Presented at the *Federal Reserve Bank of Atlanta*, 15 November 2002

Info-gap Analysis of Economic Policy

Yakov Ben-Haim
Itzhak Moda'i Chair in Technology and Economics
Faculty of Mechanical Engineering
Technion — Israel Institute of Technology
Haifa 32000 Israel
yakov@technion.ac.il
http://tx.technion.ac.il/~yakov

Contents

---

papers fed_res fed.tex 28.10.2002

1 Introduction

We call ourselves *Homo sapiens* because we value our ability to optimize. The question however is, 'What should we optimize?'. I will not address the value judgments entailed in choosing priorities. The focus of this paper is on a methodology for obtaining feasible optimality. Our sapience is not limited to the persistent pursuit of unattainable goals. Rather, what eons have taught the species is the lesson of balancing goals against the constraints of resources, knowledge and ability. Economics is the foremost science of decision under such constraints and, as Frank Knight taught long ago, it is the unmeasurable uncertainties which provide both motivation and justification for much economic activity.

In section 2 I will very briefly review some of the characterizations of dominant uncertainties in economics. This material is familiar, but it is striking to note how many leading economic thinkers, from so many diverse perspectives, have concurred in their critique of economic theory's treatment of uncertainty. Information-gap decision theory, which underlies this paper, provides a methodology for quantifying and managing these uncertainties. While it is not the intention of this paper to develop the micro-economics of info-gap uncertainty, it is relevant to note that such a theory is being developed, and constitutes one possible quantitative economic realization of Knight's unmeasurable uncertainty.

Sections 3–6 contain four rather different economic examples of info-gap analysis. They are fairly independent or each other, and the reader can choose among them those of particular interest, and skip the others without unduely jeopardizing the clarity of the exposition.

Section 3 begins our discussion with an example of economic modelling: the info-gap analysis of Cournot-Nash equilibrium in an oligopolistic market. The purpose of the example is to illustrate the basic decision function of info-gap theory: the robustness function for satisficing one's performance aspirations. This example also illustrates a general theorem which underlies all the examples in this paper: performance-optimization entails robustness-minimization. The implications of this trade-off between reward and robustness-to-uncertainty is central to the methodology of info-gap theory.

Section 4 is our first policy example. We consider the info-gap analysis and formulation of governmental choice of subsidy and tax rates for inducing investors to maximize their involvement in hi-tech ventures. We construct a simple micro-economic model for investment choice, which reflects the severe uncertainties in the future returns on the investment. We use this model to select tax and subsidy rates to maximize the incentive to invest.

Section 5 takes a rather different direction and discusses the optimal estimation, based on measurements, of a model for predicting the behavior of a system. We ordinarily evaluate a model in terms of its fidelity to the data: that model is optimal which best reproduces the measurements according to some measure of performance. In this example we illustrate the general theorem that performance-optimality is accompanied by total lack of immunity to uncertainty. In particular, if the basic structure of the model is uncertain, lacking perhaps non-linear terms, then the assessment of its performance in terms of its fidelity to data is unrealistically optimistic. Furthermore, sub-optimal models can display greater robustness to model-uncertainties than the performance-optimal model. Since robustness to uncertainty is an important measure of the feasibility or practicality of the model, this implies that sub-optimal models may be preferable to performance-optimal models.

Section 6 explores some connections between the robustness function of info-gap theory, and incentives which can influence decision makers when they are subject to limited liability. This discussion is based on the general theorem that robustness-to-uncertainty decreases as the decision maker's aspiration for reward increases. This trade-off supports the interpretation of the robustness curve in terms of a robustness-price of reward.

Two appendices provide succinct background material on info-gap decision theory. Appendix A summarizes the basic structure of info-gap models of uncertainty. Appendix B briefly describes the basic decision functions of info-gap theory: the robustness function for satisficing one's aspirations, and the opportunity function for ‘windfalling’ the aspiration.

2 Info-gap Uncertainty in Economics

There has been enormous progress in the economics of deficient information during the past half century. However, this progress has been based almost exclusively upon probabilistic models of uncertainty. Knight repeatedly argued that the uncertainties upon which entrepreneurial competition thrives are utterly different from probabilities:

> The uncertainties which persist as causes of profit are those which are uninsurable because there is no objective measure of the probability of gain or loss. ... Situations in regard to which business judgment must be exercised do not repeat themselves with sufficient conformity to type to make possible a computation of probability. [10, p.120]

The distinction between probabilistic risk and generic uncertainty has been widely accepted by economists. Samuelson notes that [15, pp.503–504]:

> Thinkers have naturally questioned whether the phenomena of *uncertainty* [original italics] can be usefully handled by the quasi-mathematical notions of ‘probability.’ Certain subsets of uncertainty — those dealing with risks, gambling, insurance, repetitive inventory and quality control of production, and even with tactics of repeated investing — are thought to lend themselves better to useful employment of probability procedures.

Nobody disputes that probability distributions reflect imperfect knowledge. The point is that real economic uncertainties, those which motivate the entrepreneur and are either a blessing or a bane, are often starker and sparser than is reflected in frequentist or Bayesian/subjectivist measure functions. Real economic uncertainty “is the complement of knowledge. It is the gap between what is known and what needs to be known to make correct decisions.” [12, p.1]. Uncertainty is an information gap: “the difference between the amount of information required to perform the task and the amount of information already possessed by the organization.” [6, p.5].

Shackle explains that info-gaps arise as a necessary epistemic consequence of intelligent learning [16]:

> This insufficiency of knowledge is permanent and part of the nature of things, for consciousness consists precisely in the continuous gaining of knowledge. (pp.3–4)

Consequently, Shackle continues, the enterpriser’s

> duty is to fill, with inventions and figments, the gap between what can be known and what needs to be known. When there is no such gap, there need be no enterpriser in the sense of policy-originator. (p.145)

Keynes stated the same idea somewhat differently [8]:

> The outstanding fact is the extreme precariousness of the basis of knowledge on which our estimates of prospective yield have to be made. Our knowledge of the factors which will govern the yield of an investment some years hence is usually very slight and often negligible. (p.149)
>
> Nor can we rationalise our behaviour by arguing that to a man in a state of ignorance errors in either direction are equally probable, so that there remains a mean actuarial expectation based on equi-probabilities. For it can easily be shown that the assumption of arithmetically equal probabilities based on a state of ignorance leads to absurdities. (p.152)

Hayek has also made this point, from the rather different perspective of modelling the naturally regulated behavior which is characteristic of competitive markets [7, p.77]:

> What is the problem we wish to solve when we try to construct a rational economic order? On certain familiar assumptions the answer is simple enough. *If* we possess all the relevant information, *if* we command complete knowledge of available means, the problem which remains is purely one of logic. ...
>
> This, however, is emphatically *not* the economic problem which society faces. ...
>
> The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess.

