Chapter 5
Info-Gap Decision Theory (IG)

Yakov Ben-Haim

Abstract

- Info-Gap (IG) Decision Theory is a method for prioritizing alternatives and making choices and decisions under deep uncertainty.
- An "info-gap" is the disparity between what *is known* and what *needs to be known* for a responsible decision.
- Info-gap analysis does not presume knowledge of a worst-case or of reliable probability distributions.
- Info-gap models of uncertainty represent uncertainty in parameters and in the shapes of functional relationships.
- IG Decision Theory offers two decision concepts: *robustness* and *opportuneness*.
- The *robustness* of an alternative is the greatest horizon of uncertainty up to which that alternative satisfies critical outcome requirements.
- The *robustness strategy* satisfices the outcome and maximizes the immunity to error or surprise. This differs from outcome optimization.
- The *robustness function* demonstrates the trade-off between immunity to error and quality of outcome. It shows that knowledge-based predicted outcomes have no robustness to uncertainty in that knowledge.
- The *opportuneness* of a decision alternative is the lowest horizon of uncertainty at which that decision enables better-than-anticipated outcomes.
- The *opportuneness strategy* seeks windfalls at minimal uncertainty.
- We discuss "innovation dilemmas" in which the decisionmaker must choose between two alternatives, where one is putatively better but more uncertain than the other.
- Two examples of info-gap analysis are presented, one quantitative that uses mathematics and one qualitative that uses only verbal analysis.

Y. Ben-Haim (✉)
Technion—Israel Institute of Technology, Haifa, Israel
e-mail: yakov@technion.ac.il

5.1   Info-Gap Theory: A First Look

Info-Gap (IG) is a non-probabilistic decision theory for prioritizing alternatives and making choices and decisions under deep uncertainty (Ben-Haim 2006, 2010). They might be operational alternatives (design a system, choose a budget, decide to launch or not, etc.) or more abstract alternatives (choose a model structure, make a forecast, formulate a policy, etc.). Decisions are based on data, scientific theories, empirical relations, knowledge, and contextual understanding, all of which I'll refer to as one's *models*, and these models often recognize and quantify uncertainty.

IG theory has been applied to decision problems in many fields, including various areas of engineering (Kanno and Takewaki 2006; Chinnappen-Rimer and Hancke 2011; Harp and Vesselinov 2013), biological conservation (Burgman 2005), economics (Knoke 2008; Ben-Haim 2010), medicine (Ben-Haim et al. 2012), homeland security (Moffitt et al. 2005), public policy (Hall et al. 2012), and more (see www.info-gap.com). IG robust satisficing has been discussed non-mathematically elsewhere (Schwartz et al. 2011; Ben-Haim 2012a, b, 2018; Smithson and Ben-Haim 2015).

Uncertainty is often modeled with probability distributions. If the probability distributions are correct and comprehensive, then one can exploit the models exhaustively to reliably achieve stochastically optimal outcomes, and one doesn't need IG theory. However, if one's models will be much better next year when new knowledge has become available (but you must decide now), or if processes are changing in poorly known ways, or if important factors will be determined beyond your knowledge or control, then one faces deep uncertainty and IG theory might help. This section presents an intuitive discussion of two basic ideas of IG theory: satisficing and robustness. A more systematic discussion of IG robustness appears in Sect. 5.2. Simple examples are presented in Sects. 5.3 and 5.4, and more detailed examples appear in Chap. 10.

Knight (1921) distinguished between what he called "risk" (for which probability distributions are known) and "true uncertainty" (for which probability distributions are not known). Knightian ("true") uncertainty reflects ignorance of many things, including underlying processes, functional relationships, strategies or intentions of relevant actors, future events, inventions, discoveries, surprises, and so on. Info-gap models of uncertainty provide a non-probabilistic quantification of Knightian uncertainty. An info-gap is the disparity between what you *do know* (or think to be true) and what you *need to know* for making a reliable or responsible decision (though what is needed may be uncertain). An info-gap is not ignorance per se, but rather those aspects of one's Knightian uncertainty that bear on a pending decision and the quality of its outcome.

An info-gap model of uncertainty is particularly suitable for representing uncertainty in the shape of a function. For instance, one might have an estimate of the stress–strain curve for forces acting on a metal, or of the supply and demand curves for a new product, or of a probability density function (pdf), but the shape of the function (e.g., the shape of the elastic–plastic transition curve in the first case, or the

shapes of the supply and demand curves, or the tails of the pdf) may be highly uncertain. Info-gap models are also widely used to represent uncertainty in parameters or vectors, or sometimes uncertainty in sets of such entities.

Decisionmakers often try to optimize the outcome of their decisions. That is usually approached by using one's best models to predict the outcomes of the various alternatives, and then choosing the one whose predicted outcome is best. The aspiration for excellence is commendable, but outcome optimization may be costly, or one may not really need the best possible outcome. Schwartz (2004) discusses the irrelevance of optimal outcomes in many situations.

Outcome optimization—using one's models to choose the decision whose predicted outcome is best—works fine when the models are pretty good, because exhaustively exploiting good models will usually lead to good outcomes.

However, when one faces major info-gaps one's models contain major errors or lacunae, and exhaustively exploiting the models can be unrealistic, unreliable, and can lead to undesired outcomes (Ben-Haim 2012a). Under deep uncertainty, it is better to ask: What outcomes are critical and must be achieved? This is the idea of *satisficing* introduced by Simon (1956): achieving a satisfactory or acceptable, but not necessarily optimal, outcome—"good enough" according to an explicitly stated set of criteria.

Planners, designers, and decisionmakers in all fields have used the language of optimization (the lightest, the strongest, the fastest) for ages. In practice, however, satisficing is very widespread, although not always recognized as such. Engineers satisfy design specifications (light enough, strong enough, fast enough). Stock brokers, business people, and investors of all sorts do not really need to maximize profits; they only need to beat the competition, or improve on last year, or meet the customer's demands. Beating the competition means satisficing a goal. The same can be said of the public official who must reduce crime below a legislative target, or the foreign aid planner who must raise public health to international standards.

Once the decisionmaker identifies the critical goals or outcomes that must be achieved, the next step is to decide about something or choose an action that will achieve those goals despite current ignorance or unknown future surprises. A decision has high *robustness* if it satisfices the performance requirements over a wide range of unanticipated contingencies. Conversely, a decision has low robustness if even small errors in our knowledge can prevent achievement of the critical goals. The *robust-satisficing* decisionmaker prioritizes the alternatives in terms of their robustness against uncertainty for achieving the critical goals. The decision methodology of IG robust satisficing is often motivated by the pernicious potential of the unknown. However, uncertainty can be propitious, and IG theory offers a method for prioritizing one's alternatives with respect to the potential for favorable surprises. The idea of "windfalling" supplements the concept of satisficing. The opportune windfalling decisionmaker prioritizes the alternatives in terms of their potential for exploiting favorable contingencies. This is illustrated in the example in Sect. 5.4.

