improving, productivity is continually rising, or security is high. No change in policy is justified.

From these two extremes we draw an important general conclusion: the right choice depends on what you need (or think you need, or perceive as the values of the options). To adopt the innovation and to change, or to stay with conventional methods, depends on what outcome is needed or deemed acceptable. There is no universal answer, like, "Always try to improve" or "If it's working, don't fix it". This is a very general property of decisions under uncertainty, and we will call it *"preference reversal"* because of its close structural relation to the psychological phenomenon known by that name.[18 ch.11, 38] One's preference between alternatives depends on what one needs.

The decision strategy that we have just described is attuned to the needs of the decision maker. The strategy attempts to satisfy the agent's critical requirements. If the status quo would reliably do that, then stay put; if not, then change. Simon called this a **satisficing decision strategy:** one that satisfies a critical requirement.[39]

The concept of a "critical requirement" requires some further explanation. In some situations the critical requirement is explicit, as when a regulatory agency requires that the probability of failure be less than 1 per million per year, or that the lifetime exceed 40 years, and so on. In other situations the agent may initially think in terms of maxima. For instance, the owner of a firm may wish to maximize the time to failure of the firm. On closer thought, and given the pressure of limited foresight, the owner might be quite satisfied with confidence in another 40 years of operation, (or maybe, "Let's just get past the current crisis."). A common example of the distinction between maximizing and satisficing is in finance. Stock brokers may advertise their ability to maximize their client's wealth. However, at the end of year, what the broker brags about is having performed better than other brokers. The broker's critical requirement is to beat the competition, and this is a satisficing requirement. The claim of maximization is stronger than what the broker needs to make in order to attract clients.

Now let's consider a different decision strategy that the agent might use. The agent has models—in the broad sense of information and understanding—of two alternatives: the standard alternative and the innovative one. The models may be data sets, or mental models, or deterministic or probabilistic models, or combinations of these or other possibilities. Given these models, the agent predicts which alternative would yield the better outcome. The decision strategy is to choose the alternative whose predicted outcome is best. We will call this decision strategy **outcome optimization.** This strategy uses the models to find the choice that—if the models are correct—will yield the best outcome. Outcome optimization (usually) gives a single "best" decision, unlike the satisficing strategy that returns different answers depending on the agent's needs.

The distinction between outcome optimization and satisficing is not in the mathematics that is used to describe or implement them. They both might be optimizing *something.* The distinction is in the agent's expectations[3] for the outcome. In outcome optimization the agent seeks the best possible result of the decision, where "best" is in the substantive context of the agent's needs or desires. In the satisficing strategy the agent looks for an outcome that is adequately good (which may be very demanding), where "good" is in the same substantive context as "best" for the optimizer. Note that outcome optimization can be thought of as a special case of satisficing: the one where the satisficer's adequate outcome is, in fact, the predicted optimum. But satisficing and outcome-optimizing are different if the satisficer aims at less than the predicted optimum.

There is an attractive logic—and even perhaps a moral imperative—to use one's models to achieve the best outcome. One should always try to do one's best. But, under severe uncertainty, the

[3]The agent needn't be aware of these expectations. The distinction between satisficing and outcome optimization is relevant also for animals or organizations.

catch in the argument for outcome optimization is that the models, even if they quantify uncertainty with probability distributions, may actually be grievously wrong. Outcome optimization ignores the agent's central dilemma under severe uncertainty: stay with the relatively well known but modest alternative, or go for the more promising but more uncertain alternative. Under severe uncertainty, our models, even if probabilistic, may not adequately quantify our ignorance.

In many situations it can rightly be pointed out that the agent's models account probabilistically for uncertainty. Probabilistic models are good to have, and if they are correct then one can realistically optimize probabilistic outcomes. But a probabilistic outcome optimization is simply one type of outcome optimization, and is subject to the same vulnerability to error. The world is full of surprises. Under severe uncertainty, the probability distributions that are used are quite likely wrong, especially in predicting the rare events that the agent is most concerned to avoid or achieve.

