34th IMAC, 25–28.1.2016, Orlando, Florida

Innovation Dilemmas, Design Optimization, and Info-Gaps

Yakov Ben-Haim
Yitzhak Moda'i Chair in Technology and Economics
Faculty of Mechanical Engineering
Technion—Israel Institute of Technology, Haifa, Israel

ABSTRACT Engineering designers seek innovative solutions to operational problems. The power of technological innovation is tremendous, and we invest great effort in training and promoting innovative design practice. However, one of the challenges in time-constrained design is that innovative ideas are usually less well understood than standard state of the art methods. This implies that innovative design solutions may contain hidden vulnerabilities that are absent from state of the art methods. While the innovative design may indeed be surprisingly better than the state of the art, it may also fail in surprising ways. We begin with a brief discussion of the non-probabilistic nature of ignorance. We then introduce the idea of an innovation dilemma and its resolution by info-gap robust-satisficing. We consider an example from the analysis of stress concentration around holes with uncertain shape.

KEYWORDS Innovation dilemma, info-gap, robustness, engineering design, stress concentration

1 What Is Info-Gap Theory?

Info-gap is a decision theory for prioritizing options and making choices under severe uncertainty [1, 2]. The options might be operational (design a system, choose a budget, decide to launch or not, etc.) or more abstract (choose a model structure, make a forecast, design a project flow chart, etc.). One uses data, scientific theories, empirical relations, knowledge and contextual understanding, all of which I'll refer to as one's *models*. Many models quantify uncertainty with probability distributions. If the models are correct (or judged to be reliable) then one can exploit them exhaustively to reliably achieve optimal outcomes, and one doesn't need info-gap. If the 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 severe uncertainties and info-gap theory might help.

2 Info-Gaps, Knightian Uncertainty, Ignorance, and Probability

Frank Knight [9] distinguished between 'risk' (for which probability distributions are known) and 'true uncertainty' (for which probability distributions are not known). Knightian ('true') uncertainty reflects ignorance of underlying processes, functional relationships, strategies or intentions of relevant actors, future events, inventions, discoveries,

[0] \papers\imac2016\in-dil-imac001.tex 17.10.2015

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* and what you *need to know* in order to make a reliable or responsible decision. 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, or of a probability density function, but the shape of the function (e.g. the shape of the elastic-plastic transition 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 even in sets of such entities.

Info-gap theory is for the epistemically impoverished. If your models (including probability models) are basically correct, then you don't need info-gap theory. Otherwise, maybe info-gap can help protect against adverse uncertainty or exploit favorable uncertainty.

3 Optimize or Satisfice?

There is a moral imperative to do one's best: to optimize the outcome of one's decisions. That usually means to use one's best models to achieve the best possible outcomes. While optimization may have prohibitive psychological costs [10], or one may simply not need the best possible mouse trap, the aspiration for excellence is usually to be commended.

Outcome-optimization is the process of using one's models to choose the decision whose outcome will be best. This works fine when the models are pretty good, because exhaustively exploiting good models will reliably lead to good outcomes.

However, when one faces major info-gaps, then one's models hide major flaws, and exhaustively exploiting the models can be unrealistic, unreliable, and can lead to major shortfalls and poor decisions [3]. Under severe uncertainty it is better to ask: what outcomes are critical and must be achieved? This is the idea of *satisficing,* introduced by Herbert Simon [12].

While engineering designers may use the language of optimization (the lightest, the strongest, the fastest, ...), in practice, engineers have been satisficing for ages: satisfying a design specification (light enough, strong enough, fast enough, ...). Furthermore, in competitive situations, one only needs to design a product that is better (lighter, stronger, faster, ...) than the competition's product. Beating the competition means satisficing a goal. Engineering progress is piecemeal, one improvement at a time, sometimes converging on an unspecified optimum, but more often diverging innovatively over the unbounded space of design possibilities.