Min-max or worst-case analysis is a widely used alternative to outcome optimization when facing deep uncertainty and bears some similarity to IG robustness. Neither min-max nor IG presumes knowledge of probabilities. The basic approach

behind these two methods is to find decisions that are robust to a range of different contingencies. Wald (1947) presented the modern formulation of min-max, and it has been applied in many areas (e.g., Hansen and Sargent 2008). The decisionmaker considers a bounded family of possible models, without assigning probabilities to their occurrence. One then identifies the model in that set which, if true, would result in a worse outcome than any other model in the family. A decision is made that minimizes this maximally bad outcome (hence "min-max"). Min-max is attractive because it attempts to insure against the worst anticipated outcome. However, min-max has been criticized for two main reasons. First, it may be unnecessarily costly to assume the worst case. Second, the worst usually happens rarely and therefore is poorly understood. It is unreliable (and perhaps even irresponsible) to focus the decision analysis on a poorly known event (Sims 2001).

Min-max and IG methods both deal with Knightian uncertainty, but in different ways. The min-max approach is to choose the decision for which the contingency with the worst possible outcome is as benign as possible: Identify and ameliorate the worst case. The IG robust-satisficing approach requires the planner to think in terms of the worst consequence that can be tolerated, and to choose the decision whose outcome is no worse than this, over the widest possible range of contingencies. Min-max and IG both require a prior judgment by the planner: Identify a worst model or contingency (min-max) or specify a worst tolerable outcome (IG). These prior judgments are different, and the corresponding policy selections may, or may not, agree. Ben-Haim et al. (2009) compare min-max and IG further.

5.2 IG Robustness: Methodological Outline

This section is a systematic description of the IG robustness methodology for decisionmaking under deep uncertainty. The following two sections present examples.

IG robustness is the attribute of satisfying critical requirements even when the situation is, or evolves, differently from what is expected. A decision is robust to uncertainty if it remains acceptable even if the understanding that was originally available turns out to be substantially wrong. The robustness is assessed by answering the question: How wrong can our current model be so that the outcome of this decision will still be acceptable? In other words, how immune to our current ignorance is this decision? A decision is highly robust if it remains acceptable throughout a wide range of deviation of reality from the original understanding. More robustness to uncertainty is better than less, so decisionmakers prefer the more robust decision over one with less robustness to uncertainty. We will now formulate these ideas more rigorously. Sections 5.3 and 5.4 present two simple examples.

5.2.1   Three Components of IG Robust Satisficing

The IG analysis of robustness to uncertainty rests on three components: the model, the performance requirements, and the uncertainty model. These components constitute the framing of the analysis. Furthermore, this framing may be done in a quick exploratory mode, or may be more fundamental and advanced. Likewise, decisions based on the analysis may be short-term actions, long-term options, or adaptive strategies that combine the short and long term.

The *model*, as defined earlier, is our understanding of the system or situation that must be influenced, its temporal dynamics, the evidence, the environment, and any other available relevant knowledge.

The *performance requirements* are specified in response to the question: What do we need to achieve in order for the outcome of the decision to be acceptable? The economist may require inflation within specified bounds; the engineer may require an operational lifetime exceeding a given value; the military commander may require a substantial decrease in insurgent violence, etc. The analyst together with the decisionmaker makes judgments of the required or acceptable levels of performance. More demanding performance entails greater vulnerability to uncertainty and hence lower levels of robustness, as we will see.

Uncertainty model. We have best estimates of our model: the knowledge, understanding, and evidence relating to the situation. However, these estimates may be wrong or incomplete. We may also be uncertain about the performance requirements (how much inflation or insurgent violence is acceptable?). There are many specific forms of info-gap models of uncertainty, encoding different information about the uncertainty. However, they all express the intuition that we do not know how wrong our best estimates are. The info-gap model also includes insights and contextual understanding about this uncertainty. An info-gap model of uncertainty expresses what we do know, as well as the unbounded horizon of uncertainty surrounding our knowledge. It expresses the idea that we cannot confidently identify a realistic worst case.

5.2.2   IG Robustness

The analyst must formulate and prioritize candidate decisions and will do so based on their robustness against uncertainty for achieving acceptable outcomes: More robustness is preferred over less robustness. We evaluate the robustness to uncertainty of any candidate decision by combining these three components: the model, the performance requirements, and the uncertainty model. We do this by addressing two questions, the first regarding the putative performance; the second addressing the robustness against uncertainty.

Putative performance. The putative performance of a decision is the outcome that is predicted by the best understanding and the evidence in hand (the model). Given

a candidate decision, we ask if its putative performance satisfies our performance requirements. That is, we ask whether or not this outcome is acceptable (according to the performance requirements) as assessed by the best understanding that we have (our model). If the answer is negative, then we reject this decision alternative. At this stage, we are ignoring uncertainty.

Robustness. Given a candidate decision that has gotten a positive answer to the question regarding putative performance, we now ask: How much could the model change without violating the performance requirements? This explicitly addresses the uncertainty. We are asking: What is the greatest horizon of uncertainty (in the uncertainty model) up to which the performance requirements are guaranteed by this decision? How much could reality deviate from our understanding and evidence (our model) so that the decision we are contemplating would still satisfy the requirements? Could the decision tolerate *any* error in the model, up to some large degree, without violating the performance requirements (implying large robustness)? Or is there *some* small error in the model that would jeopardize the requirements (implying low robustness)?

5.2.3   Prioritization of Competing Decisions

More robustness is better than less robustness. This means that, given two alternative decisions that both putatively satisfy the performance requirements, we prefer the decision that satisfies the performance requirements throughout a larger range of uncertainty. This prioritization is called *robust satisficing*, because it selects the decision that is more robust against uncertainty while also satisfying the performance requirements.

Note that the robust-satisficing decision optimizes the robustness against uncertainty rather than optimizing the substantive quality of the decision's outcome. We do not optimize the outcome. The outcome must be satisfactory, though the analyst can choose the satisficing level to be more or less demanding. We optimize the robustness against uncertainty, and we satisfice the outcome.

5.2.4   How to Evaluate Robustness: Qualitative or Quantitative?

Many decisions under uncertainty are amenable to quantitative analysis (i.e., using mathematics). Many other situations depend on conceptual models and verbal formulations that cannot be captured with equations. IG theory has been applied in both qualitative and quantitative analyses, as we illustrate in Sects. 5.3 and 5.4, respectively. These examples are independent and one can read either or both.

5.3 IG Robustness: A Qualitative Example

We begin with a simple math-free qualitative example. In this section, we examine five conceptual proxies for the concept of robustness (Ben-Haim and Demertzis 2016) and then discuss a simple example based on qualitative reasoning.

