Working Paper No. MS-3-06

The Fundamental Flaw in Info-Gap's Uncertainty Model

Moshe Sniedovich
Department of Mathematics and Statistics
The University of Melbourne
Melbourne, Vic 3001, Australia

moshe@ms.unimelb.edu.au
www.ms.unimelb.edu.au/∼moshe

December 31, 2006
Last correction: September 21, 2007

Abstract

Info-Gap is supposedly a new theory for decision making under severe uncertainty. Its claim to fame is that it is non-probabilistic in nature and thus offers an alternative to all current theories for decision making under uncertainty. In this short article I explain why *Info-Gap*'s uncertainty model is flawed and why there is no reason to believe that the solutions generated by *Info-Gap* are likely to be robust, let alone optimal and robust.

Keywords: Decision making, severe uncertainty, maximin, worst case analysis, info-gap, robust.

1 Introduction

As we should all know only too well, it is conceptually wrong (in general) to deduce conclusions about *global* behavior from a *local* search. For example, in general, the deployment of a *local* search procedure in a given small neighborhood of the feasible region is unlikely to generate a *global* optimal solution.

For similar reasons, it is very unlikely that an analysis based on a *point*[1] *estimate* of an unknown parameter will generate a robust solution, especially if the estimate is subject to *severe* uncertainty.

[1]an element of the topological space under consideration.

It is therefore rather intriguing that *Info-Gap* (Ben-Haim [2001, 2006]) claims that the solutions it generates are robust. After all, *Info-Gap*'s uncertainty model is based on a *single estimate* of a parameter whose true value is subject to *severe* uncertainty.

How could it be?

In this short discussion I explain why there is no reason to believe that the solutions generated by *Info-Gap* are likely to be robust, let alone optimal and robust.

2 Generic Info-Gap model

The generic *Info-Gap* model consists of the following ingredients:

- A decision space, $\mathbb{Q}$.
- An uncertainty space $\mathcal{U}$.
- A real valued function R on $\mathbb{Q}\times\mathcal{U}$.
- A critical reward value r_c.
- An estimate $\tilde{u}$ of a parameter $u\in\mathcal{U}$ whose true value is unknown.
- A set of nested regions of uncertainty, $\mathcal{U}(\alpha,\tilde{u})\subseteq\mathcal{U},\alpha\geq 0$ such that $\mathcal{U}(0,\tilde{u})=\{\tilde{u}\}$ and $\mathcal{U}(\alpha,\tilde{u})$ is non-decreasing with α, namely $\alpha>\alpha'$ implies $\mathcal{U}(\alpha',\tilde{u})\subseteq\mathcal{U}(\alpha,\tilde{u})$.

The *Info-Gap* recipe goes like this:

The *robustness* of decision $q\in\mathbb{Q}$ is defined as

$$\hat{\alpha}(q,r_c):=\max\left\{\alpha\geq 0: r_c\leq\min_{u\in\mathcal{U}(\alpha,\tilde{u})}R(q,u)\right\}\tag{1}$$

That is, it is the largest value of α such that the requirement $r_c\leq R(q,u)$ is satisfied for all $u\in\mathcal{U}(\alpha,\tilde{u})$.

The best decision is then one whose robustness is the largest. To find such a decision we thus have to solve the following optimization problem:

$$\begin{aligned}\hat{\alpha}(r_c) &:= \max_{q\in\mathbb{Q}}\ \hat{\alpha}(q,r_c) &(2)\\ &= \max_{q\in\mathbb{Q}}\ \max\left\{\alpha\geq 0: r_c\leq\min_{u\in\mathcal{U}(\alpha,\tilde{u})}R(q,u)\right\} &(3)\end{aligned}$$

In words, our mission is to find the largest value of α such that the constraint $r_c\leq R(q,u)$ is satisfied for all u in $\mathcal{U}(\alpha,\tilde{u})$ for some $q\in\mathbb{Q}$.

With no loss of generality we assume that $r_c\leq R(q,\tilde{u})$ for all $q\in\mathbb{Q}$. Any $q\in\mathbb{Q}$ that does not satisfy this requirement can be removed from $\mathbb{Q}$ at the outset.