When making a decision under severe uncertainty, and when outcomes must satisfy critical requirements, the strategy based on satisficing is usually more reliable than outcome optimization, in a sense that we now explain.

3.3.3 Robustness and Probability

The satisficing strategy might not use probabilistic information. Nonetheless, the satisficing strategy is often a better bet (or at least not a worse bet), probabilistically speaking, than any other strategy, including probabilistic outcome optimization,

When the satisficing decision strategy is the best bet, this is, in part, because it is more robust to uncertainty than any other strategy. A decision is robust to uncertainty if it achieves required outcomes even if adverse surprises occur. In many important situations (though not invariably), as we mentioned in section 3.2, *more robustness* to uncertainty is equivalent to being *more likely to succeed.* When this is true we say that *robustness is a proxy for probability.*

A thorough analysis of the proxy property is rather technical.[(36)] However, we can understand the gist of the idea by considering a simple special case.

Suppose we are completely confident about the future value of not making any change (not adopting the innovation). In contrast, the future value of changing is apparently better though uncertain. If not changing would satisfy our critical requirement, then we are absolutely certain of satisfying our requirements if we do not change. Not changing is completely robust to surprises so the probability of success equals 1 if we do not change, regardless of what happens with the other option. Likewise, if not changing would *not* satisfy our critical requirements, then we are absolutely certain of failure if we do not change; the probability of success equals 0 if we do not change, and changing cannot be worse. Regardless of what probability distribution describes future outcomes if we change, we can always choose the option whose likelihood of success is greater (or at least not worse). This is because not changing is either sure to succeed or sure to fail, and we know which.

This argument can be extended to the more realistic case where the outcome of rejecting the innovation is uncertain and the outcome of adopting the innovation, while seemingly better than not changing, is much more uncertain. The agent can know which option is more robust to uncertainty, without having to know probability distributions. This implies, in many situations, that the agent can choose the option that is a better bet for satisfying the outcome requirements.

3.3.4 Summing Up: Robust-Satisficing and the Innovation Dilemma

An innovation dilemma arises when one must choose between a seemingly better (innovative) option that is more uncertain, and a more thoroughly understood but possibly less attractive (tried and true) option. The first step in balancing between uncertainty and decisiveness is to identify critical or necessary outcomes. Then, in a range of situations, a satisficing decision strategy that is maximally robust to our ignorance is a better bet than other strategies, for instance outcome optimization, as we have explained above.

3.4 Robustness and Worst Cases: Two Approaches

There are many types of risk analysis partly because ignorance and uncertainty come in many forms. Probabilistic uncertainty induces probabilistic risk analysis, while starker uncertainty—for instance ignorance of relevant probability distributions—engenders other analyses of risk.

A widely occurring operational distinction between risk analyses hinges on whether or not meaningful worst cases can be identified. When one can plausibly specify the worst events that can occur (and presuming we don't know probability distributions), then one might justifiably try to ameliorate these worst contingencies. This can be done in many different ways, and we will refer to this type of strategy as min-max analysis: minimizing the maximum damage.

The ability to implement a min-max analysis depends on identifying meaningful worst contingencies. This is feasible in many situations. The concept of a "meaningful worst case" depends on knowledge and judgment that may be within the risk analyst's competence. However, it is not usually sufficient to specify a worst case in some formal or abstract sense, such as the set of all contingencies that are consistent with the laws of science. A min-max analysis based on such an inclusive formulation may be uselessly over conservative. Min-max analysis is most useful when the analyst is able to avoid vacuous specification of worst cases. However, when information is really scarce, for instance when processes are poorly understood or changing, then even typical cases cannot be reliably identified. It may then be impossible to meaningfully specify the boundary between extreme but possible occurrences, and the impossible or negligible.

