

Decision trade-offs under severe info-gap uncertainty

Yakov Ben-Haim

Faculty of Mechanical Engineering
Technion — Israel Institute of Technology
Haifa, Israel 32000
<http://aluf.technion.ac.il/~yakov>
yakov@aluf.technion.ac.il

Abstract

We are concerned in this paper with the trade-offs which confront a decision maker who deals with severely deficient information and unstructured uncertainty. We employ the theory of info-gap uncertainty, described briefly in section 1. Using info-gap models of uncertainty we derive two decision functions which express (1) immunity to failure (robustness function) and (2) immunity to windfall gain (opportunity function). These immunity functions are discussed in section 2. In this paper we will consider three types of trade-offs: robustness vs. reward, certainty vs. windfall, and opportunity vs. robustness. These are described succinctly in section 3 and illustrated with an example in section 4. In section 5 we briefly mention three methods for combining info-gap and probabilistic models of uncertainty.

1 Info-Gap Models of Uncertainty

Our quantification of uncertainty is based on non-probabilistic information-gap models. An info-gap model is a family of nested sets. Each set corresponds to a particular degree of uncertainty, according to its level of nesting. Each element in a set represents a possible realization of the uncertain event. Info-gap models, and especially convex-set models of uncertainty, have been described elsewhere, both technologically [2] and axiomatically [4].

Uncertain quantities are vectors or vector functions. Uncertainty is expressed at two levels by info-gap models. For fixed α the set $\mathcal{U}(\alpha, \tilde{u})$ represents a degree of uncertain variability of the uncertain quantity u around the centerpoint \tilde{u} . The greater the value of α , the greater the range of possible variation, so α is called the *uncertainty parameter* and expresses the information gap between what is known (\tilde{u} and the structure of the sets) and what needs to be known for an ideal solution (the exact value of u). The value of α

is usually unknown, which constitutes the second level of uncertainty: the horizon of uncertain variation is unbounded.

Let \mathfrak{R} denote the non-negative real numbers and let S be a Banach space in which the uncertain quantities u are defined. An info-gap model $\mathcal{U}(\alpha, \tilde{u})$ is a map from $\mathfrak{R} \times S$ into the power set of S .

The basic axiom, which characterizes the representation of uncertainty by info-gap models, is that the sets of an info-gap model are *nested* by the uncertainty parameter α :

$$\mathcal{U}(\alpha, \tilde{u}) \subseteq \mathcal{U}(\alpha', \tilde{u}) \quad \text{if } \alpha \leq \alpha' \quad (1)$$

In many applications it is found that the relevant info-gap models obey specific structural axioms. The most common structural axioms are [4]:

Contraction: $\mathcal{U}(0, 0)$ is a singleton set containing its centerpoint:

$$\mathcal{U}(0, 0) = \{0\} \quad (2)$$

Translation: $\mathcal{U}(\alpha, \tilde{u})$ is obtained by shifting $\mathcal{U}(\alpha, 0)$ from the origin to \tilde{u} :

$$\mathcal{U}(\alpha, \tilde{u}) = \mathcal{U}(\alpha, 0) + \tilde{u} \quad (3)$$

where $\mathcal{U} + \tilde{u}$ means that \tilde{u} is added to each element of \mathcal{U} .

Linear expansion: info-gap models centered at the origin expand linearly:

$$\mathcal{U}(\beta, 0) = \frac{\beta}{\alpha} \mathcal{U}(\alpha, 0) \quad \text{for all } \alpha, \beta > 0 \quad (4)$$

where $\frac{\beta}{\alpha} \mathcal{U}$ means that $\frac{\beta}{\alpha}$ multiplies each element of \mathcal{U} . In some situations the linear-expansion axiom is altered to include non-linear expansion properties [7].

2 Robustness and Opportunity

2.1 A First Look

The *robustness function* expresses the greatest level of info-gap uncertainty at which failure cannot occur; the *opportunity function* is the lowest info-gap which entails the possibility of sweeping success. The robustness and opportunity functions address, respectively, the pernicious and propitious facets of uncertainty [8]

Let q be a decision vector of parameters such as design variables, time of initiation, model parameters or operational options. We can verbally express the robustness and opportunity functions as the maximum or minimum of a set of values of the uncertainty parameter α of an info-gap model:

$$\hat{\alpha}(q) = \max \{ \alpha : \begin{array}{l} \text{minimal requirements are} \\ \text{always satisfied} \end{array} \} \quad \begin{array}{l} \text{(robustness)} \\ \end{array} \quad (5)$$

$$\hat{\beta}(q) = \min \{ \alpha : \begin{array}{l} \text{sweeping success is} \\ \text{sometimes enabled} \end{array} \} \quad \begin{array}{l} \text{(opportunity)} \\ \end{array} \quad (6)$$

