Info-Gap Theory: An Intuitive Overview for Engineering Design and Reliability Assessment

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

ABSTRACT: An intuitive overview of central ideas in info-gap decision theory is presented, especially as it is applied in engineering design and reliability analysis. The idea of an info-gap and its relation to Knightian uncertainty, ignorance, and probability is discussed. The methodological distinction between optimizing and satisficing the outcome of a decision is examined. The trade off and zeroing properties of the info-gap robustness function and its applications to analysis of innovation dilemmas and preference reversals are explored. Robust-satisficing and max-min are compared. We dispel the notion that info-gap robustness is strictly a local analysis. The ideas of opportune windfalling and exploiting favorable uncertainty are examined. The advantages of combining info-gap with other decision methodologies are explained. We conclude by discussing some generic aspects of how one performs an info-gap analysis.

1 WHAT IS INFO-GAP THEORY?

Info-gap is a decision theory for prioritizing options and making choices under severe uncertainty (Ben-Haim 2006, 2010). 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 (1921) 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, 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 (see section 8 for discussion of 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 (Schwartz, 2004), 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 (Ben-Haim 2012a). 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 (1956).

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.

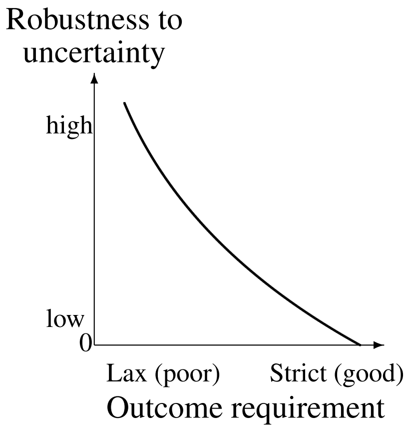

Figure 1: Robustness vs requirement.

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 probabilities 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 (Ben-Haim 2014). 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.

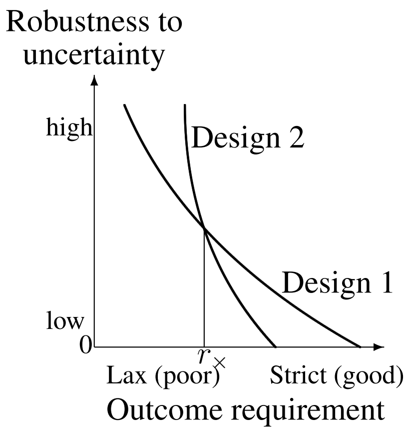

Figure 2: Preference reversal between 2 designs.

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* (Ben-Haim 2012c, Ben-Haim, Osteen & Moffitt 2013). 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 LOCAL OR GLOBAL ROBUSTNESS?

An info-gap model of uncertainty is an unbounded family of nested sets of possible occurrences. If, for instance, a probability distribution is uncertain, then the info-gap model for quantifying this uncertainty is an unbounded family of nested sets of probability distributions.

There are many different mathematical forms for info-gap models of uncertainty, but they all share two properties: nesting and contraction. *Nesting* is the property that the uncertainty sets become more inclusive as the horizon of uncertainty increases. This growing inclusiveness is what expresses the non-probabilistic uncertainty of the info-gap model. The expansion of the uncertainty sets continues without bound in the domain in which the uncertain entity is defined. *Contraction* states that the family of sets shrinks down to a singleton set in the absence of uncertainty. The single element in that contracted set is the nominal or best estimate of the uncertain entity. This nominal value constitutes (or is part of) the analyst's models upon which decisions will be based.

An info-gap analysis of robustness begins by stating a model—embodying one's best knowledge, data, insight, etc.—and by stating that this model may err in unknown and perhaps serious ways. The uncertainty is quantified with an info-gap model, and this family of nested sets is unbounded (on the domain of definition of its elements). There is no known worst case or greatest error in an info-gap model. This is deliberate, and intended to represent situations where surprise, invention, innovation and change can (and often does) occur.

The info-gap robustness question is posed in the context of this unbounded uncertainty. For a specified decision (e.g. design option or reliability assessment), one asks: what is the greatest horizon of uncertainty that does not allow violation of the performance requirements? The decision is highly robust if it can tolerate a large horizon of uncertainty without failure. Low robustness means failure can occur at low deviation of reality from the model. The robustness function is evaluated with respect to the model, but does not assume that uncertainty is limited to the near vicinity of that model. The uncertainty is global in the sense that the info-gap model of uncertainty is unbounded, while the robustness may be either small or large. The robustness analysis is not a local analysis—it is not limited in any way to the vicinity of the analyst's model—even if the value of the robustness is small. Small robustness simply means that the contemplated decision is highly vulnerable to the unbounded uncertainty of the info-gap model.