Nonetheless, even when worst cases cannot be meaningfully specified, the analyst still has data, understanding, and mathematical representations: models in the broad sense that we are using that term. It is simply that the analyst cannot responsibly specify the magnitude of error of these models. For instance, we have many models for long-range climate change, but the earnest scientific disputes over these models may preclude the ability to confidently bound the errors. Or, introducing a new species to an ecosystem, either deliberately as an genetically modified organism or inadvertently by invasion, may alter the ecosystem dynamics in unknown ways.

In such situations one can still formulate and implement a robustness analysis. Info-gap theory has been developed precisely for the task. Let's discuss min-max and info-gap concepts of robustness.

The min-max concept of robustness responds to the question: how bad is the worst case? This is valuable information for the risk analyst and decision maker because if the worst case—after amelioration by a min-max analysis—is tolerable, then one can reasonably say that the system is robust to uncertainty.

The info-gap concept of robustness responds to a different question: how wrong can the models be and still guarantee that the outcome is acceptable? This is useful for the risk analyst and decision maker because if the models can err enormously without preventing acceptable outcomes, then one can reasonably say that the system is robust to uncertainty.

These two concepts of robustness—min-max and info-gap—are different, motivated by different information available to the analyst. The min-max concept responds to severe uncertainty that nonetheless can be bounded. The info-gap concept responds to severe uncertainty that is unbounded or whose bound is unknown. It is not surprising that min-max and info-gap robustness analyses sometimes agree on their policy recommendations, and sometimes disagree, as has been discussed elsewhere.(40)

It should be pointed out that Sniedovich's tendentious discussion(41) of "voodoo decision making" and "No Man's Land" misses the point. Info-gap theory, like all theories of robustness, starts with the analyst's models, and asks: how much error in these models can be tolerated? The info-gap robustness question—how wrong can one's models be and yet the decision still yields an acceptable outcome—is pertinent when maximum error is unknown. If the robustness is large, (and this is a judgment that the analyst must make, like other judgments made by risk analysts) then one may have confidence in the decision. If the robustness is not large, and especially if the robustness is small, then confidence is not warranted. If the robustness is small then confidence is warranted only "locally", near the models, while if the robustness is large then confidence is warranted over a wide domain of deviation from the models. Info-gap theory uses the analyst's models, but this does not make it a "local" theory of robustness.

Info-gap theory provides a range of distinctive mathematical models for non-probabilistically representing uncertain information of many sorts. Scholars are just barely launched on the study of uncertainty, as witnessed by the florescence of uncertainty models in the past century.[(42)] If we remain dispassionate and abjure demagoguery, then our mastery of the unknown will continue to grow. See Snow's comments in a similar context.[(43, pp.56–58)]

3.5 The Other Side of the Coin: Opportuneness

Risk analysts participate in preventing high-consequence adverse events in critical technologies. However, uncertainty is not necessarily pernicious, and may even be propitious. Analysts of risk should at least glance in the direction of potential favorable opportunities. Info-gap theory provides a method for doing this, as we now explain.

The analyst has models for anticipating the outcome of an action. A surprise is favorable if its consequence is better than the anticipated outcome. An action or decision is opportune if it can facilitate or exploit favorable surprises.

The info-gap concept of opportuneness[(18)] responds to the question: how wrong must the models be in order for outcomes better than expected (and perhaps even wonderful) to be possible? If the models need to err only slightly in order to enable wonderful outcomes, then the decision is opportune for uncertainty.

The opportuneness question is the converse of the robustness question discussed in section 3.4: how wrong can the models be, and the outcome is still guaranteed to be acceptable?

In some domains, such as some areas of finance, risk analysts focus primarily on favorable opportunities. In other domains, such as health or technological safety, risk analysts focus primarily on robustness against unacceptable outcomes, rather than on opportuneness for wonderful outcomes. Nonetheless, the info-gap opportuneness analysis can support a robustness analysis in three ways.

First, consider the choice between two options that have equal or similar robustness against failure. In this case, the opportuneness analysis can break the tie.

Second, robustness and opportuneness are not necessarily antagonistic. There are many situations in which a change in the decision that augments the robustness also augments the opportuneness. Awareness of this possibility is relevant to a comprehensive analysis of the decision.