We can “read” eq. (5) as follows. The robustness $\hat{\alpha}(q)$ of decision vector q is the greatest value of the uncertainty parameter α for which specified minimal requirements are always satisfied. $\hat{\alpha}(q)$ expresses robustness — the degree of resistance to uncertainty and immunity against failure — so a large value of $\hat{\alpha}(q)$ is desirable. Eq. (6) states that the opportunity $\hat{\beta}(q)$ is the least level of uncertainty α which must be tolerated in order to enable the possibility of sweeping success as a result of decisions q . $\hat{\beta}(q)$ is the immunity against windfall reward, so a small value of $\hat{\beta}(q)$ is desirable. A small value of $\hat{\beta}(q)$ reflects the opportunity situation that great reward is possible even in the presence of little ambient uncertainty. The immunity functions $\hat{\alpha}(q)$ and $\hat{\beta}(q)$ are complementary and are defined in an anti-symmetric sense. Thus “bigger is better” for $\hat{\alpha}(q)$ while “big is bad” for $\hat{\beta}(q)$. The immunity functions — robustness and opportunity — are the basic decision functions in info-gap decision theory.

The robustness function in eq.(5) involves a maximization, but not of the performance or outcome of the decision. The immunity to uncertainty is maximized, while the performance is “satisficed”: a critical survival-level of performance is demanded.¹ By selecting an action q according to its robustness $\hat{\alpha}(q)$,

¹Etymologically, ‘satisfice’ is an alteration of ‘satisfy’. The word was introduced to the psychological and economic literature by Herbert Simon with the meaning: “To decide on and

the robustness function underlies a satisficing decision algorithm which optimizes the immunity to pernicious uncertainty.

The opportunity function in eq.(6) involves a minimization, however not, as might be expected, of the damage which can accrue from unknown adverse events. What is minimized is the level of uncertainty which is needed for large windfall gain to be possible. Unlike the robustness function, the opportunity function does not satisfice, it “windfalls”.² When $\hat{\beta}(q)$ is used to choose an action q , one is “windfalling” by optimizing the opportunity from propitious uncertainty in an attempt to enable highly ambitious goals or rewards.

2.2 Immunity Functions

Quite often the degree of success is assessed by a scalar reward function $R(q, u)$. The reward may be in monetary units, or it may have other dimensions expressing the performance demanded of the system. $R(q, u)$ depends on the vector q of actions or decisions as well as on an uncertain vector u whose variations are described by an info-gap model $\mathcal{U}(\alpha, \tilde{u})$, $\alpha \geq 0$. We will refer rather vaguely to u as an ‘ambient uncertainty’. It may be an outcome which depends in some way upon the decision vector q , or u may be entirely indifferent to how the decision maker acts. The uncertain u may be the essence of the outcome which the decision maker seeks (dollars of profit, or millimeters of displacement, etc.) or u may simply be an auxiliary variable of no inherent significance which nonetheless influences the overall reward.

Given a scalar reward function $R(q, u)$, the minimal requirement in eq.(5) is that the reward $R(q, u)$ be no less than a critical value r_c . Likewise, the sweeping success in eq.(6) is attainment of a “wildest dream” level of reward r_w which is much greater than r_c . Usually neither of these threshold values, r_c and r_w , is chosen irrevocably before performing the decision analysis. Rather, these parameters enable the decision maker to explore a range of options. In any case the windfall reward r_w is greater, usually much greater, than the critical reward r_c :

$$r_w > r_c \quad (7)$$

The robustness and opportunity functions of eqs.(5)

pursue a course of action that will satisfy the minimum requirements necessary to achieve a particular goal.” [11]

²While a windfall is, in its original meaning, simply something blown down by the wind, it has come to mean such a thing of value. The Oxford English Dictionary [11] gives the following quaint usage from 1705:

The grizly Boar is hunting round,
To see what Windfals may be found.

and (6) can now be expressed more explicitly:

$$\widehat{\alpha}(q, r_c) = \max \left\{ \alpha : \min_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u) \geq r_c \right\} \quad (8)$$

$$\widehat{\beta}(q, r_w) = \min \left\{ \alpha : \max_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u) \geq r_w \right\} \quad (9)$$

$\widehat{\alpha}(q, r_c)$ is the greatest level of uncertainty consistent with guaranteed reward no less than the critical reward r_c , while $\widehat{\beta}(q, r_w)$ is the least level of uncertainty which must be accepted in order to facilitate (but not guarantee) windfall as great as r_w . The complementary or anti-symmetric structure of the immunity functions is evident from eqs.(8) and (9).