5.3.1 Five Conceptual Proxies for Robustness

Like all words, "robustness" has many connotations, and its meanings overlap with the meanings of other words.[1] We discuss five concepts that overlap significantly with the concept of robustness against uncertainty, and that are useful in the qualitative assessment of decisions under uncertainty. Each of these five concepts emphasizes a different aspect of robustness, although they also overlap. The *five proxies for robustness* are resilience, redundancy, flexibility, adaptiveness, and comprehensiveness. A decision, policy, action, or system is highly robust against uncertainty if it is strong in some or all of these attributes; it has low robustness if it is weak in all of them. We will subsequently use the term "system" quite broadly, to refer to a physical or organizational system, a policy for formulating or implementing decisions, or a procedure for political foresight or clinical diagnosis, etc.

Resilience of a system is the attribute of rapid recovery of critical functions. Adverse surprise is likely when facing deep uncertainty. A system is robust against uncertainty if it can rapidly recover from adverse surprise and achieve critical outcomes.

Redundancy of a system is the attribute of providing multiple alternative solutions. Robustness to surprise can be achieved by having alternative responses available.

Flexibility (sometimes called agility) of a system is the ability for rapid modification of tools and methods. Flexibility or agility, as opposed to stodginess, is often useful in recovering from surprise. A physical or organizational system, a policy, or a decision procedure is robust to surprise if it can be modified in real time.

Adaptiveness of a system is the ability to adjust goals and methods in the mid- to long-term. A system is robust if it can be adjusted as information and understanding change. Managing Knightian uncertainty is rarely a once-through procedure. We often must re-evaluate and revise assessments and decisions. The emphasis is on the longer time range, as distinct from on-the-spot flexibility.

Comprehensiveness of a system is its interdisciplinary system-wide coherence. A system is robust if it integrates relevant considerations from technology, organizational structure and capabilities, cultural attitudes and beliefs, historical context, economic mechanisms and forces, or other factors. A robust system addresses the multi-faceted nature of the problem.

---

[1]The representation of knowledge with words is fraught with info-gaps (Ben-Haim 2006, Sect. 13.2).

5.3.2 Simple Qualitative Example: Nuclear Weapon Safety

Nuclear weapons play a role in the national security strategy of several countries today. Like all munitions, the designer must assure effectiveness (devastating explosion during wartime use) together with safety (no explosion during storage, transport, or abnormal accident conditions). The "always/never" dilemma is "the need for a nuclear weapon to be safe and the need for it to be reliable …. A safety mechanism that made a bomb less likely to explode during an accident could also, during wartime, render it more likely to be a dud. … Ideally, a nuclear weapon would always detonate when it was supposed to—and never detonate when it wasn't supposed to" (Schlosser 2013, pp. 173–174). There are many quantitative methods for assessing effectiveness, safety, and the balance between them, but there remains a great need for human judgment based on experience. We briefly illustrate the relevance of the five qualitative proxies for assessing and achieving robustness to uncertainty.

Nuclear weapon safety is assured, in part, by the requirement for numerous independent actions to arm and detonate the weapon. Safety pins must be removed, secret codes must be entered, multiple activation keys controlled by different individuals must be inserted and turned, etc. This *redundancy* of safety features is a powerful concept for assuring the safety of weapon systems. On the other hand, the wartime detonation of the weapon is prevented if any of these numerous redundant safety features gets stuck and fails to activate the device. Redundancy for safety is a primary source of the "always/never" dilemma.

Resilience of the weapon system is the ability to recover critical functions—detonation during wartime in the present example—when failure occurs. For example, resilience could entail the ability to override safety features that fail in the locked state in certain well-defined circumstances. This override capability may be based on a voting system of redundant safety features, or on human intervention, or on other functions. The robustness to uncertainty is augmented by redundant safety features together with a resilient ability to countervail those safety features in well-defined situations where safety features have failed in the locked mode.

Sometimes, the critical function of a system is not a physical act, like detonation, but rather the act of deciding. A command and control hierarchy, like those controlling nuclear weapon use, needs to respond effectively to adverse surprise. The decision to initiate the use of nuclear weapons in democratic countries is usually vested exclusively in the highest civilian executive authority. A concern here is that a surprise "decapitation" strike against that civilian authority could leave the country without a nuclear-response capability. The decisionmaking hierarchy needs *flexibility* against such a surprise: the ability to exercise the critical function of deciding to use (or not to use) nuclear weapons after a decapitating first strike by an adversary. Flexibility could be attained by a clearly defined line of succession after incapacitation of the chief executive, together with both physical separation between the successors and reliable communication among them. It is no simple matter to achieve this finely balanced combination of succession, separation, and communication. The concept

of flexibility assists in assessing alternative implementations in terms of the resulting robustness against uncertainty in hierarchical decisionmaking.

Hierarchical decisionmaking needs to be *adaptive* in response to changing circumstances in the mid- to long-term. For example, the US line of Presidential succession is specified in the US Constitution, and this specification has been altered by amendment and clarified by legislation repeatedly over time to reflect new capabilities and challenges.

The *comprehensiveness* of a decision is the interdisciplinary scope of the mechanisms and interactions that it accounts for and the implications it identifies. The uncertainties regarding nuclear weapons are huge, because many mechanisms, interactions, and implications are unknown or poorly understood. This means that the potential for adverse surprise is quite large. Comprehensiveness of the decision analysis is essential in establishing robustness against uncertainty. Thinking "outside the box" is a quintessential component in achieving comprehensiveness, and human qualitative judgment is of foremost importance here.

We now use these ideas to schematically prioritize two alternative hypothetical strategies for supervising nuclear weapons in a liberal democracy such as the USA, based on the proxies for robustness.

The first strategy is based on current state-of-the-art (SotA) technologies, and authority is vested in the President as commander in chief. Diverse mechanisms assure redundancy of safety features as well as resilience and flexibility of the systems to assure only operational detonation, and adaptability in response to changes over longer times. The system of controls is comprehensive, but primarily based on human observation, communication, and decision.

The second strategy is new and innovative (NaI) and extensively exploits automated sensor- and computer-based access, authentication, communication, control, and decision. The strategy employs "big data" and artificial intelligence in assessing threats and evaluating risks. Humans are still in the loop, but their involvement is supported to a far greater extent by new and innovative technologies.

Our best understanding of these strategies—SotA and NaI—predicts that the second strategy would provide better safety and operability. However, deep uncertainties surround both strategies, and more so for the innovative second strategy because of its newness. The *innovation dilemma* is that the putatively preferable innovative alternative is more uncertain, and hence potentially worse, than the standard alternative. Two properties of robustness assist in resolving this dilemma: zeroing and trade-off.

The *zeroing* property of an IG robustness assessment states that the predicted performances have no robustness to uncertainty. This is because even small errors or lacunae in the knowledge or understanding (upon which predictions are based) could result in outcomes that are worse than predicted. Hence, prioritizing the strategies based on the predictions is unreliable and irresponsible. We must ask what degrees of safety and operability are essential for acceptable performance. That is, we satisfice the performance, rather than trying to optimize it.