4 Robustness: Trade Off and Zeroing

It commonly happens that several or many alternative designs all satisfy the design requirements. A single option must be chosen, which may be done by optimizing some exogenous variable (such as cost). However, when dealing with major info-gaps, it is better to prioritize the putatively adequate options according to their robustness against Knightian uncertainty.

Consider a design option that satisfies the performance requirements according to one's models. One then asks: how wrong can those models be, and this design will *still* satisfy the requirements? Equivalently: how robust against (or tolerant to) uncertainty is this design? A design that is more robust against uncertainty is preferred over a design that is less robust. This prioritization of the design options is called *robust-satisficing:* one prioritizes according to the robustness for satisficing the performance requirements.

Each design is characterized by a robustness curve as shown schematically in fig. 1. The horizontal axis represents the performance requirement, where a large value is better than a small value (e.g. long life or high stiffness is better than short or low values). The vertical axis is the robustness: the degree of immunity against ignorance, surprise, or uncertainty. The robustness curve displays two universal characteristics: trade off and zeroing.

The negative slope 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: higher aspirations (more demanding requirements) are more vulnerable to surprise than lower aspirations.

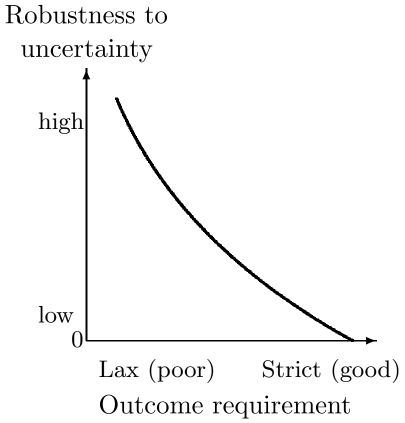

Figure 1: Robustness vs requirement.

The second property illustrated is fig. 1 is *zeroing:* The robustness precisely equals zero when the outcome requirement equals the outcome that is predicted by the models. 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 decisions, because those predictions have no robustness against errors in the models. Recall that we're discussing situations with large info-gaps. If your models are correct (no info-gaps), then you don't need robustness.

The zeroing property asserts that the predicted outcome is not a reliable characterization of the design. 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 "sold" in exchange for robustness, and the slope quantifies the cost of this exchange.

Info-gap robustness is non-probabilistic: it does not employ probability distributions and it does not explicitly provide insight about the likelihood of satisfying the design requirements. However, in many situations it is nonetheless true that a more robust design does indeed have greater probability of satisfying the requirements than a less robust design. When this is true we say that robustness is a "proxy" for probability. One can then maximize the likelihood of satisfying the requirements without knowing the probability distributions at all: one maximizes the info-gap robustness [5]. One will not know the value of this maximum probability of success because the probability distributions are unknown. However, one will nonetheless have found the design with greatest probability of success.

5 Preference Reversals and Innovation Dilemmas

Fig. 2 shows schematic robustness curves for two designs. Design 1 is predicted, by the models, to have a better outcome than design 2: the horizontal intercept for design 1 is to the right of (better than) the horizontal intercept for design 2. However, from the zeroing property we know that this is not a good basis for preferring design 1 over design 2: these predictions have zero robustness against info-gaps. Furthermore, in this particular case the robustness curves cross one another, reflecting different costs of robustness for the two designs.

Suppose that the outcome requirement is somewhere between the predicted outcome for design 1 (horizontal intercept in fig. 2) and the outcome at which the robustness curves cross one another ($r_{\times}$). Design 1 is more robust than design 2 for this outcome requirement. The robust satisficing designer will therefore prefer design 1 over design 2. Note that the outcome-optimizer, ignoring severe uncertainty, would also prefer design 1, but for a different reason: design 1 is predicted to be better than design 2. Outcome-optimization is a good basis for prioritization only if the models are reliable and we are not in a situation of severe uncertainty.

Now suppose that the outcome requirement is less than $r_{\times}$. Design 2 is more robust than design 1 for this requirement, and the robust satisficer will prefer design 2. The outcome-optimizer still prefers design 1, so these design strategies now prioritize the options differently.