Simon has stated much the same idea, expressing it against the backdrop of modern utility theory [17, p.17]:

> Global rationality, the rationality of neoclassical theory, assumes that the decision maker has a comprehensive, consistent utility function, knows all the alternatives that are available for choice, can compute the expected value of utility associated with each alternative, and chooses the alternative that maximizes expected utility. Bounded rationality, a rationality that is consistent with our knowledge of actual human choice behavior, assumes that the decision maker must search for alternatives, has egregiously incomplete and inaccurate knowledge about the consequences of actions, and chooses actions that are expected to be satisfactory (attain targets while satisfying constraints).

March has presented the same criticism of traditional decision theory in the context of organizational decision-making [13, p.12]:

> Within decision theory, preferences are treated as important but unproblematic. A decision-maker is assumed to have preferences that are consistent, stable, and exogenous to the choice process. Observations of organizations suggest that preferences are often far from consistent, stable, or exogenous.

Economists have long been aware of the critique of probability provided by the workers cited here as well as by others. Since virtually all treatments of economic uncertainty have been probabilistic, some writers have found it necessary to provide an explanation for why these treatments can be expected to work at all. Mas-Colell, Whinston and Green suggest that the Bayesian approach, based on personal subjective probabilities, has usually underwritten these attempts "by reducing all uncertainty to risk through the use of beliefs expressible as probabilities." [14, p.207]. This, however, does not answer the critique of probability, it only sidesteps it by assuming that there is actually no distinction at all between uncertainty and risk. Arrow follows Knight in accepting the reality of the distinction, and provides a revealing explanation for why uncertainty has not been explicitly included in economic analysis. Arrow accepts in a "fundamental sense" that the "seemingly mechanical nature of the probability calculus ... [leads to its] failure to reflect the tentative, creative nature of the human mind in the face of the unknown." Arrow suggests, however, that this "seems to lead only to the conclusion that no theory can be formulated for this case." [1, p.19].

This paper is based on info-gap models of uncertainty, which are a quantification of Knight's idea of 'uncertainty', as distinct from what he called 'risk'. Appendix A presents a concise formal definition of info-gap models. However, we begin our discussion with a simple economic example.

3 Oligopolistic Competition: Satisficing and Feasible Maximization

We begin by illustrating the info-gap approach to economic analysis by considering the Cournot-Nash equilibrium of a single homogeneous commodity in an oligopolistic market. The purpose of this example is to illustrate the info-gap realization of the idea of robust satisficing. The traditional Cournot-Nash result is obtained as a non-unique equilibrium for price and output, but only by exposing each firm to extreme vulnerability about the other firms' behavior. We will see that the firm has strong motivation *not* to adopt the direct profit-maximizing strategy, but rather to maximize its robustness and to satisfice its profit with less-than-maximal profit aspiration. This example illustrates the general theorem that performance-optimization entails robustness-minimization. This has profound and far-reaching implications for decision under uncertainty.

3.1 Formulation

I will employ Dixon's formulation of the problem [4, pp.130–132] in which the demand and cost functions are:

$$P = 1 - \sum_{i=1}^{N} x_i \tag{1}$$

$$c(x_i) = c x_i \tag{2}$$

where x_i is the output of the ith firm, $c(x_i)$ is the firm's production cost, and P is the unit cost to the consumer.

Let x_{-i} be the vector of outputs of the $N-1$ firms other than firm i. Assuming the market clears, the ith firm's profit is:

$$\pi_i(x_i, x_{-i}) = x_i \left(1 - \sum_{j=1}^{N} x_j \right) - c x_i \tag{3}$$

$$= (1-c)x_i - x_i^2 - x_i \sum_{j \neq i}^{N} x_j \tag{4}$$

3.2 Robustness Function

Now I depart from the traditional Cournot analysis, to which we will return in section 3.5. We will obtain the traditional Cournot-Nash equilibrium as a special case whose practical infeasibility will be evident.

The firm indeed wants more profit rather than less, but its aspirations are tempered by the need for feasibility, for reliable decision-making. In the Cournot case, which is driven by the choices of outputs by the firms, each firm faces uncertainty about what other firms will do. Firm i will choose an output level, x_i, which **satisfices** the firm's aspiration for profit (attempts to guarantee profit no less than a specified level), while maximizing the firm's immunity to error in its anticipation of the competitors' production levels. In order to implement this I first need to define an info-gap model of uncertainty, and then define the robustness function [3].

x_{-i} is unknown to firm i, but let $\widetilde{x}_{-i}$ denote firm i's best estimate of x_{-i}. Let $\mathcal{U}(\alpha, \widetilde{x}_{-i})$ be a set of output vectors x_{-i}, containing i's nominal estimate $\widetilde{x}_{-i}$. An **info-gap model** for firm i's uncertainty about x_{-i} is a *family of nested sets* $\mathcal{U}(\alpha, \widetilde{x}_{-i})$, $\alpha \geq 0$. The value of α is unknown, and as α grows the sets become more inclusive:

$$\alpha \leq \alpha' \quad \Longrightarrow \quad \mathcal{U}(\alpha, \widetilde{x}_{-i}) \subseteq \mathcal{U}(\alpha', \widetilde{x}_{-i}) \tag{5}$$

This nesting of the uncertainty-sets embues α with its meaning of 'horizon of uncertainty'. A large α entails great variability of actual output vectors x_{-i} around firm i's nominal estimate $\widetilde{x}_{-i}$. We will see an example of an info-gap model of uncertainty shortly.