We then note that more demanding performance requirements can fail in more ways and thus are more vulnerable to uncertainty. This implies a *trade-off* between performance and robustness to uncertainty: Greater robustness is obtained only by

accepting more modest performance. In high-consequence systems, such as nuclear weapons, the performance requirements are very demanding. Nonetheless, the trade-off is irrevocable and it is wishful thinking to ignore it. The robustness of each strategy is assessed by its strength in the conceptual proxies. The *robust-satisficing* preference is for the strategy that satisfies the performance requirements at greater robustness.

Suppose that the predicted performance of the SotA strategy only barely satisfies the performance requirements. The proxies for robustness of the SotA will have low strength, because small errors can jeopardize the adequacy of the performance. This may in fact have motivated the search for the NaI strategy whose predicted performance exceeds the requirements. In this case, the robust preference will be for NaI, although consideration must be given to the strength of its proxies for robustness. If the proxies for robustness of NaI are also weak, then neither alternative may be acceptable.

Alternatively, suppose that the SotA satisfies the performance requirements by a wide margin. Its proxies for robustness will be strong and probably stronger than for NaI. In this case, the robust preference is for SotA, illustrating the potential for a *reversal of preference* between the strategies: NaI is putatively preferred (based on predicted outcomes), but SotA is more robust, and hence SotA is preferred (based on robustness) over NaI. This emphasizes the difference between robustly satisficing the performance requirements (which leads to either SotA or NaI, depending on the requirements) as distinct from prioritizing based on predicted outcomes (which leads to the putatively better alternative, NaI).

5.4 IG Robustness and Opportuneness: A Quantitative Example

We have discussed qualitative concepts for assessing deep uncertainty and for prioritizing the options facing a decisionmaker: trade-off between robustness and performance requirements, zero robustness of predicted outcomes, innovation dilemmas, and preference reversals. These concepts are embodied in mathematical theorems of IG theory, and they have quantitative realizations, as we illustrate in this section with a simple mechanical example. The one-dimensional linear model of a gap-closing electrostatic actuator with uncertainty in a single parameter stands in for other systems, often significantly more complex and uncertain in their models and predictions.

The nonlinear force–displacement relation for the gap-closing electrostatic actuator (a type of electric switch) in Fig. 5.1 is fairly well represented by:

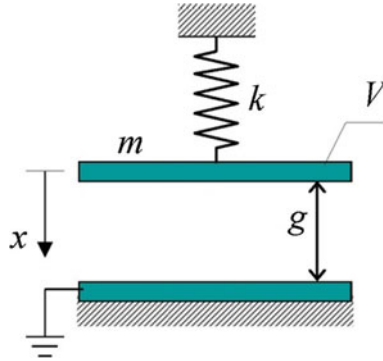$$F = kx - \frac{\varepsilon A V^2}{2(g-x)^2} \qquad (5.1)$$

Fig. 5.1 Gap-closing electrostatic actuator. The figure is reproduced here with the permission of Prof. David Elata, head, Mechanical Engineering Micro Systems (MEMS) lab, Technion—Israel Institute of Technology
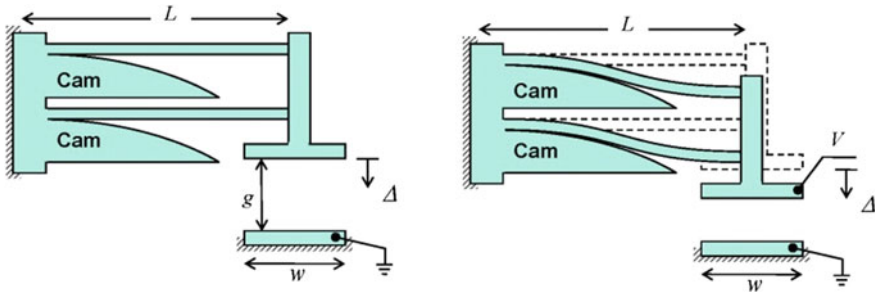

Fig. 5.2 Mechanically linearized gap-closing electrostatic actuator. The figure is reproduced here with the permission of Prof. David Elata

where F is the applied force, x is the displacement, ε is the dielectric constant, A is the area of the plates, V is the electric potential on the device, k is the spring stiffness, and g is the initial gap size.

Clever mechanical design can circumvent the complex nonlinearity of Eq. (5.1). Figure 5.2 shows a mechanically linearized modification of the device for which the force–displacement relation is, putatively, linear:

$$F = Kx \tag{5.2}$$

where K is a constant stiffness coefficient whose value is uncertain because it depends on the precise shapes of the cams that may vary in manufacture. We will explore the robustness to uncertainty in the stiffness coefficient of the linearized device. We will also explore robustness to uncertainty in a probabilistic model. Finally, we will consider opportuneness. We assume F and x to be positive.

In our first approach to this problem, we suppose that our knowledge of the stiffness coefficient, K, is quite limited. We know an estimated value, $\widetilde{K}$, and we have an estimate of the error, s, but the most we can confidently assert is that the true stiffness, K, deviates from the estimate by $\pm s$ or more, although K must be positive. We do not know a worst-case or maximum error, and we have no probabilistic information about K.

There are many types of info-gap models of uncertainty (Ben-Haim 2006). A fractional-error info-gap model is suitable to this state of knowledge:

$$U(h) = \left\{ K \colon K > 0, \quad \left| \frac{K - \widetilde{K}}{s} \right| \le h \right\}, \quad h \ge 0 \tag{5.3}$$

The info-gap model of uncertainty in Eq. (5.3) is an unbounded family of sets of possible values of the uncertain entity, which is the stiffness coefficient K in the present case. For any non-negative value of h, the set $U(h)$ is an interval of K values. Like all info-gap models, this one has two properties: *nesting* and *contraction*. "Nesting" means that the set $U(h)$ becomes more inclusive (containing more and more elements) as h increases. "Contraction" means that $U(h)$ is a singleton set containing only the known putative value $\widetilde{K}$ when $h = 0$. These properties endow h with its meaning as a "horizon of uncertainty."

5.4.1   IG Robustness

IG robustness is based on three components: a system model, an info-gap uncertainty model, and one or more performance requirements. In this present case, Eq. (5.2) is the system model and Eq. (5.3) is the uncertainty model. Our performance requirement is that the displacement, x, be no less than the critical value x_c.
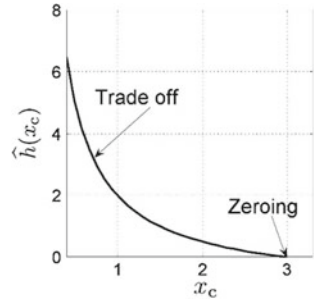
The IG robustness is the greatest horizon of uncertainty h up to which the system model obeys the performance requirement:

$$\hat{h}(x_c) = \max\left\{ h : \left(\min_{K \in U(h)} x \right) \ge x_c \right\} \tag{5.4}$$

Reading this equation from left to right, we see that the robustness $\hat{h}$ is the maximum horizon of uncertainty h up to which all realizations of the uncertain stiffness K in the uncertainty set $U(h)$ result in displacement x no less than the critical value x_c.