3 Local vs Global analysis

It is very unfortunate that the *Info-Gap* notation hides the fact that the analysis is LOCAL *par excellence*. To fix this let us re-write the model properly:

$$\hat{\alpha}(q, r_c, \tilde{u}) := \max\left\{\alpha \geq 0 : r_c \leq \min_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u)\right\} \tag{4}$$

$$\hat{\alpha}(r_c, \tilde{u}) := \max_{q \in \mathbb{Q}} \ \hat{\alpha}(q, r_c) \tag{5}$$

$$= \max_{q \in \mathbb{Q}} \ \max\left\{\alpha \geq 0 : r_c \leq \min_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u)\right\} \tag{6}$$

The correction involves the inclusion of the estimate $\tilde{u}$ as an argument in the expressions for the robustness of a decision, namely $\hat{\alpha}(q, r_c, \tilde{u})$ and the robustness for the maximal robustness $\hat{\alpha}(r_c, \tilde{u})$.

This correction is a reminder that the analysis is based on a single point estimate of the unknown value of the parameter of interest. That is, the analysis is based only on $\tilde{u}$ and its immediate neighborhood.

In any case, the question is this: how should we interpret these measures of robustness? Are they of a global or local nature with regard to the region of uncertainty $\mathcal{U}$? Do they incorporate features of possible values of u in specific neighborhoods of $\mathcal{U}$? Or are they based on representative values of u appropriately selected from $\mathcal{U}$ to ensure that they adequately represent the whole of $\mathcal{U}$?

This, of course, is a rhetorical question. The *Info-Gap* analysis is based on regions of $\mathcal{U}$ in the immediate neighborhood of the estimate $\tilde{u}$. The correction above is just a reminder that this is the case.

So, how should we interpret $\hat{\alpha}(q, r_c, \tilde{u})$?

However we interpret it, clearly, to make sense the interpretation should be *local* in nature, that is, local to the region of uncertainty $\mathcal{U}(\alpha^\circ, \tilde{u})$ where $\alpha^\circ = \hat{\alpha}(q, r_c, \tilde{u})$. And since the analysis is under severe uncertainty, we can assume that in all likelihood the true value of u is far away from the estimate $\tilde{u}$.

A schema of this nature is shown in Figure 1, where the region of uncertainty $\mathcal{U}(\alpha^\circ, \tilde{u})$ is shown as a circle centered at $\tilde{u}$.

To be blunt, given that the true value of u is far away from the region $\mathcal{U}(\alpha^\circ, \tilde{u})$, why should we care at all about the analysis conducted in this region? And why should we bother at all about the fact that the decision q is robust in that region?

The short answer is: we should not.

In other words, associated as it is with a region in the neighborhood of the estimate $\tilde{u}$, the robustness considered by the *Info-Gap* model is *local in nature*. It is local by design, and should therefore be interpreted as such. The trouble is that under *severe uncertainty* of u's true value, such a local notion of robustness is a very poor indication of the true robustness and can be substantially wrong.

I should add that what is shown in Figure 1 is "on scale" and is confirmed by numerical experiments with simple problems (Sniedovich [2006]).

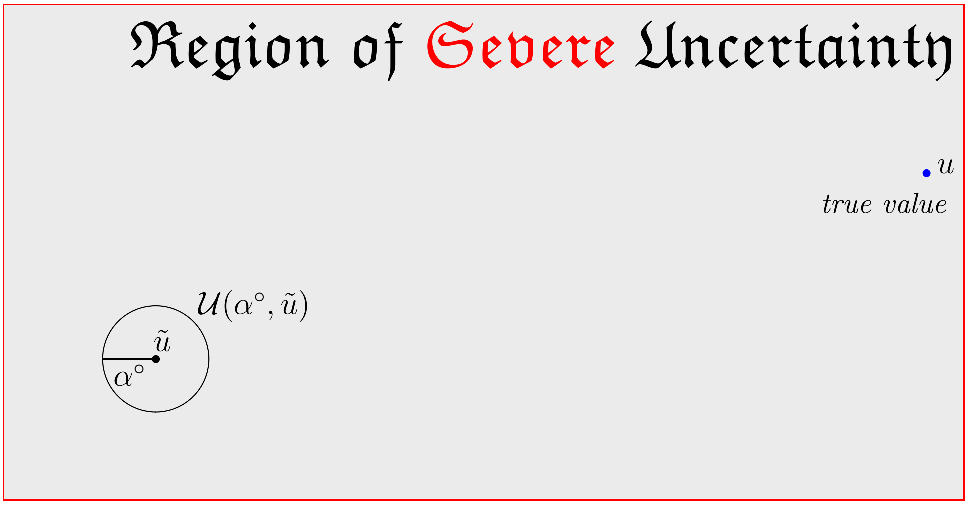

Figure 1: The region of uncertainty

The puzzling thing about this entire business is that *Info-Gap* itself warns us against the use of a point estimate in decision making under severe uncertainty. Indeed, *Info-Gap* goes out of its way – and rightly so – to do so. It stresses that under SEVERE uncertainty it is "wrong" to base our decision on a *single point estimate* of the uncertainty parameter under consideration. The argument is simple: under severe uncertainty estimates are of poor quality and are likely to be substantially wrong.