Further insight is obtained by distinguishing between two distinct questions: "How wrong is the model?" and "How much error in the model can be tolerated?" The latter is the robustness question, and it can be answered quantitatively, producing the info-gap robustness function. The former question is interesting and important but unanswerable with our current state of knowledge. Further research may lead to at least a partial answer.

7 ROBUST-SATISFICING AND MAX-MIN

Abraham Wald (1945) 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 (Ben-Haim *et al* 2009).

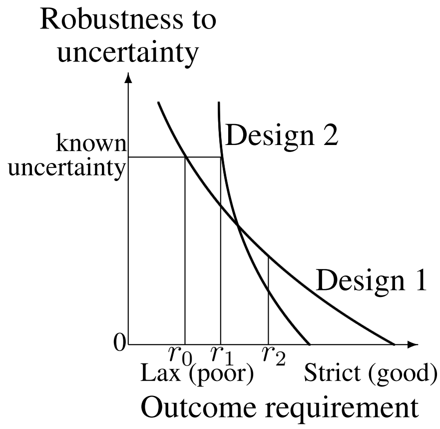

Figure 3: Preference reversal between 2 designs.

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.

8 WINDFALLING AND EXPLOITING UNCERTAINTY

The info-gap robustness function is used to prioritize options so as to protect the decision maker against pernicious uncertainty. However, uncertainty can be propitious and surprise can be beneficial. We now consider the exploitation of propitious uncertainty.

Opportune windfalling begins by aspiring to better-than-anticipated outcomes. If there is no uncertainty then only the anticipated outcome will occur. Wonderful windfalls are possible (but not guaranteed) only if there is uncertainty. When considering a specific decision option, the opportuneness question is: How much uncertainty is needed in order to enable (but not necessarily guarantee) the windfall outcome to which we aspire? The option is highly opportune if windfall is possible even at a low level of uncertainty. If great uncertainty is needed in order to enable the windfall then the option has low opportuneness.

Opportuneness is a desirable quantity, so the info-gap opportuneness function prioritizes the available options. This prioritization may or may not agree with the robustness prioritization. Furthermore, the trade off and zeroing properties of the robustness function have analogs with the opportuneness function.

9 HYBRIDIZATION: TWO IS BETTER THAN ONE (SOMETIMES)

"Two is better than one, ... and the three-fold cord isn't quickly broken." (*Kohelet,* 4: 9, 12).

I stressed earlier that the relevance of info-gap theory for supporting decisions depends on what the analyst knows and on the degree of severity of the uncertainties. If the models are reliable then info-gap decision theory may not be needed. One might base the analysis on probabilistic, statistical or perhaps even deterministic methods.

Sometimes one's models are pretty good in some respects and more uncertain in others. For instance, Wald considered statistical decisions in which prior probabilities are uncertain while conditional probabilities are known. In such situations one can robustify one's decision methodology by embedding it in an info-gap analysis. For example, a Bayesian decision can be embedded in an info-gap robustness analysis to strengthen its immunity against uncertainty in the prior probabilities (Burgman *et al* 2010). Methodological 'hybridization' like this is quite common.

In some situations one has no probabilistic or other uncertainty models, and one uses info-gap as a stand-alone tool for supporting decision making. This is widely done to handle parameter uncertainty and uncertainties in functional forms: material constitutive relations (Ben-Haim 2012b); supply and demand curves (Stranlund & Ben-Haim 2008); model-form in empirical up-dating (Atamturktur *et al* 2014) and more (many citations on info-gap.com).

The criterion for formulating one's decision methodology needs to be epistemic (what do I know and where are my info-gaps) and not scholastic (what is *my* decision theory). Kohelet was right (sometimes).

10 HOW DOES ONE 'DO' INFO-GAP?

We have not addressed the 'how do you do it?' question in this paper because there is a huge literature of examples, methodological discussions and theoretical explorations. Two books (Ben-Haim 2006, 2010) deal extensively with methodology and are chock full of examples. Burgman (2005) has a very accessible chapter on info-gap. The website info-gap.com has numerous citations and abstracts (with some links to articles) of work by scholars from around the globe dealing the info-gap theory and applications in many disciplines.