Robustness is a useful decision support tool because more robustness against uncertainty is better than less. Given two options that are approximately equivalent in other respects but one is more robust than the other, the robust-satisficing decisionmaker will prefer the more robust option. In short, "bigger is better" when prioritizing decision options in terms of robustness.

Fig. 5.3 Robustness curve of Eq. (5.5). $F/s = 3$, $\widetilde{K}/s = 1$

Derivation of the robustness function is particularly simple in this case. From the system model, we know that $x = F/K$. Let $m(h)$ denote the inner minimum in Eq. (5.4), and note that this minimum occurs, at horizon of uncertainty h, when $K = \widetilde{K} + sh$. The robustness is the greatest value of h up to which $m(h)$ is no less than x_c:

$$m(h) = \frac{F}{\widetilde{K} + sh} \geq x_c \Rightarrow \hat{h}(x_c) = \frac{1}{s}\left(\frac{F}{x_c} - \widetilde{K}\right) \tag{5.5}$$

or zero if this is negative which occurs if the performance requirement x_c is too large to be obtained with the putative system.

5.4.2 Discussion of the Robustness Results

The robustness function in Eq. (5.5) demonstrates two fundamental properties that hold for all IG robustness functions: *trade-off* and *zeroing*, illustrated in Fig. 5.3.

The performance requirement is that the displacement x be no less than the critical value x_c. This requirement becomes more demanding as x_c increases. We see from Eq. (5.5) and Fig. 5.3 that the robustness decreases as the requirement becomes more demanding. That is, robustness trades off against performance: The robustness can be increased only by relaxing the performance requirement. The negative slope in Fig. 5.3 represents the *trade-off between robustness and performance*: Strict performance requirements, demanding very good outcome, are less robust against uncertainty than lax requirements. This trade-off quantifies the intuition of any healthy pessimist: More demanding requirements are more vulnerable to surprise and uncertainty than lower requirements.

The second property illustrated in Fig. 5.3 is *zeroing*. Our best estimate of the stiffness is $\widetilde{K}$, so the predicted displacement is $x = F/\widetilde{K}$. Equation (5.5) shows that the robustness becomes zero precisely at the value of the critical displacement x_c that is predicted by the putative model: $x_c = F/\widetilde{K}$, which equals 3 for the parameter values in Fig. 5.3. Stated differently, the zeroing property asserts that best-model

predictions have no robustness against error in the model. Like trade-off, this is true for all info-gap robustness functions.

Models reflect our best understanding of the system and its environment. Nonetheless, the zeroing property means that model predictions are not a good basis for design or planning decisions, because those predictions have no robustness against errors in the models. Recall that we are discussing situations with large info-gaps as represented in Eq. (5.3): The putative value of the stiffness $\widetilde{K}$ is known, but the size of its deviation from the true stiffness K is unknown. If your models are correct (no info-gaps), then you do not need robustness against uncertainty. However, robustness is important when facing deep uncertainty.

The zeroing property asserts that the predicted outcome is not a reliable characterization of the system. The trade-off property quantifies how much the performance requirement must be reduced in order to gain robustness against uncertainty. The slope of the robustness curve reflects the cost of robustness: What decrement in performance "buys" a specified increment in robustness. Outcome quality can be exchanged for robustness, and the slope quantifies the cost of this exchange. In Fig. 5.3, we see that the cost of robustness is very large at large values of x_c and decreases as the performance requirement is relaxed.

The robustness function is a useful response to the pernicious side of uncertainty. In contrast, the opportuneness function is useful in exploiting the potential for propitious surprise. We now discuss the info-gap opportuneness function.

5.4.3   IG Opportuneness

IG opportuneness is based on three components: a system model, an info-gap model of uncertainty, and a performance aspiration. The performance aspiration expresses a desire for a better-than-anticipated outcome resulting from propitious surprise. This differs from the performance requirement for robustness, which expresses an essential or critical outcome without which the result would be considered a failure.

We illustrate the opportuneness function with the same example, for positive F and x. The robust-satisficing decisionmaker requires that the displacement be no less than x_c. The opportune windfalling decisionmaker recognizes that a larger displacement would be better, especially if it exceeds the anticipated displacement, $F/\widetilde{K}$. For the opportune windfaller, the displacement would be wonderful if it is as large as x_w, which exceeds the anticipated displacement. The windfaller's aspiration is not a performance requirement, but it would be great if it occurred.

Achieving a windfall requires a favorable surprise, so the windfaller asks: What is the lowest horizon of uncertainty at which windfall is possible (though not necessarily guaranteed)? The answer is the opportuneness function, defined as:

$$\hat{\beta}(x_w) = \min\left\{h : \left(\max_{K \in U(h)} x\right) \geq x_w\right\} \tag{5.6}$$

Reading this equation from left to right, we see that the opportuneness $\hat{\beta}$ is the minimum horizon of uncertainty h up to which at least one realization of the uncertain stiffness K in the uncertainty set $U(h)$ results in displacement x at least as large as the wonderful windfall value x_w. The opportuneness function $\hat{\beta}(x_w)$ is the complement of the robustness function $\hat{h}(x_c)$ in Eq. (5.4). The min and max operators in these two equations are reversed. This is the mathematical manifestation of the inverted meaning of these two functions. Robustness is the *greatest* uncertainty that *guarantees* the *required* outcome, while opportuneness is the *lowest* uncertainty that *enables* the *aspired* outcome.

Opportuneness is useful for decision support, because a more opportune option is better able to exploit propitious uncertainty than a less opportune option. An option whose $\hat{\beta}$ value is small is opportune, because windfall can occur even at low horizon of uncertainty. The opportune windfaller prioritizes options according to the smallness of their opportuneness function values: An option with small $\hat{\beta}$ is preferred over an option with large $\hat{\beta}$. That is, "smaller is better" for opportuneness, unlike robustness for which "bigger is better." We again note the logical inversion between robustness and opportuneness.

Whether a decisionmaker prioritizes the options using robustness or opportuneness is a methodological decision that may depend on the degree of risk aversion of the decisionmaker. Furthermore, these methodologies may or may not prioritize the options in the same order.

The opportuneness function is derived in a manner analogous to the derivation of Eq. (5.5), yielding:

$$\hat{\beta}(x_w) = \frac{1}{s}\left(\widetilde{K} - \frac{F}{x_w} \right) \tag{5.7}$$

or zero if this is negative, which occurs when x_w is so small, modest, and unambitious that it is possible even with the putative design and does not depend on the potential for propitious surprise.

5.4.4 Discussion of Opportuneness Results

The robustness and opportuneness functions, Eqs. (5.5) and (5.7), are plotted in Fig. 5.4. The opportuneness function displays zeroing and trade-off properties, whose meanings are the reverse of those for robustness. The opportuneness function equals zero at the putative outcome $x = F/\widetilde{K}$ like the robustness function. However, for the opportuneness function this means that favorable windfall surprise is not needed for enabling the predicted outcome. The positive slope of the opportuneness function means that greater windfall (larger x_w) is possible only at larger horizon of uncertainty.