Now I'll define the **info-gap robustness function**. Firm i's aspiration for profit is $\pi_{\rm c}$. That is, i needs to earn no less than $\pi_{\rm c}$; more is fine, but not needed for 'survival'. i attempts to earn at least $\pi_{\rm c}$ by producing x_i but, since the competitions' output is unknown when x_i is chosen, the profit from x_i is also unknown since the sale price depends upon the total industry output. The robustness of action x_i is the greatest horizon of uncertainty, α, at which firm i's profit is guaranteed to be no less than its aspiration:

$$\widehat{\alpha}(x_i, \pi_{\rm c}) = \max\left\{\alpha : \min_{x_{-i}\in\mathcal{U}(\alpha,\widetilde{x}_{-i})} \pi_i(x_i, x_{-i}) \geq \pi_{\rm c}\right\} \tag{6}$$

We can 'read' this linearly from left to right: The robustness $\widehat{\alpha}$ of action x_i, given aspiration $\pi_{\rm c}$, is the maximum horizon of uncertainty α such that the lowest profit $\pi(x_i, x_{-i})$, for any realization x_{-i} of the competitors' outputs up to α, is no less than $\pi_{\rm c}$.

A number of properties of the robustness function, $\widehat{\alpha}(x_i, \pi_{\rm c})$, can be demonstrated generically. For instance, the robustness decreases with increasing aspiration: as $\pi_{\rm c}$ gets larger, the immunity to failure $\widehat{\alpha}$ gets smaller. However, we will skip generalities and proceed with an example, by choosing a specific info-gap model. The subsequent development is valid only if $N > 1$, so our results will not include monopoly as a special case.

3.3 Robustness Function: Example

I will use the ellipsoid-bound info-gap model of uncertainty, consisting of a family of nested ellipsoids of x_{-i} vectors, all centered on the nominal estimate $\widetilde{x}_{-i}$:

$$\mathcal{U}(\alpha, \widetilde{x}_{-i}) = \left\{x_{-i} = \widetilde{x}_{-i} + y : \ y^T W y \leq \alpha^2\right\}, \quad \alpha \geq 0 \tag{7}$$

$\mathcal{U}(\alpha, \widetilde{x}_{-i})$, $\alpha \geq 0$, is a family of ellipsoids of output vectors x_{-i}, each ellipsoid being centered on the nominal output $\widetilde{x}_{-i}$. This family of sets represents output-uncertainty on two levels. The first level of uncertainty is that, at fixed horizon of uncertainty α, the actual realization x_{-i} is unknown but constrained within an ellipsoid. The second level of uncertainty is that the horizon of uncertainty, α, is unknown. Plenty of uncertainty here, but no measure functions. (W is a real, symmetric, positive definite matrix which determines the shape of the ellipsoids. For most of what follows we do not actually need to choose W.)

Let us define $\mathbf{1}$ as an $N-1$ vector of ones, and $\widetilde{\pi}$ as the nominal profit of firm i which, from eq.(4), can be written:

$$\widetilde{\pi}(x_i) = (1-c)x_i - x_i^2 - (\widetilde{x}_{-i}^T \mathbf{1})x_i \tag{8}$$

With some Lagrangian analysis one finds the robustness function to be:

$$\widehat{\alpha}(x_i, \pi_{\rm c}) = \frac{\widetilde{\pi}(x_i) - \pi_{\rm c}}{x_i\sqrt{\mathbf{1}^T W^{-1}\mathbf{1}}} \tag{9}$$

$\widehat{\alpha}(x_i, \pi_{\rm c})$ is defined to be zero if this expression if negative. A negative value comes about in eq.(9) if the demanded profit, $\pi_{\rm c}$, exceeds the nominal profit, $\widetilde{\pi}$. A negative value means that even without uncertainty the aspiration cannot be achieved. In this case the robustness to uncertainty vanishes and we define $\widehat{\alpha} = 0$.

3.4 Robust-optimal Output

Recall that the robustness, $\widehat{\alpha}(x_i, \pi_{\rm c})$, is the greatest horizon of uncertainty α at which firm i is guaranteed to earn a profit no less than $\pi_{\rm c}$, regardless of what the competition does. Clearly, more robustness is better than less, so the firm should choose its output to maximize the robustness at its specified profit-aspiration. The robust-optimal output of firm i, denoted $\widehat{x}_i$, maximizes the firm's robustness function at fixed aspiration:

$$\widehat{\alpha}(\widehat{x}_i, \pi_{\rm c}) = \max_{x_i \geq 0} \widehat{\alpha}(x_i, \pi_{\rm c}) \tag{10}$$

Using eq.(9), one can readily find the robust-optimal output to be:

$$\widehat{x}_i = \sqrt{\pi_{\rm c}} \tag{11}$$

Significantly, this optimal choice for firm i is independent of the outputs of the competition. Conveniently, the optimal output is also independent of the shape of the uncertainty-ellipsoids which is determined by the matrix W in the info-gap model of eq.(7).

Given that firm i's aspiration for profit is $\pi_{\rm c}$, the firm has no motivation for choosing an output other than $\sqrt{\pi_{\rm c}}$, which maximizes the firm's immunity to the choices of all $N-1$ other firms. We now employ the logic of symmetry used in the classical Cournot-Nash reasoning. As Dixon explains [4, p.131]:

> Each firm has a similar reaction function [in our case, read: robustness function], and the Nash equilibrium occurs when each firm is on its reaction function (i.e. choosing its optimal output given the output of the other firms). There will be a symmetric and unique Cournot-Nash equilibrium which is obtained by solving the n equations (6.5) for outputs (which are all equal by symmetry).

This translates simply to the result that each firm chooses its output to maximize its robustness. Thus each firm's output is the square root of its aspiration, as in eq.(11). If we furthermore interpret the similarity of the firms to mean that they all have the same profit aspirations $\pi_{\rm c}$, (which seems no further stretched than that they all have the same reaction functions and all aspire to maximum profit), then firm i's nominal estimate of the competitions' output becomes:

$$\widetilde{x}_{-i} = \sqrt{\pi_{\rm c}}\mathbf{1} \tag{12}$$

The robustness in eq.(9) depends on $\widetilde{x}_{-i}$ through the nominal profit in eq.(8). Using eq.(12), the robustness at optimal output becomes:

$$\widehat{\alpha}(\widehat{x}_i, \pi_{\rm c}) = \frac{1 - c - \sqrt{\pi_{\rm c}}(N+1)}{\sqrt{\mathbf{1}^T W^{-1} \mathbf{1}}} \tag{13}$$