To illustrate the issues involved, suppose that you have to choose an option, or alternative, from a given collection of n options or alternatives. Let v_i denote the value of option i, $i = 1, \ldots, n$ and assume that "larger is better" so ideally, under conditions of strict certainty, you would select the option whose v_i value is largest.

But what should we do in cases where the exact values of v_i, $i = 1, 2 \ldots, n$ are unknown and are subject to SEVERE uncertainty? For example, consider the following concrete case:

option, i	v_i	$\hat{v}_i$
1	?	17
2	?	21
3	?	18

assuming that the estimates $\hat{v}_i, i = 1, 2, 3$, are subject to severe uncertainty.

Here is *Info-Gap*'s expert advice on the fundamental issue represented by such simple decision making problems:

> The value of v_i is highly uncertain and possibly varying in time, so that historical evidence is of limited utility. The best estimate of the value of option i is $\tilde{v}_i$. For instance, this might be an historical mean, perhaps over a limited time window, and perhaps with temporal lag. Since things change, or since the long-range mean deviates greatly from the mean on short time intervals, the estimate is a poor indication of the true value that will accrue from option i the next time a choice is made.
>
> Ben-Haim [2006, p. 280]

In other words, *Info-Gap* argues the obvious: estimates obtained under severe uncertainty should be regarded as POOR approximations of the true values they represent.

Now, let $\tilde{v}_i$ denote the best estimate we have for the true value of v_i and let i^* denote the option whose $\tilde{v}_i$ value is largest, namely let $i^* = \arg max\{\tilde{v}_i : i = 1, \ldots, n\}$.

Here is what *Info-Gap* says about the choice of option i^* as the best (optimal) option:

> The large value of $\tilde{v}_{i^*}$ is desirable. But $\tilde{v}_{i^*}$ is only an estimate of the value of option i, and this estimate is likely to be substantially wrong. An additional reason that large $\tilde{v}_{i^*}$ is attractive is the implicit assumption that, since $\tilde{v}_{i^*}$ is large, then the actual value of option i^* is also large even if $\tilde{v}_{i^*}$ errs. This of course is not necessarily true.
>
> Ben-Haim [2006, p. 281]

In other words, *Info-Gap* warns us against the simplistic policy of ranking alternatives on the basis of poor estimates resulting from severe uncertainty. The reason: these estimates are likely to be SUBSTANTIALLY WRONG.

Who can argue against this sound advice?

THEOREM 1. *Info-Gap's uncertainty model is fundamentally flawed and there is no reason to believe that the solutions it generates are likely to be robust.*

Proof. *Info-Gap* correctly argues that under severe uncertainty $\tilde{u}$ should be regarded as a poor estimate of the true value of u and is very likely to be substantially wrong. Since the robustness indices $\hat{\alpha}(r_c, \tilde{u})$ and $\hat{\alpha}(r_c, \tilde{u})$ are based on this poor estimate, they themselves must be regarded as poor indicators of the true robustness indices (around the true value of u). Hence, there is no reason to believe that solutions it generated are likely to be robust. $\mathfrak{QED}$

One of the consequences of the local nature of *Info-Gap*'s uncertainty model is that the generic *Info-Gap* model is *completely oblivious* to the "size" of the total region of uncertainty, call it $\mathfrak{U}$, in relation to the "size", $\hat{\alpha}(r_c)$, of the optimal region of uncertainty $\mathcal{U}(\hat{\alpha}(r_c), \tilde{u})$. More precisely,

THEOREM 2. *Info-Gap does not deal with severe uncertainty, it simply ignores it. More precisely, the generic Info-Gap model is invariant with the "size" of total region of uncertainty $\mathfrak{U}$: the value of $\hat{\alpha}(r_c)$ does not vary with $\mathfrak{U}$ for all $\mathfrak{U}$ such that $\mathcal{U}(\hat{\alpha}(r_c)+\varepsilon, \tilde{u}) \subseteq \mathfrak{U}$ for some $\varepsilon > 0$.*

PROOF. Let $\alpha^* := \hat{\alpha}(r_c)$ and $\mathfrak{U}^* := \mathcal{U}(\alpha^* + \varepsilon, \tilde{u}), \varepsilon > 0$. We have to show that α^* does not vary with $\mathfrak{U}$ for all $\mathfrak{U}$ such that $\mathfrak{U}^* \subseteq \mathfrak{U}$. This follows immediately from the nesting property of the regions of uncertainty $\mathcal{U}(\alpha, \tilde{u}), \alpha \geq 0$ and the worst-case characteristic of robustness stipulated in the definition of $\hat{\alpha}(r_c)$. $\mathfrak{QED}$

This point is illustrated in Figure 2 where three regions of uncertainty are displayed, $\mathfrak{U} \subset \mathfrak{U}' \subset \mathfrak{U}''$. The same solution, α^*, is obtained for any region of uncertainty containing the set $\mathcal{U}(\alpha^*, \tilde{u})$ represented by the circle.