The definitions of robustness and opportunity functions in eqs.(8) and (9) assume that the sets of α -values are not empty. We denote these sets as:

$$\mathcal{A}(q, r_c) = \left\{ \alpha : \min_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u) \geq r_c \right\} \quad (10)$$

$$\mathcal{B}(q, r_w) = \left\{ \alpha : \max_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u) \geq r_w \right\} \quad (11)$$

$\mathcal{A}(q, r_c)$ is the set of α -values whose least upper bound is the robustness $\widehat{\alpha}(q, r_c)$. If $\mathcal{A}(q, r_c)$ is empty then the decision is completely vulnerable — no realization of the uncertain u can lead to obtaining the demanded reward — and we define the robustness function as zero. Likewise, $\mathcal{B}(q, r_w)$ is the set of α -values whose greatest lower bound is the opportunity $\widehat{\beta}(q, r_w)$. If $\mathcal{B}(q, r_w)$ is empty then the uncertain variation entails no opportunity for windfall, and we ascribe to the opportunity function the value of infinity. That is:

$$\widehat{\alpha}(q, r_c) = 0 \quad \text{if } \mathcal{A}(q, r_c) = \emptyset \quad (12)$$

$$\widehat{\beta}(q, r_w) = \infty \quad \text{if } \mathcal{B}(q, r_w) = \emptyset \quad (13)$$

In some situations the “natural” reward requirement is that the performance function $R(q, u)$ *must not exceed* a specified value r_c , rather than being required to be no less than r_c as in eq.(8). For instance, if $R(q, u)$ represents a measure of instability of the system then a small value may be preferred rather than a large value. In this case, eq.(8) is modified so that the robustness is the greatest value of the uncertainty parameter such that the maximum reward is no *greater* than r_c :

$$\widehat{\alpha}(q, r_c) = \max \left\{ \alpha : \max_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u) \leq r_c \right\} \quad (14)$$

This is still consistent with the verbal formulation of the robustness in eq.(5). In like manner, the opportunity function is the least value of α so that the reward can possibly be as small as r_w , so that eq.(9) is modified to:

$$\widehat{\beta}(q, r_w) = \min \left\{ \alpha : \min_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q, u) \leq r_w \right\} \quad (15)$$

where r_w is less, usually much less, than r_c . The anti-symmetric relation between robustness and opportunity is retained and, as before, “bigger is better” for $\widehat{\alpha}(q, r_c)$ while “big is bad” for $\widehat{\beta}(q, r_w)$.

2.3 Preferences

The immunity functions, $\widehat{\alpha}(q, r_c)$ and $\widehat{\beta}(q, r_w)$, are the basic tools in info-gap decision theory. For given values of critical or windfall reward, r_c or r_w , each immunity function induces a preference ranking on the set of available decisions. More importantly, the immunity functions enable the decision maker to explore the desirability of different options q and different requirements, r_c and r_w , and thus to alter earlier preferences.

The robustness $\widehat{\alpha}(q, r_c)$ is the greatest level of uncertainty at which action q guarantees reward no less than r_c . As we have noted before, this means that “bigger is better” for the robustness function.³ Consequently, a decision maker will usually prefer a decision-option q over an alternative decision q' if the robustness of q is greater than the robustness of q' at the same value of critical reward r_c . We can express this preference more succinctly as:

$$q \succ q' \quad \text{if } \widehat{\alpha}(q, r_c) > \widehat{\alpha}(q', r_c) \quad (16)$$

Let \mathcal{Q} be the set of all available or feasible decision vectors q . A robust-optimal decision is one which maximizes the robustness on the set \mathcal{Q} of available q -vectors. We denote the robust-optimal action by $\widehat{q}_c(r_c)$, noting that usually the robust-optimal action depends on the critical reward. $\widehat{q}_c(r_c)$ is defined implicitly from the following optimization:

$$\widehat{\alpha}(\widehat{q}_c(r_c), r_c) = \max_{q \in \mathcal{Q}} \widehat{\alpha}(q, r_c) \quad (17)$$

It must be stressed that, nonetheless, the robustness function does not determine the decision maker’s behavior, since both $\widehat{\alpha}(q, r_c)$ and $\widehat{q}_c(r_c)$ depend on the critical reward r_c , which is a free parameter. That is, $\widehat{\alpha}(q, r_c)$ does not establish a unique preference ordering on the set \mathcal{Q} of available actions. It often happens

³This is true also when small reward is sought, as in eq.(14).

that the decision maker chooses both r_c and the optimal action $\hat{q}_c(r_c)$ in an iterative (and introspective) fashion from consideration of the robustness function. In this paper we examine methods by which the decision maker uses the robustness function, sometimes together with the opportunity function, to explore the implications of alternative scenarios.