3.5 Relation to Cournot-Nash Equilibrium

We have now reached the info-gap version of the Cournot-Nash equilibrium, with each firm choosing aspiration $\pi_{\rm c}$ and output $\sqrt{\pi_{\rm c}}$. Since this output maximizes each firm's robustness at the chosen profit-aspiration, no firm has any incentive to alter its choice. This equilibrium, however, is not unique, since $\pi_{\rm c}$ is indeterminate and any positive value would produce equilibrium. In fact, we have said nothing about how the firms choose their (identical) profit aspiration $\pi_{\rm c}$. From eq.(13) we see that the optimal robustness decreases monotonically as the aspiration increases (an example of the general trade-off mentioned earlier). This is plotted in fig. 1 for $N = 4$ firms, unit production cost of $c = 0.2$, and spherical uncertainty sets expressed by choosing the shape matrix W as the identity matrix in the info-gap model of eq.(7). High or demanding profit-aspiration entails low robustness to uncertainty in the competitions' behavior. Low robustness is unfeasible since it means that the firm cannot rely upon achieving its aspiration. Conversely, the robustness is augmented by reducing the aspiration for profit. The firm is constrained to operate along this curve, unless structural modifications are considered (such as technological changes, collusion, etc.). The firm

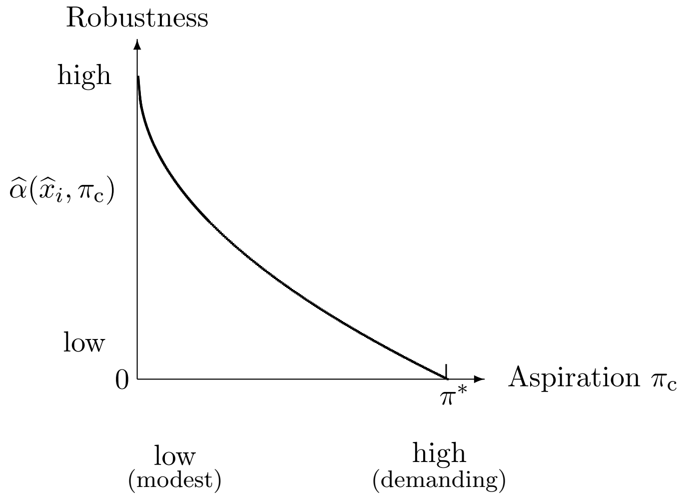

Figure 1: Optimal robustness curve, $\widehat{\alpha}(\widehat{x}_i, \pi_c)$, versus profit-aspiration π_c, eq.(13). $c = 0.2$, $N = 4$, $W = I$.

uses this quantitative robustness-vs.-profit curve to choose a value of π_c, and hence of $\widehat{x}_i$, which is a satisfactory compromise between profit and feasibility.

The profit-aspiration at which the robustness equals zero, denoted in fig. 1 by π^*, is of particular interest in this example. From eq.(13), this aspiration is seen to be:

$$\pi^* = \left(\frac{1-c}{N+1} \right)^2 \tag{14}$$

This is clearly not a feasible aspiration for a firm which wishes to plan: since $\widehat{\alpha}(\widehat{x}_i, \pi^*) = 0$ we conclude that failure to achieve this aspiration can result from infinitesimal deviation, by the competition, from firm i's estimate of the aggregate output. The firm would be advised to moderate its aspirations below the value π^* so as to increase its robustness.

From eq.(11) we see that the output induced by the aspiration π^* is precisely the classical Cournot-Nash output, [4, eq.(6.6), p.131]:

$$x_i = \frac{1-c}{N+1} \tag{15}$$

In other words, the classical Cournot-Nash equilibrium is extremely vulnerable to the assumption of identical behavior by the firms. Arbitrarily small deviations by any firm can cause all other firms to profit less than their aspiration. It is true that π^* is a high aspiration, so small short-falls may (or may not) be inconsequential, but it is also true that π^* is not a reliable choice since it is highly vulnerable to incurring such short-falls.

3.6 Implications

One prescriptive implication of this example is that firms which plan for the future are ill-advised to pursue direct profit maximization without considering the feasibility of their contemplated actions. Feasibility in this example has been measured in terms of robustness to the firm's ignorance about the other firms' intentions. (Feasibility can be measured in other ways as well in info-gap decision theory.) What the robustness function does is *maximize immunity* to uncertainty and *satisfice profit.* The Cournot-Nash solution is based on *maximizing profit,* and turns out to *minimize immunity* to uncertainty.

From a descriptive point of view, it is interesting to note that the info-gap Cournot-Nash equilibrium is not unique. The solution $x_i = \sqrt{\pi_c}$ is info-gap-Nash-stable for *any* choice of profit aspiration

π_c chosen by all firms. It seems questionable whether all firms would in fact choose the same aspiration. It is in the nature of aspirations to be highly subjective. This suggests that the analysis should be extended in two directions. First, each firm can have its own aspiration. Second, through some signalling process, each firm learns the aspirations of the competitors, either individually or in the aggregate. As this knowledge emerges, each firm can then calculate its own optimal output. (The above example would be modified at eq.(8) where the individual firm must know the value of the aggregate output of the competition, $\widetilde{x}_{-i}^T \mathbf{1}$, in order to calculate its own nominal profit, $\widetilde{\pi}$.)

4 Tax and Subsidy for Encouraging Hi-tech Investment

We now consider an example of info-gap analysis and formulation of economic policy: governmental choice of subsidy and tax rates for inducing investors to maximize their involvement in hi-tech ventures. We construct a simple micro-economic model for investment choice, and use this model to select tax and subsidy rates.

There are many crucial uncertainties in this problem which can, and should, be treated in a full info-gap analysis. For simplicity I will focus only on severely deficient knowledge about the marginal return on investment.

4.1 Traditional Approach

From traditional investment theory we know that the optimal size of investment, I, which maximizes the net present value (NPV), is the solution for I obtained by equating $K'(I)$, the marginal return on investment, to a known function [5]:

$$K'(I) = \beta + \frac{(1-\phi)\,(1-\tau P_{\mathrm{s}} \mathrm{e}^{-r})}{(1-\tau) P_{\mathrm{s}} \mathrm{e}^{-r}} \tag{16}$$

where β is the fraction of the government subsidy which the firm must return to the government, ϕ is the fraction of the investment which is covered by the subsidy, r is the interest rate, τ is the corporate tax rate, and P_{s} is the probability of success of the venture. A governmental policy maker could use a model such as this to choose β and ϕ to create incentives for maximizing investment.