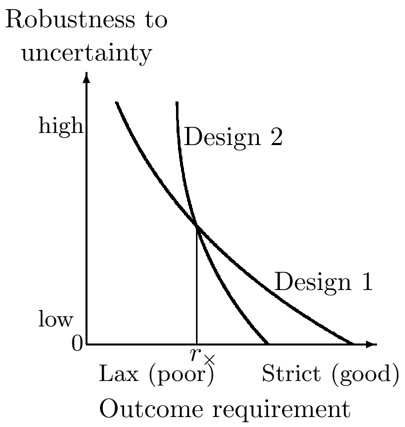

Figure 2: Preference reversal between 2 designs.

The crossing robustness curves in fig. 2 reflect the possibility of a *preference reversal:* design 1 is preferred for larger outcome requirements, while design 2 is preferred for less demanding requirements.

Preference reversals are particularly common when facing an *innovation dilemma* [4, 6]. Consider the choice between a new, innovative and highly promising design, and a more standard state-of-the-art design. The innovation is less familiar and more uncertain because it is new. Prioritizing these two design options is an innovation dilemma: should one choose the putatively better but more uncertain innovation, or the less uncertain standard option?

In fig. 2, design 1 is the innovative option and design 2 is the state of the art. Design 1 is putatively better (horizontal intercept further right), but more uncertain (lower slope; higher cost of robustness). The dilemma is expressed by the intersection of the robustness curves, and resolved as a preference reversal. One chooses the option that is more robust for the specified performance requirements.

6 Robust-Satisficing and Max-Min

Abraham Wald [13] studied statistical decisions when prior probability distributions are uncertain but bounded within a set of possible distributions. Based on identifying the most pernicious distribution, Wald proposed the idea of min-max: minimize the maximum possible loss. Wald's contribution was seminal by dealing non-probabilistically with uncertainty in probability distributions. (Note that max-min and min-max differ only in whether one is concerned with gain or loss respectively.)

Decisions based on info-gap robust-satisficing sometimes agree with Wald's max-min and sometimes do not, as explained with fig. 3 and elsewhere [7].

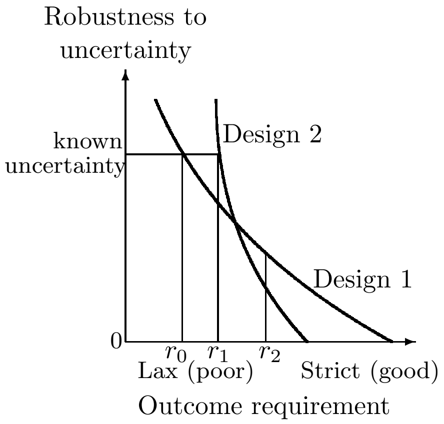

Figure 3: Preference reversal between 2 designs.

In order to compare max-min with info-gap robustness in fig. 3 one must understand that the vertical axis—robustness to uncertainty—can be thought of as a level of uncertainty (the greatest tolerable level). The method of max-min presumes knowledge of the level of uncertainty (which underlies identification of the worst case) and this known level of uncertainty is marked on the vertical axis. Design 2 is the max-min choice between the 2 designs in fig. 3 at the specified known level of uncertainty. This is because outcome r_1 is the worst that can occur with design 2 and it is better than r_0 (the worst with design 1 at this known level of uncertainty).

Robust-satisficing agrees with the max-min choice if the outcome requirement is for the value r_1 (or other values less than the curve intersection, $r_\times$ in fig. 2). They agree on the design, but for different reasons. Max-min ameliorates the worst case at known uncertainty. Robust-satisficing maximizes immunity against unbounded uncertainty while satisfying a design specification. Their reasons differ because their information differs. Min-max uses knowledge of the worst case. Robust-satisficing does not use a worst case (the info-gap model is unbounded so a worst case is unknown) but does use knowledge of an outcome requirement.