The opportunity function $\hat{\beta}(q, r_w)$ generates a preference ranking on the available actions in a similar way, though the resulting ranking is usually different. $\hat{\beta}(q, r_w)$ is the lowest level of uncertainty which must be accepted in order to facilitate windfall reward as great as r_w . Thus, unlike the robustness function, “big is bad” for $\hat{\beta}(q, r_w)$. Consequently, a decision maker who chooses to focus on windfall opportunity will prefer a decision q over an alternative q' if the opportunity with q exceeds the opportunity with q' at the same level of reward r_w . Formally:

$$q \succ q' \quad \text{if} \quad \hat{\beta}(q, r_w) < \hat{\beta}(q', r_w) \quad (18)$$

The opportunity-optimal decision, $\hat{q}_w(r_w)$, *minimizes* the opportunity function on the set of available decisions:

$$\hat{\beta}(\hat{q}_w(r_w), r_w) = \min_{q \in \mathcal{Q}} \hat{\beta}(q, r_w) \quad (19)$$

The two preference-rankings, eqs.(16) and (18), are usually different, as are the optimal decisions $\hat{q}_c(r_c)$ and $\hat{q}_w(r_w)$. The resolution of this conflict between robustness and opportunity rankings is usually simple: most decision makers simply concentrate on enhancing their robustness while guaranteeing a specified level of performance, and thus adopt the robustness-ranking of eq.(16). The opportunity function then becomes a secondary decision-support device useful in fine-tuning ones’ choice between alternative options.

3 Trade-Offs

In this paper we discuss three types of trade-offs which face the decision maker: robustness vs. reward, certainty vs. windfall, and opportunity vs. robustness. In this section we briefly describe the first two trade-offs, and consider the third in the example in section 4.

Fig. 1 illustrates the assertion (which can be proven mathematically, [6]) that the robustness function of eq.(8), $\hat{\alpha}(q, r_c)$, decreases monotonically as the minimal required reward r_c is increased. Recalling that bigger values of $\hat{\alpha}$ are preferred over smaller ones, this expresses the trade-off between demanded-reward and immunity-to-uncertainty: if large reward is required for survival then only low immunity to uncertainty is possible [3]. In addition, as also illustrated in fig. 1, the opportunity function of eq.(9), $\hat{\beta}(q, r_w)$, increases

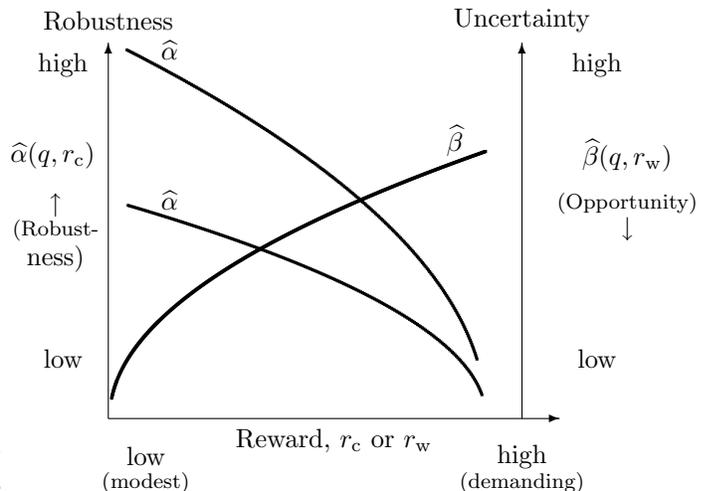


Figure 1: Two robustness curves ($\hat{\alpha}$) and one opportunity curve ($\hat{\beta}$).

monotonically with increasing wildest-dream reward r_w : sweeping success cannot be attained at low levels of ambient uncertainty. This is also a trade-off, since “big is bad” for $\hat{\beta}$. (The immunity functions in eqs.(14) and (15) are monotonic in reversed directions, though they represent the same trade-offs.)