4.2 Info-gap Model of Uncertainty

The critical point, in the context of this example, is that the firm's NPV-maximizing investment decision depends on knowledge of the marginal return, $K'(I)$. Since firms cannot know this function, especially in high-tech ventures of uncertain outcome, they must make their decisions in some other way than by solving eq.(16).

Furthermore, the example in section 3 illustrated the info-gap theorem that profit-maximization entails robustness-minimization. Given severe uncertainties about future returns, this is highly significant for how high-tech investors make their decisions. There is evidence that firms satisfice profit, rather than directly maximizing it. Info-gap theory provides an explicit framework for analyzing such decisions, which I will outline here in the context of the current case. Specifically, in section 4.3, I will derive a robustness function, $\widehat{\alpha}(I, \pi_c)$, which depends upon the investment I and the aspiration for profit π_c. For any level of profit to which the firm aspires, π_c, the firm chooses the investment I to make that profit-aspiration as feasible as possible. The firm does so by choosing I to maximize the robustness to uncertainty, $\widehat{\alpha}(I, \pi_c)$. Note that the profit-aspiration π_c need not be modest: the firm may aspire to great success, or may adopt less ambitious aspirations.

We begin an info-gap analysis by asking: what do we know about the uncertain elements, in this case, about the return function $K(I)$? It is from this knowledge, and the gaps in this knowledge, that we construct an info-gap model for the uncertainty.[1] There are many types of info-gap models, each suitable for different prior information [3]. We will consider a simple example.

Suppose we have a best estimate of the return function, denoted $\widetilde{K}(I)$, and suppose that we anticipate that the fractional deviation of the actual marginal return, $K'(I)$, from the nominal marginal return, $\widetilde{K}'(I)$, will be bounded, but that the value of this bound is unknown. Let us also suppose that there are no sunk costs, so the actual return on zero investment is zero. An **info-gap model** quantifies this partial knowledge of $K(I)$. An info-gap model is a family of nested sets, which in the present case is:

$$\mathcal{U}(\alpha, \widetilde{K}) = \left\{ K(I) : \ K(0) = 0, \ \left| K'(I) - \widetilde{K}'(I) \right| \leq \alpha \widetilde{K}'(I) \right\}, \quad \alpha \geq 0 \tag{17}$$

For any given value of α, we see that $\mathcal{U}(\alpha, \widetilde{K})$ is a set of return functions: those with no sunk cost whose fractional deviation from the nominal marginal return is no greater than α. For given α, we do not know which return function in $\mathcal{U}(\alpha, \widetilde{K})$ will occur. But we also do not know the value of the fractional variation, α, so **the info-gap model is a family of nested sets.** α is the **horizon of uncertainty** according to which the uncertainty-sets are nested.

4.3 Robustness Function

I have stressed that we do not know the value of α; that is, the horizon of uncertainty is unknown. It would be artificial to pick a specific value of α and consider only a single uncertainty set. We simply don't know what is the greatest fractional variation of future reality, $K'(I)$, from our anticipation, $\widetilde{K}'(I)$. However, we can answer the following question: what is the largest fractional variation at which the firm's aspiration for NPV is guaranteed to be obtained? The answer to this question is the robustness:

$$\widehat{\alpha}(I, \pi_{\rm c}) = \max \left\{ \alpha : \min_{K \in \mathcal{U}(\alpha, \widetilde{K})} \pi_{\rm npv}(K) \geq \pi_{\rm c} \right\} \tag{18}$$

where $\pi_{\rm npv}(K)$ is the NPV, whose dependence upon the future-return function $K(I)$ is stated explicitly. $\widehat{\alpha}(I, \pi_{\rm c})$ is the greatest horizon of uncertainty, α, at which the NPV (with tax and subsidy) is no less than the aspired-for profit $\pi_{\rm c}$, regardless of what return function $K(I)$ occurs. If the set of α-values in eq.(18), whose least upper bound equals the robustness, is an empty set, then the robustness is zero.

$\widehat{\alpha}(I, \pi_{\rm c})$ is the robustness to uncertainty in $K(I)$. Note that $\widehat{\alpha}(I, \pi_{\rm c})$ depends both on the size of the investment I, and on the firm's aspiration for profit $\pi_{\rm c}$. If $\widehat{\alpha}(I, \pi_{\rm c})$ is large, then the investment I is highly immune to uncertainty in the return function, while if $\widehat{\alpha}(I, \pi_{\rm c})$ is small then the investment is very vulnerable. Clearly, large values of the robustness function are better than small values. In other words, $\widehat{\alpha}(I, \pi_{\rm c})$ establishes a preference relation on possible investments. The firm will prefer I_1 over I_2 if the former is more robust to uncertainty than the latter at the same aspiration:

$$I_1 \succ I_2 \quad \text{if} \quad \widehat{\alpha}(I_1, \pi_{\rm c}) > \widehat{\alpha}(I_2, \pi_{\rm c}) \tag{19}$$

Furthermore, the firm can choose its aspiration, $\pi_{\rm c}$, to reflect more (or less) ambition, greed, caution, love or aversion for risk. A basic trade-off theorem in info-gap theory asserts that the robustness decreases monotonically as the aspiration increases: high aspirations are vulnerable to uncertainty, low aspirations are immune. In short, the firm can use the robustness function to seek the greatest **feasible** NPV. Feasibility, here, is assessed in terms of robustness to the firm's knowledge deficiency. This is a non-probabilistic assessment which is implementable in the presence of severely deficient information.

[1] By 'uncertainty' we explicitly mean unmeasurable uncertainty in Knight's sense, rather than structured, measurable probabilistic uncertainty.

Let us proceed to the evaluation of the robustness function. A standard expression for the NPV with tax and subsidy is [5]:

$$\pi_{\text{npv}}(K) = (1-\tau)P_{\text{s}}\mathrm{e}^{-r}K(I) + g(\phi,\beta)I \tag{20}$$

where I have defined the function:

$$g(\phi,\beta) = -\left[(1-\phi)\left(1-P_{\text{s}}\tau\mathrm{e}^{-r}\right) + \beta P_{\text{s}}(1-\tau)\mathrm{e}^{-r}\right] \tag{21}$$

which is always negative.