The robustness and opportuneness functions may respond differently to proposed changes in the design solution, as we now illustrate. From Eq. (5.5), we note that

Fig. 5.4 Robustness and opportuneness curves of Eqs. (5.5) and (5.7)
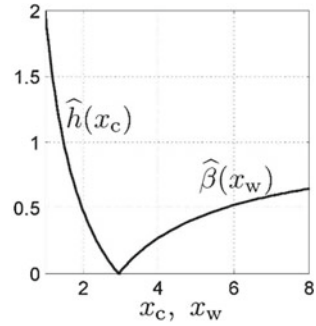
$\hat{h}$ decreases as the putative stiffness $\widetilde{K}$ increases. From Eq. (5.7), we see that $\hat{\beta}$ increases as $\widetilde{K}$ increases:

$$\frac{\partial \hat{h}}{\partial \widetilde{K}} < 0, \quad \frac{\partial \hat{\beta}}{\partial \widetilde{K}} > 0 \tag{5.8}$$

Recall that "bigger is better" for robustness while "smaller is better" for opportuneness. We see that any increase in $\widetilde{K}$ will make both robustness and opportuneness worse, and any decrease in $\widetilde{K}$ will improve them both. In summary, robustness and opportuneness are *sympathetic* with respect to change in stiffness.

Now consider the estimated error, s in the info-gap model of Eq. (5.3). A smaller value of s implies greater confidence in the estimate $\widetilde{K}$ while a larger s implies a greater propensity for error in the estimate. From Eq. (5.5), we see that robustness improves ($\hat{h}$ increases) as s decreases: Better estimate of $\widetilde{K}$ implies greater robustness against uncertainty in s. In contrast, from Eq. (5.7) we see that opportuneness gets worse ($\hat{\beta}$ increases) as s decreases: a lower opportunity for windfall as the uncertainty of the estimate declines. In short:

$$\frac{\partial \hat{h}}{\partial s} < 0, \quad \frac{\partial \hat{\beta}}{\partial s} < 0 \tag{5.9}$$

A change in the estimated error acts differently on robustness and opportuneness: By reducing the error of the estimated stiffness, one increases the robustness but diminishes the opportuneness; increasing the error acts in the reverse. In short, robustness and opportuneness are *antagonistic* with respect to the error in the estimated stiffness.

5.4.5 An Innovation Dilemma

An innovation dilemma occurs when the decisionmaker must choose between two options, where one is putatively better but more uncertain than the other.

Fig. 5.5 Robustness curves for innovative and state-of-the-art options
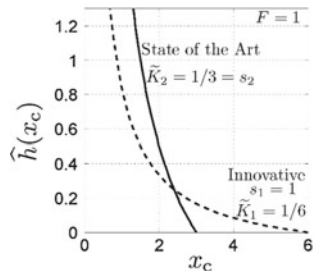
Technological innovations provide the paradigm for this dilemma. An innovation is supposedly better than the current state of the art, but the innovation is new so there is less experience with it and in practice it may turn out worse than the current state of the art. We will illustrate an innovation dilemma with the previous example, demonstrating its resolution using the robustness functions of the two options.

Consider two alternative designs (Option 1 and Option 2) of the linear elastic system (Eq. 5.2), one of which has lower estimated stiffness than the other:

$$\widetilde{K}_1 < \widetilde{K}_2 \tag{5.10}$$

Both designs will operate under the same positive force F, so the predicted displacement, $x = F/\widetilde{K}$, is greater with Option 1. Thus, Option 1 is preferred based on the estimated stiffnesses and the requirement for large displacement.

However, the putatively better Option 1 is based on innovations for which the actual stiffness in operation is more uncertain than for Option 2, which is the state of the art. Referring to the uncertainty estimate, s, in the info-gap model of Eq. (5.3), we express this as:

$$s_1 > s_2 \tag{5.11}$$

The dilemma is that Option 1 is putatively better (Eq. 5.10) but more uncertain (Eq. 5.11). This dilemma is manifested in the robustness functions for the two options, which also leads to a resolution, as we now explain. To illustrate the analysis, we evaluate the robustness function for each option (Eq. 5.5) with the following parameter values: $F = 1$, $\widetilde{K}_1 = 1/6$, $s_1 = 1$, $\widetilde{K}_2 = 1/3 = s_2$. The robustness curves are shown in Fig. 5.5.

The innovative Option 1 in Fig. 5.5 (dashed curve) is putatively better than the state-of-the-art Option 2 (solid curve), because the predicted displacement of Option 1 is $F/\widetilde{K}_1 = 6$, while the predicted displacement for Option 2 is only 3. However, the greater uncertainty of Option 1 causes a stronger trade-off between robustness and performance than for Option 2. The cost of robustness is greater for Option 1 for x_c values exceeding about 2, causing the robustness curves to cross one another at about $x_c = 2.4$.

Fig. 5.6 Robustness and
opportuneness curves for
innovative (dashed) and
state-of-the-art (solid)
options
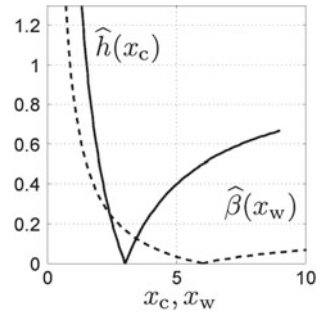
The innovation dilemma is manifested graphically by the intersection
of the robustness curves in Fig. 5.5, and this intersection is the basis for the res-
olution. Option 1 is more robust than Option 2 for highly demanding requirements
($x_c > 2.4$), and hence, Option 1 is preferred for this range of performance require-
ments. Likewise, Option 2 is more robust for more modest requirements ($x_c < 2.4$),
and hence, Option 2 is preferred for this lower range of performance requirements.

The robust-satisficing designer will be indifferent between the two options for
performance requirements at or close to the intersection value of $x_c = 2.4$. Con-
siderations other than robustness can then lead to a decision. Figure 5.6 shows the
robustness curves from Fig. 5.5 together with the opportuneness curves (Eq. 5.7)
for the same parameter values. We note that the innovative (dashed) Option is more
opportune (smaller $\hat{\beta}$) than the state-of-the-art Option (solid) for all values of x_w
exceeding the putative innovative value. Designers tend to be risk averse and to pre-
fer robust satisficing over opportune windfalling. Nonetheless, opportuneness can
"break the tie" when robustness does not differentiate between the options at the
specified performance requirement.