The location of the robustness and opportunity curves on the uncertainty-vs.-reward plane, as in fig. 1, reveals one type of gambling which is expressed by these trade-offs. Consider the uppermost of the two robustness curves, $\hat{\alpha}(q, r_c)$, which falls to low and vulnerable levels of immunity only at high demanded reward. Different prior information leads to the lower robustness curve which, though still decreasing monotonically with r_c , runs more closely to the origin. The upper robustness curve represents bolder behavior than the lower curve: at any given level of demanded reward (r_c) a greater level of ambient uncertainty ($\hat{\alpha}$) is tolerable according to the upper curve. Conversely, at fixed ambient uncertainty the upper robustness curve allows greater demanded reward than the lower curve. Ascribing these two robustness curves to two different decision makers operating with different information, we can say that the lower decision maker is more sensitive to uncertainty than the upper decision maker. Equivalently, the upper curve will lead the decision maker to behavior which would look risky or rash when viewed through the strategy of the lower robustness curve. The trade-offs portrayed in fig. 1 demonstrate gambling-like behavior in the sense that the decision maker must choose a position on the immunity curves: this requires deciding how much security can be exchanged for reward, or how much reward can be relinquished in return for security. The choice reflects the extent of the decision maker’s propensity to gam-

ble, even though no concepts of chance are involved in the evaluation of the immunity functions.

The monotonic increase of the opportunity function portrayed in fig. 1 also shows a gamble-like trade-off: the decision maker’s anticipation of greater windfall reward r_w must be accompanied by acceptance of greater ambient uncertainty.

The trade-offs illustrated in fig. 1 show a particular type of coherence between the robustness and opportunity functions. As the decision maker’s expectations are reduced, whether they be for windfall reward r_w or for critical survival-level return r_c , both $\hat{\alpha}(q, r_c)$ and $\hat{\beta}(q, r_w)$ indicate a rosier picture of the effect of uncertainty. The robustness function gets larger and indicates greater immunity to failure as r_c is reduced, and the opportunity function gets smaller and shows less immunity to windfall as r_w gets smaller. $\hat{\alpha}(q, r_c)$ and $\hat{\beta}(q, r_w)$ are ‘cooperative’ or ‘sympathetic’ in the sense that they share the same trends with varying expectations r_c and r_w .

However, the variation of robustness and opportunity with varying decision q need not be sympathetic at all. A change in the choice of q which enlarges $\hat{\alpha}(q, r_c)$ need not simultaneously decrease $\hat{\beta}(q, r_w)$. These immunities may be either sympathetic or antagonistic as a function of the actions available to the decision maker.

4 Portfolio Investment

A typical simplified portfolio investment problem requires the decision maker to choose the dollar amount to buy or sell for each of a number of options, where the future values of these options are uncertain. If the (unknown) future unit value of the i th option is u_i and the dollar amount purchased or sold is q_i (positive for purchase, negative for sale), then the net change in the worth of the portfolio after the transaction is:

$$R(q, u) = \sum_{i=1}^N q_i u_i = q^T u \quad (20)$$

The question is how to choose the investment vector q given uncertainty in the future option-value vector u , as well as constraints such as budget limitations. Furthermore, one may be able to consider alternative investment portfolios: different sets of options with different uncertainties. How does one assess the relative riskiness of such investment alternatives?

Uncertainty model. For a given investment scenario we know the nominal future values of the options, $\tilde{u}_1, \dots, \tilde{u}_N$, which we combine in a nominal vector \tilde{u} . Furthermore, we may have information indi-

cating the relative degree of variability of the options, and we may also have information on the propensity for correlated or anti-correlated variation. We can use this information to formulate an ellipsoid-bound info-gap model for the uncertain variation of the option values. Let W be a real, symmetric, positive definite matrix. If we know only the relative propensities for variation of the options, without correlation data, then we choose W to be diagonal and the i th diagonal element, w_{ii} , is greater or less than unity in proportion to the tendency for the i th option to vary less or more than the norm. If we have data on the correlations between the options then we choose the eigen-structure of W to tilt the ellipsoid so as to reflect this information. In any case, an info-gap model of uncertainty is less informative than a probabilistic model (so its use is motivated by severe uncertainty) since it entails no information about likelihood or frequency of occurrence of u -vectors.