Third, consider a situation where robustness and opportuneness are antagonistic: improving one of them causes deterioration of the other. The risk analyst may demand large robustness against pernicious uncertainty. Now suppose the maximum robustness is quite large. Variation of the decision around the robust maximum will result in small change in robustness (if the maximum is a stationary point). But such variation of the decision can lead to substantial improvement in the opportuneness (if its optimum is separate from the robust optimum). This means that a small amount of robustness can, sometimes, be traded for substantial improvement in the ability to exploit opportune surprises. The risk analyst may wish to consider this trade off.

4 Conclusion

Many challenges face the risk analyst, and in this essay we have touched on some that arise from the limitation of knowledge and the richness of the unknown. We have discussed some tools for the analysis and management of severe uncertainty, focussing on concepts of robustness.

Within the context of the epistemological considerations raised here, a great pitfall to be avoided by the risk analyst is rigid adherence to any one methodology. The risk analyst is butting against the endless diversity of the unknown. Sometimes one conceptual approach is appropriate, and sometimes another. The need for methodological pluralism has implications for both education and legislation.

Another great pitfall to avoid is the exclusive devotion to scientific analysis. Science is the basis of our analysis, but "it is anything but a pure scientific world in which [we] live."[44, p.88] The risk analyst must understand the historical, philosophical, social and economic contexts and adapt the scientific analysis to them. Concepts and modes of thought from the humanistic and social disciplines are invaluable in this enterprize.

Acknowledgement

The author is indebted to Louis Anthony (Tony) Cox, jr. and to Mark A. Burgman for valuable comments and suggestions.

5 References

1. Heath T. A History of Greek Mathematics. Vol.1: From Thales to Euclid. Vol.2: From Aristarchus to Diophantus. Oxford: Clarendon Press. 1921, New York: Dover reprint, 1981.

2. Klein J. Greek Mathematical Thought and the Origin of Algebra. New York: Dover Publications, 1992.

3. Franklin J. The Science of Conjecture: Evidence and Probability before Pascal. Baltimore: The Johns Hopkins University Press, 2001.

4. Hacking I. The Emergence of Probability: A Philosophical Study of Early Ideas About Probability, Induction and Statistical Inference. Cambridge: Cambridge University Press, 1975.

5. Bayes T. An essay towards solving a problem in the doctrine of chances. Phil. Trans. Royal Society London, Dec 23, 1763, 53: 370–418.

6. Porter TM. The Rise of Statistical Thinking: 1820–1900, Princeton: Princeton University Press, 1986.

7. Stigler SM. The History of Statistics: The Measurement of Uncertainty before 1900. Boston: The Belknap Press of Harvard University Press, 1986.

8. Kolmogorov AN. Grundbegriffe der Wahrscheinlichkeitsrechnung. Berlin: Springer, 1933.

9. Lejewski C, Lukasiewicz J. in Edwards, P., ed., The Encyclopedia of Philosophy. New York: Simon and Schuster Macmillan, vol. 5, 1996.

10. Zadeh LA. Fuzzy sets. Information and Control, 1965; 8: 338–353.

11. Wald A. Statistical decision functions which minimize the maximum risk, Annals of Mathematics, 1945; 46(2): 265–280.

12. Ferson S. RAMAS Risk Calc 4.0 Software: Risk Assessment with Uncertain Numbers. Boca Raton: Lewis Publishers, 2002.

13. Walley P. Measures of uncertainty in expert systems, Artificial Intelligence, 1996; vol.83: 1–58.

14. Miranda E. A survey of the theory of coherent lower previsions. Intl J Approximate Reasoning, 2008; 48: 628–658.

15. Dempster AP. Upper and lower probabilities induced by multivalue mappings. Annals of Mathematical Statistics, 1967; 38: 325–339.

16. Shafer G. A Mathematical Theory of Evidence. Princeton: Princeton University Press, 1976.

17. Klir GJ. Uncertainty and Information: Foundations of Generalized Information Theory, Hoboken, NJ: Wiley, 2006.

18. Ben-Haim Y. 2006, Info-gap Decision Theory: Decisions Under Severe Uncertainty. 2nd ed., London: Academic Press, 2006.