Now suppose that the robust-satisficer's outcome requirement is r_2 in fig. 3. In this case the robust-satisficing choice is for design 1, even if the designer agrees that uncertainty could be (or is) as large as what the max-min analyst knows it to be. The max-min choice is still for design 2, while the robust-satisficer disagrees.

Max-min and info-gap robust-satisficing share much in common, especially the non-probabilistic representation of uncertainty. However, they may or may not agree prescriptively, as we have just explained.

7 Innovation Dilemma in Design Against Stress Concentration

In this section we present a brief example of designing the pre-treatment of a drilled hole in an infinite plate, loaded at infinity with uniform radial stress, to deal with stress concentration on the boundary of the hole. This example is based on Sha'eir [11, chapter 9].

7.1 Formulation

The hole is designed to be precisely circular with radius R. The actual radius, at azimuthal angle θ with respect to a reference direction, is $r(\theta)$. The fractional radius-deviation-function is defined as:

$$g(\theta) = \frac{r(\theta) - R}{R} \tag{1}$$

This is expanded in a truncated Fourier series as:

$$g(\theta) = \sum_{m=1}^{M} (A_m \cos m\theta + B_m \sin m\theta) \tag{2}$$

The Fourier coefficients are represented as a vector: $c^T = (A_1, \ldots, A_M, B_1, \ldots, B_M)$. The stress concentration factor at angle θ on the boundary of the hole, $\kappa(c, \theta)$, is a function of the shape of the hole [8] as expressed by the Fourier coefficients c.

The shape of the hole may deviate from the circular design specification due to uncertain manufacturing. Hence c is uncertain. We use an ellipsoid-bound info-gap model [1] to represent this uncertain shape:

$$\mathcal{U}(h) = \left\{ c : \ (c - \widetilde{c})^T W (c - \widetilde{c}) \le h^2 \right\}, \quad h \ge 0 \tag{3}$$

$\widetilde{c}$ is the designer's best estimate of the Fourier coefficients. W is a real, symmetric, positive definite matrix representing the relative variability of the uncertain Fourier coefficients. Both $\widetilde{c}$ and W are known. $\widetilde{c}$ determines the centerpoint of the ellipsoid, and W determines its shape. The size of the ellipsoid is determined by h whose value is not known. Thus the info-gap model is not a single set of coefficients, but rather an unbounded family of nested sets. h represents the unknown horizon of uncertainty in the Fourier coefficients.

All info-gap models, of which eq.(3) is a special case, have two properties. *Nesting* is the property that the uncertainty sets, $\mathcal{U}(h)$, become more inclusive as h increases. The nesting property endows h with its meaning as an horizon of uncertainty. *Contraction* is the property that $\mathcal{U}(0)$ is a singleton set: there is no uncertain variability of the Fourier coefficients when the horizon of uncertainty is zero.

The performance requirement, that the designer must satisfy, is that the stress concentration at the boundary of the hole not exceed a specified critical value, κ_c:

$$\max_{0 \le \theta \le 2\pi} \kappa(\theta, c) \le \kappa_c \tag{4}$$

However, the Fourier coefficients c are not known, so this equation, by itself, is not a operationally implemental basis for design. Nonetheless, the concept of robustness against uncertainty in the shape provides a solution.

We begin by requiring that the putative design—a hole with estimated Fourier coefficients $\widetilde{c}$—satisfies the performance requirement in eq.(4). The *robustness question* is: how large a deviation of the Fourier coefficients, c, from their estimated values $\widetilde{c}$, can be tolerated without violating the performance requirement? Referring to the info-gap model in eq.(3), the robustness question can be stated as: what is the maximum horizon of uncertainty, h, up to which all Fourier coefficients c in $\mathcal{U}(h)$, satisfy the performance requirement? The answer to this question is the robustness function, whose formal definition is:

$$\widehat{h}(\kappa_c) = \max\left\{h : \left(\max_{c\in\mathcal{U}(h)} \max_{0\le\theta\le 2\pi} \kappa(\theta, c)\right) \le \kappa_c\right\} \tag{5}$$

Reading this equation from left to right: the robustness (against shape uncertainty), $\widehat{h}$, of performance requirement κ_c is the maximum horizon of uncertainty, h, up to which all shape realizations c in the set $\mathcal{U}(h)$ result in maximal stress concentration around the boundary that does not exceed κ_c.