The ellipsoid-bound info-gap model for uncertain variation of the actual option-value vector u around the nominal value vector \tilde{u} is the following family of nested sets:

$$\mathcal{U}(\alpha, \tilde{u}) = \{u = \tilde{u} + v : v^T W v \leq \alpha^2\}, \quad \alpha \geq 0 \quad (21)$$

Robustness function. The decision vector q is chosen to guarantee that the change in the portfolio worth, $R(q, u)$, is no less than a minimum critical reward r_c , often called a minimum attractive rate of return (MARR) [9]. The robustness of the portfolio investment q for critical reward r_c is the greatest value of the uncertainty parameter α such that any vector u in $\mathcal{U}(\alpha, \tilde{u})$ results in a net worth $R(q, u)$ which is no less than r_c . This is precisely the robustness in eq.(8). The least reward up to uncertainty α is readily found to be:

$$\min_{u \in \mathcal{U}(\alpha, \tilde{u})} q^T u = q^T \tilde{u} - \alpha \sqrt{q^T W^{-1} q} \quad (22)$$

Equating this minimum reward to the critical value r_c and solving for the uncertainty parameter α results in the robustness:

$$\hat{\alpha}(q, r_c) = \frac{q^T \tilde{u} - r_c}{\sqrt{q^T W^{-1} q}} \quad (23)$$

if this expression is non-negative. The robustness is zero otherwise.

Robust-optimal investment. The robustness-strategy for choosing the investment portfolio is to select q to maximize $\hat{\alpha}(q, r_c)$. The choice of the portfolio is subject to many different possible constraints. For instance, some options may be accessible only if some other options are purchased as well. Or, the

quantity of an option sold may be limited by the decision maker's holdings. In addition, overall budgetary constraints may limit the total purchasing power. For simplicity, we will consider only the last constraint. Let \mathcal{Q} , the set of feasible investment vectors, be the set of all q -vectors which exactly meet the budget, Q :

$$\sum_{i=1}^N q_i = Q \quad (24)$$

where q_i is negative if option i is sold, and positive otherwise. To represent this constraint more conveniently let $\mathbf{1}$ denote the N -vector whose elements are all ones. The budget constraint is:

$$q^T \mathbf{1} = Q \quad (25)$$

The robust-optimal investment for critical reward r_c is the vector q which maximizes $\hat{\alpha}(q, r_c)$, as in eq.(17).

To simplify matters we will consider a special case: the nominal values of all the options are the same, though of course their uncertainties may be different. That is, the nominal vector \tilde{u} is:

$$\tilde{u} = u_o \mathbf{1} \quad (26)$$

where u_o is a known constant. With this simplification, the robustness in eq.(23) becomes:

$$\hat{\alpha}(q, r_c) = \frac{u_o q^T \mathbf{1} - r_c}{\sqrt{q^T W^{-1} q}} \quad (27)$$

$$= \frac{u_o Q - r_c}{\sqrt{q^T W^{-1} q}} \quad (28)$$

where in eq.(28) we have employed the budget constraint of eq.(25).

Examining eq.(28) we see that the robust-optimal investment \hat{q}_c , which maximizes $\hat{\alpha}(q, r_c)$, is the vector which minimizes $q^T W^{-1} q$ subject to the constraint in eq.(25). Using Lagrange optimization one readily finds the robust optimal investment vector to be:

$$\hat{q}_c = \frac{Q}{\mathbf{1}^T W \mathbf{1}} W \mathbf{1} \quad (29)$$

This means that, when the nominal values of the options are equal but their uncertainties are possibly different, the robust-optimal investment in the i th option is proportional to the sum of the i th row of the uncertainty shape matrix W . The meaning of this becomes particularly transparent in the further special case that W is diagonal, so that the investment in the i th option becomes:

$$\hat{q}_{c,i} = \frac{w_{ii}}{\sum_{j=1}^N w_{jj}} Q \quad (30)$$

The investment in an option is inversely proportional to its relative propensity for variation. In both cases, eq.(29) and (30), we see that the robust-optimal investment in options with equal nominal values is controlled entirely by the info-gap uncertainty, and is independent of the demanded critical return r_c .

Substituting the robust-optimal investment \hat{q}_c of eq.(29) into the robustness function of eq.(28) we obtain the maximal robustness:

$$\hat{\alpha}(\hat{q}_c, r_c) = \frac{(u_o Q - r_c) \sqrt{\mathbf{1}^T W \mathbf{1}}}{Q} \quad (31)$$

(recalling that $\hat{\alpha}(\hat{q}_c, r_c) = 0$ if the righthand side is negative.) Eq.(31) shows that $\hat{\alpha}(\hat{q}_c, r_c)$ decreases as r_c increases. This shows the trade-off between immunity to uncertainty (large $\hat{\alpha}$) and reward (large r_c): the decision maker can confidently demand great reward only in exchange for low immunity against failure due to uncertain fluctuations in the option values.