5.4.6 Functional Uncertainty

We have discussed the IG robustness function and its properties of trade-off, zeroing,
and cost of robustness. We have illustrated how these concepts support the decision
process, especially when facing an innovation dilemma. We have described the IG
opportuneness function and its complementarity to the robustness function. These
ideas have all been illustrated in the context of a one-dimensional linear system with
uncertainty in a single parameter. In most applications with deep uncertainty, the
info-gaps include multiple parameters as well as uncertainty in the shapes of func-
tional relationships. We now extend the previous example to illustrate the modeling
and management of functional or structural uncertainty in addition to the paramet-
ric uncertainty explored so far. This will also illustrate how uncertain probabilistic
models can be incorporated into an IG robust-satisficing analysis.

Let the stiffness coefficient K in Eq. (5.2) be a random variable whose estimated probability density function (pdf) $\tilde{p}(K)$ is normal with mean μ and variance σ^2. We are confident that this estimate is accurate for K within an interval around μ of known size $\pm\delta_s$. However, outside of this interval of K values, the fractional error of the pdf is unknown. In other words, we are highly uncertain about the shape of the pdf outside of the specified interval. This uncertainty derives from lack of data with extreme K values and absence of fundamental understanding that would dictate the shape of the pdf. We face "functional uncertainty" that can be represented by info-gap models of many sorts, depending on the type of information that is available. Given the knowledge available in this case, we use the following fractional-error info-gap model:

$$U(h) = \left\{ p(K) : \int_{-\infty}^{\infty} p(K)\mathrm{d}K = 1, \quad p(K) \geq 0 \text{ for all } K, \right.$$

$$p(K) = \tilde{p}(K) \text{ for} |K - \mu| \leq \delta_s,$$

$$\left. \left| \frac{p(K) - \tilde{p}(K)}{\tilde{p}(K)} \right| \leq h \text{ for} |K - \mu| > \delta_s \right\}, \quad h \geq 0 \qquad (5.12)$$

The first row of this info-gap model states that the set $U(h)$ contains functions $p(K)$ that are normalized and non-negative (namely, mathematically legitimate pdfs). The second line states that these functions equal the estimated pdf in the specified interval around the mean, μ. The third line states that, outside of this interval, the functions in $U(h)$ deviate fractionally from the estimated pdf by no more than h. In order to avoid some technical complications, we assume that the pdfs in $U(h)$ are non-atomic: containing no delta functions. In short, this info-gap model is the unbounded family of nested sets, $U(h)$, of pdfs that are known within the interval $\mu \pm \delta_s$ but whose shapes are highly uncertain beyond it. This is one example of an info-gap model for uncertainty in the shape of a function.

The system fails if $x < x_c$ where $x = F/K$ and F is a known positive constant. x is now a random variable (because K is random) so the performance requirement is that the probability of failure not exceed a critical value P_c. We will explore the robustness function. We consider the special case that $F/x_c > \mu \pm \delta_s$, meaning that the failure threshold for K lies outside the interval in which the pdf of K is known. The probability of failure is:

$$P_{\mathrm{f}}(p) = \mathrm{Prob}(x < x_c) = \mathrm{Prob}(K > F/x_c) = \int_{F/x_c}^{\infty} p(K)\mathrm{d}K \qquad (5.13)$$

For the estimated pdf, $\tilde{p}(K)$, one finds the following expression for the estimated probability of failure:

Fig. 5.7 Robustness curve
for Eq. (5.16)
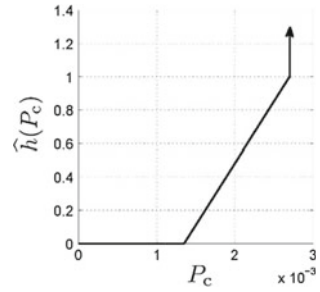
$$P_{\mathrm{f}}(\tilde{p}) = 1 - \Phi\left(\frac{(F/x_c) - \mu}{\sigma}\right) \tag{5.14}$$

where $\Phi(\cdot)$ is the cumulative pdf of the standard normal variate.

The robustness function, $\hat{h}(P_c)$, is the greatest horizon of uncertainty h up to which all pdf's $p(K)$ in the uncertainty set $U(h)$ do not have failure probability $P_{\mathrm{f}}(p)$ in excess of the critical value P_c:

$$\hat{h}(P_c) = \max\left\{h : \left(\max_{p \in U(h)} P_{\mathrm{f}}(p)\right) \leq P_c\right\} \tag{5.15}$$

After some algebra, one finds the following expression for the robustness:

$$\hat{h}(P_c) = \begin{cases} 0 & \text{if } 0 \leq P_c < P_{\mathrm{f}}(\tilde{p}) \\ \frac{P_c}{P_{\mathrm{f}}(\tilde{p})} - 1 & \text{if } P_{\mathrm{f}}(\tilde{p}) \leq P_c \leq 2P_{\mathrm{f}}(\tilde{p}) \\ \infty & \text{otherwise} \end{cases} \tag{5.16}$$

The robustness function in Eq. (5.16) is illustrated in Fig. 5.7 for $[(F/x_c) - \mu]/\sigma = 3$, meaning that the failure threshold is 3 standard deviations above the mean. Hence, the estimated probability of failure is $P_{\mathrm{f}}(\tilde{p}) = 1.35 \times 10^{-3}$. The trade-off property is evident in this figure: Lower (better) required probability of failure P_c entails lower (worse) robustness $\hat{h}(P_c)$. Note the discontinuous jump of robustness to infinity at $P_c = 2P_{\mathrm{f}}(\tilde{p})$. This is because the actual probability of failure $P_{\mathrm{f}}(p)$ cannot exceed more than twice the estimated value $P_{\mathrm{f}}(\tilde{p})$ resulting from constraints on the pdfs in the info-gap model of Eq. (5.12). The zeroing property is expressed by robustness becoming zero when the performance requirement P_c equals the estimated value $P_{\mathrm{f}}(\tilde{p})$.

5.5 Conclusion and Future Challenges

We live in an innovative world. Our scientific optimism embodies the belief that knowledge and understanding will continue to grow, perhaps at an ever-increasing rate. An inescapable implication of scientific optimism is that we are currently quite ignorant, and that we will be repeatedly and profoundly surprised in the near and not-so-near future. Because of this uncertainty, the planner, designer, or decisionmaker faces a profound and unavoidable info-gap: the disparity between what is currently known (or thought to be true) and what needs to be known for making a responsible decision (but is still hidden in the future). IG theory provides two complementary methodologies for managing this uncertainty. Robust satisficing helps in protecting against pernicious surprise and in achieving critical outcomes. Opportune windfalling helps in exploiting favorable surprise and in facilitating windfall outcomes.

We have illustrated the info-gap analysis in two examples, one quantitative and using mathematics, and one qualitative and using only verbal analysis. We discussed the trade-off between robustness and outcome requirements, showing that enhanced robustness against uncertainty is obtained only by relaxing the outcome requirements. In quantitative analysis, this allows an explicit assessment of the cost (in terms of reduced robustness) of making the requirements more demanding.