19. Shackle GLS. Epistemics and Economics: A Critique of Economic Doctrines. London: Transaction Publishers, 1992, originally published by Cambridge University Press, 1972.

20. Popper K. The Open Universe: An Argument for Indeterminism. From the Postscript to The Logic of Scientific Discovery. London: Routledge 1982.

21. Ben-Haim Y. Peirce, Haack and Info-gaps, in Susan Haack, A Lady of Distinctions: The Philosopher Responds to Her Critics. Cornelis de Waal, ed., Amherst, NY: Prometheus Books, 2007.

22. Hume D. An Inquiry Concerning Human Understanding. 1748, edited by Antony Flew. New York: Collier Books, 1962.

23. Ben-Haim Y. The end of science? http://decisions-and-info-gaps.blogspot.com/2011/10/end-of-science.html. 2011.

24. Goodman N. Fact, Fiction, and Forecast. 4th ed., Boston: Harvard University Press, 1983.

25. Locke J. An Essay Concerning Human Understanding. 5th edition, 1706. Roger Woolhouse, editor. London: Penquin Books, 1997.

26. Oxford English Dictionary. 2012. Online version.

27. Huber PJ. Robust Statistics. New York: John Wiley, 1981.

28. Zhou K, Doyle JC. Essentials of Robust Control. Upper Saddle River, NJ: Prentice Hall, 1997.

29. Lempert RJ, Popper SW, Bankes SC. Shaping the Next 100 Years: New Methods for Quantitative, Long-Term Policy Analysis, Santa Monica: RAND Corp., 2003.

30. Rosenhead J. Robustness analysis: Keeping your options open, in Jonathan Rosenhead, ed. Rational Analysis For a Problematic World: Problem Structuring Methods For Complexity, Uncertainty and Conflict. New York: John Wiley, 1989.

31. Hansen LP, Sargent TJ. Robustness, 2008. Princeton: Princeton University Press 2008.

32. Carmel Y, Ben-Haim Y. Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice? American Naturalist, 2005; 166: 633–641.

33. Ben-Haim Y, Jeske K. Home-bias in financial markets: Robust satisficing with info-gaps, Federal Reserve Bank of Atlanta, Working Paper Series, 2003–35, Dec. 2003.

34. Ben-Haim Y. Doing Our Best: Optimization and the Management of Risk, Risk Analysis, to appear.

35. Ben-Haim Y. 2011, When is non-probabilistic robustness a good probabilistic bet? http://www.technion.ac.il/yakov/prx27a.pdf

36. Ben-Haim Y. Robust satisficing and the probability of survival, Intl J System Science, to appear.

37. Ben-Haim Y. The innovation dilemma, http://decisions-and-info-gaps.blogspot.com/2011/08/innovation-dilemma.html

38. Tversky A, Thaler RH. Anomalies: Preference Reversals, J Economic Perspectives, 1990; 4(2): 201–211.

39. Simon HA. Rational choice and the structure of the environment, Psychological Review, 1956; 63(2), 129–138.

40. Ben-Haim Y, Dacso CC, Carrasco J, Rajan N. Heterogeneous uncertainties in cholesterol management, Intl J Approximate Reasoning, 2009; 50: 1046–1065.

41. Sniedovich M. Fooled by local robustness, Risk Analysis, 2012; DOI: 10.1111/j.1539-6924.2011.01772.x

42. Ben-Haim Y. We're just getting started: A glimpse at the history of uncertainty, 2012; http://decisions-and-info-gaps.blogspot.com/2012/03/were-just-getting-started-glimpse-at.html

43. Snow CP. The Two Cultures: And A Second Look. Cambridge: Cambridge University Press, 1964.

44. Adams H. The Education of Henry Adams, edited by Ernest Samuels, Boston: Houghton Mifflin Co., 1918.