The math needed to do an info-gap robustness or opportuneness analysis is mostly a function of the math you need to describe your system and your performance requirements.

The foremost challenge of an info-gap analysis is in formulating the info-gap model of uncertainty. As the fictional detective Nero Wolfe might have said, "judgment based on intelligence guided by experience" (Rex Stout 1950) is needed in choosing the structure of the info-gap model. Furthermore, intellectual modesty and an openness to the vast unknown are needed in realistically assessing our ignorance and the potential for surprise. *"History* [meaning the future] *has a habit of being richer and more ingenious than the limited imaginations of most scholars or lay-*

men." (Kahn 1961, p.137, italics in the original).

The second major challenge in successfully using info-gap is in incorporating robustness (and possibly also opportuneness) functions in deliberation and decision making. Info-gap supports the assessment of trade offs (and trade ons) of various sorts. In section 4 we discussed how robustness against uncertainty trades off against the quality of the outcome that can be demanded. Analogously, the opportuneness function (section 8) expresses a trade off between the amount of uncertainty one must accept and the wonderfulness of the windfall to which one aspires. Furthermore, robustness may trade off (or trade on) with opportuneness from uncertainty. If they trade *off,* then some robustness can be 'sold' in exchange for some opportuneness. If they trade *on,* then an improvement in robustness also improves the opportuneness (and you can even optimize them both). An info-gap robustness analysis also supports deliberations over dilemmas of innovation as discussed in section 5, leading perhaps to reversal of one's preference from the initial choice based on the models and ignoring severe uncertainty. Robustness functions support other deliberative issues, such as the value of new (or potentially new) information. If the past is any indication of the future, new uses of info-gap, and new types of insights, will emerge as new applications are explored.

ACKNOWLEDGEMENT

The author is indebted to Dr. Ullrika Sahlin for valuable and insightful comments on this paper.

REFERENCES

Atamturktur, S., Z.Liu, S.Cogan, S. & H.Juang (2014). Calibration of imprecise and inaccurate numerical models considering fidelity and robustness: a multi-objective optimization-based approach, *Structural and Multidisciplinary Optimization,* published online 23 August 2014.

Ben-Haim, Yakov (2006). *Info-Gap Decision Theory: Decisions Under Severe Uncertainty,* 2nd ed., Academic Press, London.

Ben-Haim, Yakov (2010). *Info-Gap Economics: An Operational Introduction,* Palgrave-Macmillan, London.

Ben-Haim, Yakov (2012a). Doing Our Best: Optimization and the Management of Risk, *Risk Analysis,* 32(8): 1326–1332.

Ben-Haim, Yakov (2012b). Modeling and design of a hertzian contact: An info-gap approach, *J. Strain Analysis for Engineering Design,* 47(3): 153–162.

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

Ben-Haim, Yakov (2014). Robust satisficing and the probability of survival, *Intl. J. of System Science,* 45: 3–19

Ben-Haim, Yakov, Craig D. Osteen & L. Joe Moffitt (2013). Policy dilemma of innovation: An info-gap approach, *Ecological Economics,* 85: 130–138.

Ben-Haim, Yakov, Clifford C. Dacso, Jonathon Carrasco & Nithin Rajan (2009). Heterogeneous uncertainties in cholesterol management, *Intl. J. of Approximate Reasoning,* 50: 1046–1065.

Burgman, Mark (2005). *Risks and Decisions for Conservation and Environmental Management,* Cambridge University Press, Cambridge.

Burgman, M.A., B.A.Wintle, C.A.Thompson, A.Moilanen, M.C.Runge & Yakov Ben-Haim (2010). Reconciling uncertain costs and benefits in Bayes nets for invasive species management, *Risk Analysis,* 30(2): 277–284.

Kahn, Herman (1961). *On Thermonuclear War,* 2nd ed., Princeton University Press, Princeton, New Jersey.

Knight, Frank H. (1921). *Risk, Uncertainty and Profit,* Houghton Mifflin Co., Re-issued by University of Chicago Press, 1971.

Schwartz, Barry (2004). *Paradox of Choice: Why More Is Less,* Harper Perennial, New York.

Simon, Herbert (1956). Rational choice and the structure of the environment, *Psychological Review,* 63(2): 129–138.

Stout, Rex (1950). *In the Best Families,* Viking Press.

Stranlund, John K. & Yakov Ben-Haim, (2008). Price-based vs. quantity-based environmental regulation under Knightian uncertainty: An info-gap robust satisficing perspective, *J. Environmental Management,* 87: 443–449.

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