Comparing portfolios. Now consider the choice between two different portfolios, each with its own set of options, its own nominal values and its own ellipsoid-bound info-gap model of uncertainty. The two portfolios may contain different numbers of options. Assuming eq.(26) holds, separately, for each set of options, let $u_{o,1}$ be the nominal value of each option in one portfolio and let $u_{o,2}$ be the nominal value of the options in the other. Likewise, let W_1 and W_2 be the shape matrices for the two uncertainty models, whose dimensions match the number of options in the corresponding portfolio. The maximum-robustness functions for the two portfolios are each described by eq.(31), as shown schematically in fig. 2 versus the critical reward r_c .

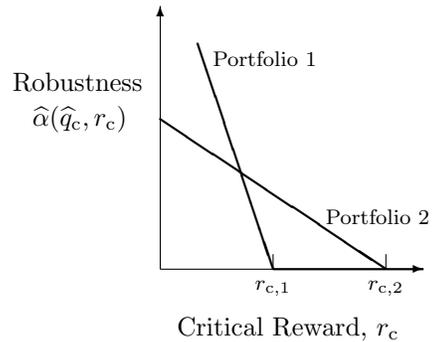


Figure 2: Robustness functions for two different portfolio investment alternatives.

Fig. 2 assists the decision maker to assess the relative riskiness of the two portfolios. Both robustness functions vanish for critical rewards in excess of $r_{c,2}$, so neither portfolio is acceptable if rewards this large are needed. For critical rewards between $r_{c,1}$ and $r_{c,2}$, and especially in the vicinity of $r_{c,1}$, the second port-

folio is the clear favorite over the first, since the second portfolio has finite robustness while the first has no immunity to uncertainty whatsoever. The riskiness of the two portfolios becomes equal when the robustness curves cross, and if values of r_c less than the intersection point are acceptable then the first alternative becomes increasingly preferable because it affords greater immunity at the same level of guaranteed return.

Opportunity function. We now consider the opportunity function $\hat{\beta}(q, r_w)$, which is the least level of uncertainty needed to sustain the possibility of reward as large as r_w , as expressed in eq.(9). The opportunity function assesses the immunity to windfall gain r_w , so a small value of $\hat{\beta}$ — low immunity to windfall — is desirable, unlike the robustness function for which a large value is needed to assure large immunity to failure. Windfalling, upon which the opportunity function is based, is different from satisficing which underlies the robustness function, though on the surface the mathematics looks quite similar.

To evaluate the opportunity function we need the greatest possible reward up to uncertainty α , which is found to be:

$$\max_{u \in \mathcal{U}(\alpha, \tilde{u})} q^T u = q^T \tilde{u} + \alpha \sqrt{q^T W^{-1} q} \quad (32)$$

whose similarity to the minimum reward in eq.(22) is evident. The opportunity function is obtained by equating this maximum to the windfall reward r_w and solving for the uncertainty parameter α , leading to:

$$\hat{\beta}(q, r_w) = \frac{r_w - q^T \tilde{u}}{\sqrt{q^T W^{-1} q}} \quad (33)$$

(or zero if this expression is negative.) This relation displays the usual trade-off between opportunity (small $\hat{\beta}$) and windfall reward (large r_w): large windfall is obtained only at the expense of accepting large ambient uncertainty.

If we impose the budget constraint of eq.(25) and if, as in eq.(26), we assume that the nominal option-values are all equal, then the opportunity function becomes:

$$\hat{\beta}(q, r_w) = \frac{r_w - u_o Q}{\sqrt{q^T W^{-1} q}} \quad (34)$$

which is similar to the robustness function of eq.(28).