We also showed that predicted outcomes—based on the best available models and understanding—have no robustness against uncertainty in that knowledge. This has a profound implication for the decisionmaker and for the conventional conception of optimization. One's alternatives cannot be responsibly prioritized with their predicted outcomes, because those predictions have no robustness against uncertainty and surprise. Attempting to optimize the outcome, based on zero-robustness predictions, is not recommended when facing deep uncertainty. Instead, one should prioritize the alternatives according to their robustness for achieving critical outcomes, supplemented perhaps by analysis of opportuneness for windfall. This is particularly pertinent when facing an innovation dilemma: the choice between a new and putatively better alternative that is more uncertain due to its newness, and a more standard state-of-the-art alternative. The info-gap analysis of robustness enables the decisionmaker to assess the implications of uncertainty and to prioritize the alternatives to robustly achieve critical goals. We showed that this can lead to a reversal of the preference from the putatively optimal choice.

Many challenges remain. The quantitative analysis of robustness and opportuneness of large and complicated systems often faces algorithmic or numerical difficulties resulting from high dimensionality of the computations. Another challenge arises in response to new types of information and new forms of uncertainty. Many different mathematical forms of info-gap models of uncertainty exist (here we have examined only a few). However, analysts sometimes need to construct new types of info-gap models of uncertainty.

Another challenge is in bridging the gap between mathematics and meaning: between quantitative and qualitative analysis. Mathematics is a powerful tool that has facilitated the exploration of everything under the sun. However,

the incorporation of mathematics is difficult when knowledge is predominantly verbal, and when the meanings of subtle concepts are crucial to the decisionmaking process. A mathematical equation expresses a structural relationship between abstract entities. Meaning can be attributed to an equation, but meaning is not inherent in the equation. Witness the fact that the same equation can describe diverse and unrelated phenomena. Bridging the divide between mathematics and meaning is challenging in both directions. Quantitative analysts are often challenged to appreciate the limitations of their tools, while qualitative analysts often find it difficult to appreciate the contribution that mathematics can make.

The analysis and management of deep uncertainty faces many challenges, but our scientific optimism, tempered by recognition of our persistent ignorance, will carry us through as we acquire new understanding and face new surprises. As John Wheeler wrote (1992): "We live on an island of knowledge surrounded by a sea of ignorance. As our island of knowledge grows, so does the shore of our ignorance."

References

Ben-Haim, Y. (2006). *Info-gap decision theory: Decisions under severe uncertainty* (2nd ed.). London: Academic Press.

Ben-Haim, Y. (2010). *Info-gap economics: An operational introduction*. London: Palgrave-Macmillan.

Ben-Haim, Y. (2012a). Doing our best: Optimization and the management of risk. *Risk Analysis, 32*(8), 1326–1332.

Ben-Haim, Y. (2012b). Why risk analysis is difficult, and some thoughts on how to proceed. *Risk Analysis, 32*(10), 1638–1646.

Ben-Haim, Y. (2018). *Dilemmas of wonderland: Decisions in the age of innovation*. Oxford: Oxford University Press.

Ben-Haim, Y., Dacso, C. C., Carrasco, J., & Rajan, N. (2009). Heterogeneous uncertainties in cholesterol management. *International Journal of Approximate Reasoning, 50,* 1046–1065.

Ben-Haim, Y., & Demertzis, M. (2016). Decision making in times of Knightian uncertainty: An info-gap perspective, *Economics*, The Open-Access, Open Assessment e-journal, Special Issue Radical Uncertainty and Its Implications for Economics. No. 2016-23: 1–29. http://www.economics-ejournal.org/economics/journalarticles/2016-23.

Ben-Haim, Y., Zetola, N. M., & Dacso, C. (2012). Info-gap management of public health policy for TB with HIV-prevalence. *BMC Public Health, 12,* 1091. https://doi.org/10.1186/1471-2458-12-1091.

Burgman, M. (2005). *Risks and decisions for conservation and environmental management*. Cambridge: Cambridge University Press.

Chinnappen-Rimer, S., & Hancke, G. P. (2011). Actor coordination using info-gap decision theory in wireless sensor and actor networks. *International Journal of Sensor Networks, 10*(4), 177–191.

Hall, J. W., Lempert, R. J., Keller, K., Hackbarth, A., Mijere, C., & McInerney, D. J. (2012). Robust climate policies under uncertainty: A comparison of robust decision making and info-gap methods. *Risk Analysis, 32*(10), 1657–1672.

Hansen, L. P., & Sargent, T. J. (2008). *Robustness*. Princeton: Princeton University Press.

Harp, D. R., & Vesselinov, V. V. (2013). Contaminant remediation decision analysis using information gap theory. *Stochastic Environmental Research and Risk Assessment, 27*(1), 159–168.

Kanno, Y., & Takewaki, I. (2006). Robustness analysis of trusses with separable load and structural uncertainties. *International Journal of Solids and Structures, 43*(9), 2646–2669.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin Co. (Re-issued by University of Chicago Press, 1971).

Knoke, T. (2008). Mixed forests and finance—Methodological approaches. *Ecological Economics, 65*(3), 590–601.

Moffitt, L. J., Stranlund, J. K., & Field, B. C. (2005). Inspections to avert terrorism: Robustness under severe uncertainty. *Journal of Homeland Security and Emergency Management, 2*(3). http://www.bepress.com/jhsem/vol2/iss3/3.

Schlosser, E. (2013). *Command and control: Nuclear weapons, the Damascus accident, and the illusion of safety*. New York: Penguin Books.

Schwartz, B. (2004). *Paradox of choice: Why more is less*. New York: Harper Perennial.

Schwartz, B., Ben-Haim, Y., & Dacso, C. (2011). What makes a good decision? Robust satisficing as a normative standard of rational behaviour. *The Journal for the Theory of Social Behaviour, 41*(2), 209–227.

Simon, H. (1956). Rational choice and the structure of the environment. *Psych. Rev., 63*(2), 129–138.

Sims, C. A. (2001). Pitfalls of a minimax approach to model uncertainty. *American Economic Review, 91*(2), 51–54.

Smithson, M., & Ben-Haim, Y. (2015). Reasoned decision making without math? Adaptability and robustness in response to surprise. *Risk Analysis, 35*(10), 1911–1918.

Wald, A. (1947). *Sequential analysis*. J. Wiley & Sons (re-issued by Dover Publications, 1973).

Wheeler, J. A. (1992). Quoted in *Scientific American*, December, 1992, p. 20.

Prof. Yakov Ben-Haim (Technion—Israel Institute of Technology) is a professor of mechanical engineering and holds the Yitzhak Moda'i Chair in Technology and Economics. He initiated and developed Info-Gap (IG) Decision Theory for modeling and managing deep uncertainty. IG theory has impacted the fundamental understanding of uncertainty in human affairs and is applied by scholars and practitioners around the world in engineering, biological conservation, economics, project management, climate change, natural hazard response, national security, medicine, and other areas (see info-gap.com). He has been a visiting scholar in many countries and has lectured at universities, technological and medical research institutions, public utilities, and central banks. He has published more than 100 articles and 6 books.