In order to evaluate the robustness, eq.(18), we need an expression for the lowest possible NPV, up to horizon of uncertainty α. This minimum is:

$$\min_{K\in\mathcal{U}(\alpha,\widetilde{K})} \pi_{\text{npv}}(K) = g(\phi,\beta)I + (1-\tau)P_{\text{s}}(1-\alpha)\mathrm{e}^{-r}\widetilde{K}(I) \tag{22}$$

The robustness, $\widehat{\alpha}(I,\pi_{\text{c}})$, is the greatest horizon of uncertainty at which an NPV no less than π_{c} is guaranteed. Thus $\widehat{\alpha}(I,\pi_{\text{c}})$ is the greatest value of α at which the minimum in eq.(22) just equals π_{c}. Equating eq.(22) to π_{c} and solving for α yields the robustness:

$$\widehat{\alpha}(I,\pi_{\text{c}}) = 1 + \frac{\mathrm{e}^{r}}{(1-\tau)P_{\text{s}}}\frac{g(\phi,\beta)I-\pi_{\text{c}}}{\widetilde{K}(I)} \tag{23}$$

unless this expression is negative, in which case the robustness is zero. Note that the robustness decreases as the aspiration increases, which is the basic trade-off between aspiration and feasibility mentioned before. Any realization of $K(I)$ up to uncertainty $\widehat{\alpha}(I,\pi_{\text{c}})$ will result in an NPV no less than π_{c}. For greater horizons of uncertainty there are return functions which cause a shortfall from the aspired NPV.

4.4 Optimal Investment and Policy Implications

The optimal investment, if there is one, maximizes the robustness for specified aspiration. Differentiating $\widehat{\alpha}(I,\pi_{\text{c}})$ with respect to I and equating to zero yields the following relation which is solved to find the optimal investment:

$$\frac{\widetilde{K}'(I)}{\widetilde{K}(I)} = \frac{g(\phi,\beta)}{g(\phi,\beta)I-\pi_{\text{c}}} \tag{24}$$

Whether or not the solution, for I, of this relation yields a maximum robustness, depends on the functions involved. We will consider a special case in which $\widehat{\alpha}(I,\pi_{\text{c}})$ has an unconstrained maximum. Let the nominal return function, $\widetilde{K}(I)$, be positive, increasing, with decreasing marginal return: $\widetilde{K}(I)>0$, $\widetilde{K}'(I)>0$, $\widetilde{K}''(I)<0$. These conditions assure that the function on the lefthand side of eq.(24), $\widetilde{K}'/\widetilde{K}$, decreases as I increases. Also, let $\Lambda(I)$ denote the function on the righthand side of eq.(24). Let us suppose that I and π_{c} are non-negative. Since $g(\phi,\beta)$ is negative, we see that $\Lambda(I)$ decreases as I increases. We will assume that $\widetilde{K}'(I)/\widetilde{K}(I)$ has more negative slope than $\Lambda(I)$, as in fig. 2. The value of I at which these curves cross maximizes the robustness. We will denote this value $\widehat{I}$, or sometimes $\widehat{I}(\pi_{\text{c}})$ to denote its dependence upon the profit-aspiration π_{c}. $\widehat{\alpha}(\widehat{I},\pi_{\text{c}})$ is an unconstrained maximum robustness for profit-aspiration π_{c}.

To re-iterate, the value of I which satisfies eq.(24) yields the greatest possible immunity to uncertainty in the return function $K(I)$, while attempting to achieve an NPV no less than π_{c}. This relation is clearly different from eq.(16), which is a direct optimization of the NPV (even if in eq.(16) we use the known nominal return $\widetilde{K}(I)$ rather than the unknown actual return $K(I)$). The difference between eqs.(16) and (24) arises from the difference in approach. Eq.(24) is the result of satisficing

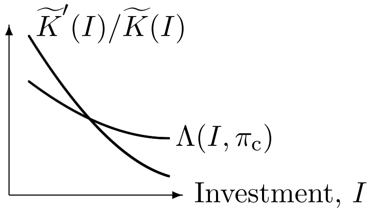

Figure 2: $\widetilde{K}'(I)/\widetilde{K}(I)$ and $\Lambda(I)$ vs. I.

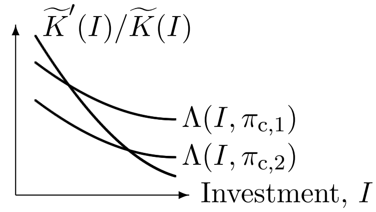

Figure 3: $\widetilde{K}'(I)/\widetilde{K}(I)$ and $\Lambda(I)$ vs. I. Demonstration that $\widehat{I}$ increases with aspiration π_c: $\pi_{c,1} < \pi_{c,2}$.

the reward at the level π_c, while seeking the investment which makes the achievement of this reward as immune as possible to fluctuation in the unknown future return function.

We can immediately draw some policy conclusions. Note that $\Lambda(I)$ decreases as π_c increases. Thus, for $\pi_{c,2} > \pi_{c,1}$, we see that the curve $\Lambda(I,\pi_{c,2})$-vs.-I lies below the curve $\Lambda(I,\pi_{c,1})$-vs.-I, as shown in fig. 3. Recall that the intersection of $\Lambda(I)$ with $\widetilde{K}'/\widetilde{K}$ determines the robustness-maximizing investment, $\widehat{I}$. Hence we conclude that the robust-optimal investment increases as the aspiration increases.