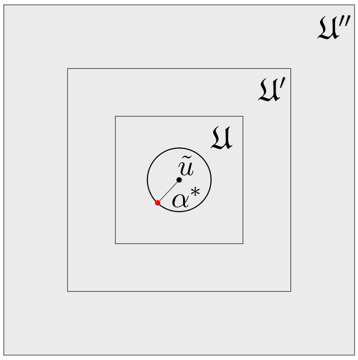

Figure 2: Illustration of Theorem 2

To appreciate the implication of this fact consider the following case: you have just solved a decision making problem under severe uncertainty using *Info-Gap* and obtain an optimal decision q^* whose robustness is $\alpha^* = \hat{\alpha}(r_c) = \hat{\alpha}(q^*, r_c)$. Then you discover the bad news that you actually severely underestimated the severity of the uncertainty associated with your problem: the level of uncertainty is 1000-fold larger! That is, the true total region of uncertainty $\mathfrak{U}$ is 1000-fold larger.

Since this news means that the updated total region of uncertainty contains the old one, there is no change in the *Info-Gap*'s analysis and the same results will be generated: there is no change is the optimal decision and there is no change in its robustness.

Isn't this ridiculous?

4 Robust Optimization

The good news is that there has been a lot of progress in the area of *Robust Optimization* over the past thirty years so there are today methods and techniques for obtaining robust optimal solutions in the framework of decision making under severe uncertainty.

By coincidence, a special issue of the journal *Mathematical Programming* dedicated to *Robust Optimization* was published this year (Ben-Tal et al [2006]).

It is unfortunate that *Info-Gap* seems to be unaware of the extensive body of knowledge available in this area of optimization, which is very relevant to what *Info-Gap* is attempting to do.

5 The (missing) Maximin Connection

For some strange reason, Info-Gap is completely oblivious to the fact that its generic model is an instance of the famous Wald's Maximin model (Wald [1945, 1950]). What is so special about Info-Gap's deployment of Wald's old and very established model is that the analysis is conducted in the immediate neighborhood of a poor estimate of the parameter of interest.

More on this bizarre and problematic aspect of Info-Gap can be found in Sniedovich [2006, 2006a].

6 Conclusion

The notion of robustness deployed by *Info-Gap* is – by design – local in nature relative to the total region of uncertainty. Therefore, this notion should be used with care in decision making under severe uncertainty where – as a rule – the estimates are poor and very likely to be substantially wrong.

In short, Info-Gap's uncertainty model is fundamentally flawed because it does not deal with severe uncertainty, it simply ignores it. Info-Gap's recipe involves then two ingredients:

- Replacing severe uncertainty by a poor estimate of the parameter of interest, knowing full well that it is likely to be substantially wrong.

- Conducting a standard Maximin analysis in the immediate neighborhood of this estimate.

The first amounts to practicing voodoo decision making. Figure 1 speaks for itself.

References

[1] Ben-Haim Y, *Information Gap Decision Theory*, Academic Press, 2001.

[2] Ben-Haim Y, *Info-Gap Decision Theory*, Elsevier, 2006.

[3] Ben-Tal A, El Ghaoui L, Nemirovski A, *Mathematical Programming*, Special issue on *Robust Optimization*, Volume 107(1-2), 2006.

[4] Sniedovich[1] M, *What's Wrong with Info-Gap? An Operational Research Perspective*, a paper presented at the *ASOR Recent Advances 2006* mini-conference, December 1, 2006, Melbourne, Australia.

[5] Sniedovich[1] M, *Eureka! Info-Gap is Worst Case Analysis (Maximin) in Disguise!*, Working Paper No. MS-2-2006, Department of Mathematics and Statistics, The University of Melbourne, Melbourne, Australia, 2006a.

[6] Wald A, Statistical decision functions which minimize the maximum risk, *The Annals of Mathematics,* 46(2), 265-280, 1945.

[7] Wald, A, *Statistical Decision Functions*, John Wiley, 1950.

[1]Copies can be found at `www.ms.unimelb.edu.au/~moshe/frame_maximin.html`