Because windfalling is different from satisficing, and because opportunity is different from robustness, we can now see that optimizing $\hat{\beta}(q, r_w)$ is very different from optimizing $\hat{\alpha}(q, r_c)$. The opportunity function $\hat{\beta}(q, r_w)$ is minimized (optimized) by maximizing $q^T W^{-1} q$, while the robustness is optimized (maximized) by minimizing this same quadratic term. First

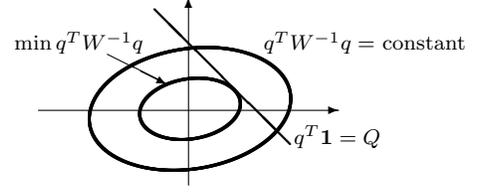


Figure 3: Schematic illustration of constrained optimization of $q^T W^{-1} q$.

of all, we obviously cannot do both optimizations simultaneously. Furthermore, the first — optimizing $\hat{\beta}(q, r_w)$ — cannot be done at all if only the budget constraint of eq.(25) is imposed. There simply is no maximum of $q^T W^{-1} q$ subject to $q^T \mathbf{1} = Q$. This is illustrated in fig. 3. No matter how large we make the quadratic term (which defines an ellipsoid) it still intersects the plane defined by the constraint. This is unlike the minimization of $q^T W^{-1} q$, which occurs when any further constriction of the ellipsoid would cause it to disconnect from the constraint plane.

In practice of course the budget limitation of eq.(25) is not the only constraint. Additional constraints become active as the investment vector q ranges further from the origin: limitations in the supply of options which can be purchased or constraints on the quantity of holdings which can be sold. Nonetheless, this example demonstrates some of the fundamental differences between windfalling with the opportunity function and satisficing with the robustness function.

Let us leave the attempt to optimize robustness and opportunity, and note that any improvement in one function is obtained at the expense of deteriorating the other. Comparing the robustness and opportunity functions in eqs.(28) and (34) we note that any change in the investment vector q which increases one will increase the other, and likewise any decrease in one function will be accompanied by a decrease in the other. However, “big is better” for $\hat{\alpha}$ while “big is bad” for $\hat{\beta}$. These immunity functions are antagonistic: either immunity can be improved only at the expense of the other. Tantalizingly, it can be proven [6] (and examples can be found which show) that robustness and opportunity *can* be sympathetic rather than antagonistic.

5 Hybrid Uncertainties

The goal of this paper has been to develop some of the implications of info-gap uncertainty for decisions with severe lack of information. We have exclusively used info-gap models to represent uncertainty. Nonetheless, one sometimes has information which is

amenable to probabilistic representation. Such information is valuable and should be exploited. Often, however, the probabilistic information is insufficient to cover all facets of the problem, in which case one can combine it with an info-gap model to create a hybrid decision algorithm. In this section we will very briefly refer to several approaches to combining probabilistic and info-gap models of uncertainty.

Three approaches have been studied in some depth [8]. One concept which has proven quite fruitful is to examine a Poisson process whose individual events entail uncertainties which are represented by an info-gap model [5, 10]. In this way we represent situations in which complicated and incompletely understood events recur Poisson-randomly in time. In a different approach we reverse the situation and consider an imperfectly known probability distribution embedded in an info-gap model [1]. Finally, a third type of hybridization is to let the info-gap uncertainty parameter α be a random variable with a known probability distribution [3].

References

1. Y. Ben-Haim, A non-probabilistic concept of reliability, *Structural Safety*, 14: 227–245 (1994).
2. Y. Ben-Haim, 1996, *Robust Reliability in the Mechanical Sciences*, Springer-Verlag, Berlin.
3. Y. Ben-Haim, 1998, Reliability of Vibrating Structures With Uncertain Inputs. *Shock and Vibration Digest*, 30: 106–113.
4. Y. Ben-Haim, 1999, Set-models of information-gap uncertainty: Axioms and an inference scheme, *Journal of the Franklin Institute*, 336: 1093–1117.
5. Y. Ben-Haim, 1999, Design certification with information-gap uncertainty. *Structural Safety*. 21:269–289.
6. Y. Ben-Haim, 2000, Robust rationality and decisions under severe uncertainty. *Journal of the Franklin Institute*, 337: 171–199.
7. Y. Ben-Haim, Robustness of model-based fault diagnosis: Decisions with set-models of uncertainty, *International Journal of Systems Science*, 31: 1518–1511.
8. Y. Ben-Haim, 2001, *Information-Gap Decision Theory: Decisions Under Severe Uncertainty*, to be published in the Academic Press Book Series on Decision and Risk.
9. W.J. Fabrycky, G.J. Thuesen and D. Verma, 1998, *Economic Decision Analysis*, 3rd ed., Prentice-Hall.
10. K.W. Hipel and Y. Ben-Haim, 1999, Decision making in an uncertain world: Information-gap modelling in water resources management, *IEEE Trans., Systems, Man and Cybernetics*, Part C: Applications and Reviews, 29: 506–517.
11. Oxford English Dictionary, 1999, 2nd edition on CD-ROM, Version 2.0.