Also, note the following derivatives of the righthand side of eq.(24):

$$\frac{\partial}{\partial\phi}\left(\frac{g(\phi,\beta)}{g(\phi,\beta)I-\pi_c}\right) = \frac{-\pi_c\,(1-P_s\tau e^{-r})}{(g(\phi,\beta)I-\pi_c)^2} < 0 \tag{25}$$

$$\frac{\partial}{\partial\beta}\left(\frac{g(\phi,\beta)}{g(\phi,\beta)I-\pi_c}\right) = \frac{\pi_c P_s(1-\tau)e^{-r}}{(g(\phi,\beta)I-\pi_c)^2} > 0 \tag{26}$$

From eq.(25) we see that an increase in ϕ, the fractional contribution of the government, depresses the $\Lambda(I)$-vs.-I curve, and causes an increase in the size of the firm's robust-optimal investment. Likewise, we learn from eq.(26) than an increase in β, the fraction of the government subsidy which must be paid back, raises the $\Lambda(I)$-vs.-I curve, and causes a decrease in the size of the optimal investment. That is, denoting the info-gap-robust-optimal investment by $\widehat{I}(\pi_c,\phi,\beta)$, we have learned that:

$$\frac{\partial\widehat{I}}{\partial\pi_c} > 0, \quad \frac{\partial\widehat{I}}{\partial\phi} > 0, \quad \frac{\partial\widehat{I}}{\partial\beta} < 0 \tag{27}$$

The direction of variation of the info-gap-optimal investment with ϕ and β is the same as the direction of variation of the direct optimization in eq.(16). Of course the specific solutions are different, which is the important policy conclusion. If in fact firms maximize *feasible* profit aspirations by robustly satisficing their profit aspirations, then the current analysis is the right direction for the policy-maker to pursue.

There is no classical analog of the variation of $\widehat{I}$ with π_c, since π_c arises here in the context of profit-satisficing aspirations rather than profit-maximization.

5 Estimation

5.1 Optimal Estimation

We begin by formulating a fairly typical framework for optimal estimation of a model for predicting the observable behavior of a system. We then consider an example.

Let y_i be a vector of measurements of the system in state i, for $i = 1, \ldots, N$. Let $f_i(q)$ denote the model-prediction of the system in state i, which should match the measurements if the model is good. The vector q, containing real and linguistic variables, denotes the parameters and properties of the model which can be modified to bring the model into agreement with the measurements. We will denote the set of measurements by $Y = \{y_1, \ldots, y_N\}$ and the set of corresponding model-predictions by $F(q) = \{f_1(q), \ldots, f_N(q)\}$.

The overall performance of the predictor is assessed by a function $R[Y, F(q)]$. For example, this might be a mean-squared prediction error:

$$R[Y, F(q)] = \frac{1}{N} \sum_{i=1}^{N} \| f_i(q) - y_i \|^2 \tag{28}$$

An optimal model, q^*, minimizes the performance-measure:

$$R[Y, F(q^*)] = \min_q R[Y, F(q)] \tag{29}$$

5.2 Uncertainty and Robustness

The model $f_i(q)$ is undoubtedly wrong, perhaps fundamentally flawed in its structure. There may be basic mechanisms which act on the system but which are not represented by $f_i(q)$. Let us denote more general models, some of which may be more correct, by:

$$\phi_i = f_i(q) + u_i \tag{30}$$

where u_i represents the unknown corrections to the original model, $f_i(q)$. We have very little knowledge about u_i; if we had knowledge of u_i we would most likely include it in $f_i(q)$. So, let us represent the unknown variation of possible models with an info-gap model of uncertainty:

$$\phi_i \in \mathcal{U}(\alpha, f_i(q)), \quad \alpha \geq 0 \tag{31}$$

The centerpoint of the info-gap model, $f_i(q)$, is the known model, parameterized by q. The horizon of uncertainty, α, is unknown. This info-gap model is a family of nested sets of models. These sets of models becomes ever more inclusive as the horizon of uncertainty increases.

As before, the model-prediction of the system output in state i is $f_i(q)$, and the set of model-predictions is denoted $F(q) = \{f_1(q), \ldots, f_N(q)\}$. More generally, the set of model-predictions with unknown terms $u_1, \ldots, u_N$ is denoted $F_u(q) = \{f_1(q) + u_1, \ldots, f_N(q) + u_N\}$.

We wish to choose a model, q, for which the performance index, $R[Y, F_u(q)]$, is small. Let $r_{\rm c}$ represent an acceptably small value of this index. We would be willing, even delighted, if the prediction-error is smaller, but an error larger than $r_{\rm c}$ would be unacceptable.

The **robustness to model-uncertainty,** of model q with error-aspiration $r_{\rm c}$, is the greatest horizon of uncertainty, α, within which all models provide prediction-error no greater than $r_{\rm c}$:

$$\widehat{\alpha}(q, r_{\rm c}) = \max \left\{ \alpha : \max_{\substack{\phi_i \in \mathcal{U}(\alpha, f_i(q)) \\ i=1, \ldots, N}} R[Y, F_u(q)] \leq r_{\rm c} \right\} \tag{32}$$

When $\widehat{\alpha}(q, r_{\rm c})$ is large, the model $f_i(q)$ may err fundamentally to a great degree, without jeopardizing the accuracy of its predictions; the model is robust to info-gaps in its formulation. When $\widehat{\alpha}(q, r_{\rm c})$ is small, then even small errors in the model result in unacceptably large prediction errors.

Let q^* be an optimal model, which minimizes the prediction-error as defined in eq.(29), and let $r_{\rm c}^*$ be the corresponding optimal prediction error: $r_{\rm c}^* = R(Y, F(q^*))$. By using model q^*, we can achieve prediction error as small as $r_{\rm c}^*$, and no value of q can produce a model $f_i(q)$ which performs better.

However, it can be proven[2] that the robustness to model-uncertainty, of this optimal model, is zero:

$$\widehat{\alpha}(q^*, r_c^*) = 0 \tag{33}$$

This is a special case of the general result that, by optimizing the performance, one minimizes the robustness to info-gaps. By optimizing the performance of the model-predictor, $f_i(q)$, we make this predictor maximally sensitive to errors in the basic formulation of the model.

One implication of this general theorem is that eq.(33) is a special case of the following proposition. For any q, let $r_c = R[Y, F(q)]$ be the prediction-error of model $f_i(q)$. Then it can be shown that:

$$\widehat{\alpha}(q, r_c) = 0 \tag{34}$$

That is, the robustness of *any* model, $f_i(q)$, to uncertainty in the structure of that model, is precisely equal to zero, if the error-aspiration r_c equals the value of the performance function of that model. No model can be relied upon to perform at the level indicated by its performance function, if that model is subject to errors in its structure or formulation. $R[Y, F(q)]$ is an unrealistically optimistic assessment of model $f_i(q)$, unless we have reason to believe that no auxiliary uncertainties lurk in the mist of our ignorance.

5.3 Example

A simple example will illustrate the previous general discussion.