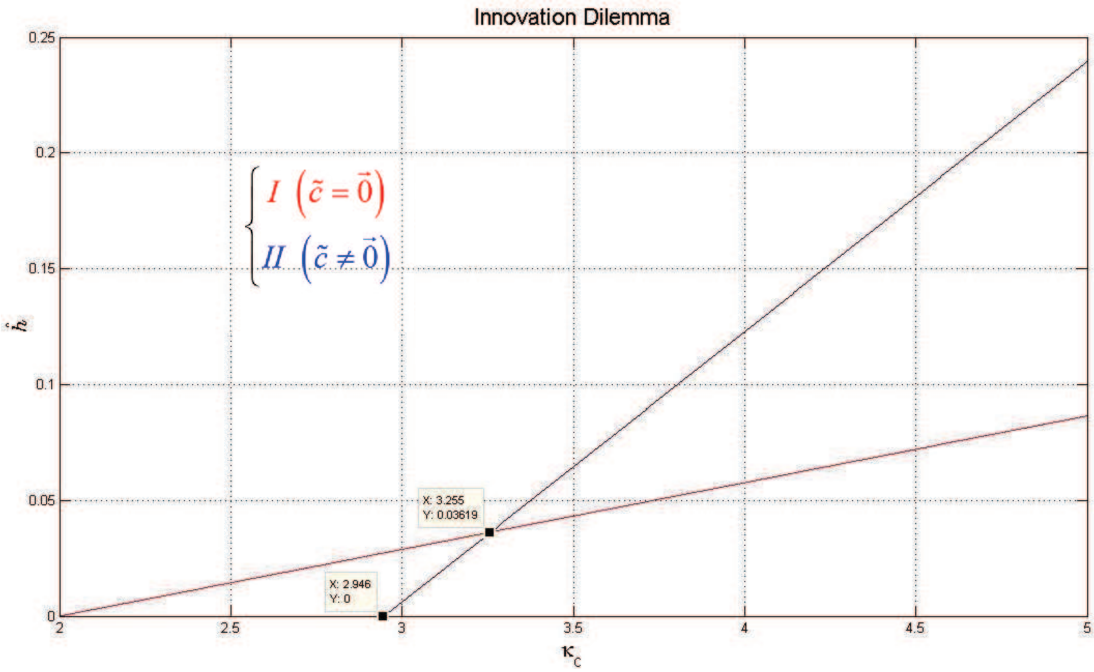

Figure 4: Robustness, $\widehat{h}(\kappa_c)$ vs. critical stress concentration, κ_c.

7.2 Numerical Example

We now present a numerical example based on Sha'eir [11, chapter 9]. The designer can choose between two alternatives: either *do* or *do not* polish the hole after drilling. The unpolished hole (design 1) has a putatively circular shape, but with the possibility of shape imperfections at high spatial frequency (entailing large stress concentration) due to manufacturing blemishes. The polished hole (design 2) is free of high-frequency shape imperfections, but the putative shape deviates from the nominal circular form due to small eccentricity of the polishing process.

We use the info-gap model of eq.(3) where the shape matrix, W, is diagonal: $W = \pi\mathrm{diag}(s_1^{-2}, \ldots, s_M^{-2}, s_1^{-2}, \ldots, s_M^{-2})$. The two designs are specified as follows:

$$\text{Design 1:} \qquad \widetilde{c}^T = 0, \quad M = 100, \quad s_m = \frac{3}{2}(m+2),\ m = 1, \ldots, M \tag{6}$$

$$\text{Design 2:} \qquad \widetilde{c}^T \neq 0, \quad M = 20, \quad s_m = m,\ m = 1, \ldots, M \tag{7}$$

The elements of $\widetilde{c}$ for design 2 are all zero except: $A_2 = 0.01$, $A_3 = 0.005$, $B_2 = -0.005$ and $B_5 = -0.055$. Thus design 1 is putatively better than design 2 ($\widetilde{c} = 0$ is better than $\widetilde{c} \neq 0$). However, design 1 is more uncertain than design 2 (large M and s_m entails greater uncertainty than small M and s_m). We have here an innovation dilemma, as we now explain.

The innovation dilemma is embodied in the robustness curves of the two designs, shown in fig. 4. Design 1 (red) has a putative stress concentration factor (SCF) of 2.0, as seen from the horizontal intercept of its robustness curve. Design 2 (blue) has a putative SCF of 2.9, as seen from its horizontal intercept. Thus design 1 is putatively better than design 2.

We note, however, that the zeroing property asserts that the robustness for the designs to obtain these putative SCF's is zero, as seen from the value of the robustness at the horizontal intercepts.

Furthermore, the robustness trade offs are different for the two designs. The slope of the design-1 robustness curve is less than the slope for design 2. This implies that the cost of robustness is greater for design 1: the performance must be relaxed more with design 1 than with design 2 in order to achieve the same increase in robustness.

The innovation dilemma is that design 1 is putatively better but more uncertain than design 2. This is manifested in fig. 4 by the intersection of the two robustness curves. The curve for design 1 is putatively better (horizontal intercept further to the left) but design 1 is more uncertain (lower slope); hence the curves intersect.

The resolution of the innovation dilemma is also represented by the intersecting robustness curves. The curves intersect at a critical SCF of 3.3. Design 1 is more robust, and hence preferred, if the designer requires an SCF less than 3.3, while design 2 is more robust, and hence preferred, if $\kappa_c > 3.3$ is acceptable.

REFERENCES

1. Ben-Haim, Yakov (2006). *Info-Gap Decision Theory: Decisions Under Severe Uncertainty,* 2nd ed., Academic Press, London.
2. Ben-Haim, Yakov (2010). *Info-Gap Economics: An Operational Introduction,* Palgrave-Macmillan, London.
3. Ben-Haim, Yakov (2012a). Doing our best: Optimization and the management of risk, *Risk Analysis,* 32(8): 1326–1332.
4. Ben-Haim, Yakov (2012c). Why risk analysis is difficult, and some thoughts on how to proceed, *Risk Analysis,* 32(10): 1638–1646.
5. Ben-Haim, Yakov (2014). Robust satisficing and the probability of survival, *Intl. J. System Science,* 45: 3–19.
6. Ben-Haim, Yakov, Craig D. Osteen & L. Joe Moffitt (2013). Policy dilemma of innovation: An info-gap approach, *Ecological Economics,* 85: 130–138.
7. Ben-Haim, Yakov, Clifford C. Dacso, Jonathon Carrasco & Nithin Rajan (2009). Heterogeneous uncertainties in cholesterol management, *Intl. J. of Approximate Reasoning,* 50: 1046–1065.
8. Givoli, Dan and Isaac Elishakoff (1992). Stress concentration at a nearly circular hole with uncertain irregularities, ASME *J. Applied Mech.,* 59: S65–S71.
9. Knight, Frank H. (1921). *Risk, Uncertainty and Profit,* Houghton Mifflin Co., Re-issued by University of Chicago Press, 1971.
10. Schwartz, Barry (2004). *Paradox of Choice: Why More Is Less,* Harper Perennial, New York.
11. Sha'eir, Or'el (2014). *Analysis of Stress Concentration Around a Hole with Uncertain Shape, Based on Info-gap Theory,* (in Hebrew). MSc thesis, Technion—Israel Institute of Technology.
12. Simon, Herbert (1956). Rational choice and the structure of the environment, *Psych.Rev.,* 63(2): 129–138.
13. Wald, A. (1945). Statistical decision functions which minimize the maximum risk, *Annals of Mathematics,* 46(2), 265–280.