We begin by formulating a **mean-squared-error estimator** for a 1-dimensional linear model. The measurements y_i are scalars, and the model to be estimated is:

$$f_i(q) = iq \tag{35}$$

The performance function is the mean-squared error between model and measurements, eq.(28), which becomes:

$$R[Y, F(q)] = \frac{1}{N}\sum_{i=1}^{N}(iq - y_i)^2 \tag{36}$$

$$= \underbrace{\frac{1}{N}\sum_{i=1}^{N} y_i^2}_{\eta_2} - 2q \underbrace{\frac{1}{N}\sum_{i=1}^{N} i y_i}_{\eta_1} + q^2 \underbrace{\frac{1}{N}\sum_{i=1}^{N} i^2}_{\eta_0} \tag{37}$$

which defines the quantities η_0, η_1 and η_2. The optimal model defined in eq.(29), which minimizes the mean-squared error, is:

$$q^* = \frac{\eta_1}{\eta_0} \tag{38}$$

Now we introduce **uncertainty in the model.** The model which is being estimated is linear in the 'time' or 'sequence' index i: $f_i = iq$. How robust is the performance of our estimator, to modification of the structure of this model? That is, how much can the model err in its basic structure, without jeopardizing its predictive power?

Consider the simplest possible non-linear modification:

$$\phi_i = iq + i^2 u \tag{39}$$

[2] The main assumption needed in the proof is that the performance function, $R(Y, F_u(q))$, is 'upper unsatiated' in the info-gap:

$$\alpha < \alpha' \implies \max_{\substack{\phi_i \in \mathcal{U}(\alpha, f_i(q)) \\ i=1,\ldots,N}} R(Y, F_u(q)) < \max_{\substack{\phi_i \in \mathcal{U}(\alpha', f_i(q)) \\ i=1,\ldots,N}} R(Y, F_u(q))$$

This means that, as the horizon of uncertainty grows, the maximum of the performance function strictly increases.

where the value of u is unknown. The uncertainty in the quadratic model is represented by the info-gap model:

$$\mathcal{U}(\alpha, iq) = \left\{ \phi_i = iq + i^2 u : \ |u| \le \alpha \right\}, \quad \alpha \ge 0 \tag{40}$$

The robustness of nominal model $f_i(q)$, with performance-aspiration r_c, is the greatest value of the horizon of uncertainty α at which the mean-squared error of the prediction is no greater than r_c for any model in $\mathcal{U}(\alpha, iq)$:

$$\widehat{\alpha}(q, r_c) = \max \left\{ \alpha : \ \max_{|u| \le \alpha} R[Y, F_u(q)] \le r_c \right\} \tag{41}$$

The mean-squared error of a model with non-linear term $i^2 u$ is:

$$R[Y, F_u(q)] = \frac{1}{N} \sum_{i=1}^{N} (iq + i^2 u - y_i)^2 \tag{42}$$

$$= \underbrace{\frac{1}{N} \sum_{i=1}^{N} (iq - y_i)^2}_{\xi_2} - 2u \underbrace{\frac{1}{N} \sum_{i=1}^{N} i^2 (iq - y_i)}_{\xi_1} + u^2 \underbrace{\frac{1}{N} \sum_{i=1}^{N} i^4}_{\xi_0} \tag{43}$$

which defines ξ_0, ξ_1 and ξ_2.

Some manipulations show that the maximum mean-squared error, for all non-linear models up to horizon of uncertainty α, is:

$$\max_{|u| \le \alpha} R[Y, F_u(q)] = \xi_2 + 2\alpha |\xi_1| + \alpha^2 \xi_0 \tag{44}$$

Referring to eq.(41), the robustness, to an unknown quadratic non-linearity $i^2 u$, of the linear model $f_i(q)$, is the greatest value of α at which this maximum error is no greater than r_c.

First we note that the robustness is zero if r_c is small:

$$\widehat{\alpha}(q, r_c) = 0, \quad r_c \le \xi_2 \tag{45}$$

This is because, if $r_c \le \xi_2$, then $\max R$ in eq.(44) exceeds r_c for any positive value of α. One implication of eq.(45) is that some non-linear models have prediction errors in excess of ξ_2. *If* it is required that the fidelity between model and measurement be as good as or better than ξ_2, *then* no modelling errors of the quadratic type represented by the info-gap model of eq.(40) can be tolerated. Comparing eqs.(36) and (43) we see that $\xi_2 = R[Y, F(q)]$, that is, ξ_2 is the mean-squared error of the linear predictor. This means that there is no robustness to model-uncertainty, if the performance-aspiration r_c is stricter or more exacting than the performance of the nominal, linear model.

For $r_c \ge \xi_2$, the robustness is obtained by equating the righthand side of eq.(44) to r_c and solving for α:

$$\widehat{\alpha}(q, r_c) = \frac{|\xi_1|}{\xi_0} \left(-1 + \sqrt{1 + \frac{r_c - \xi_2}{\xi_1^2}} \right), \quad \xi_2 \le r_c \tag{46}$$

This relation is plotted schematically in fig. 4, which shows the robustness to model-uncertainty against the aspiration for prediction error, r_c. The robustness increases as greater error is tolerated. Two curves are shown, one for the optimal linear model, q^* in eq.(38), whose mean-squared error r_c^* is the lowest obtainable with any linear model. Any other model, such as q', has a greater mean-squared error $r_c' = R[Y, F(q')]$, so $r_c' > r_c^*$. However, the best performance (the smallest r_c-value) with each of these models, q' and q^*, has no robustness to model-uncertainty. More importantly, the robustness curves can cross at a higher value of r_c (corresponding to lower aspirations for prediction-fidelity). If prediction-error r_c° is tolerable, then the sub-optimal model q' is more robust than, and hence preferable over, the mean-squared 'optimal' model q^*, at the same performance-aspiration.

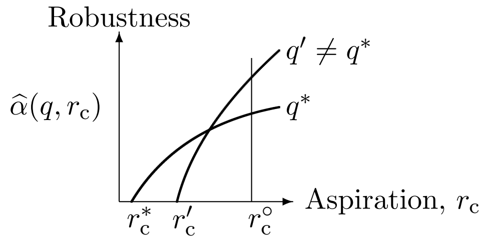

Figure 4: Robustness vs. prediction-aspiration, eq.(46).

In summary, we have established the following conclusions.

First, the optimal model, $f_i(q^*)$, has no immunity to error in the basic structure of the model. The model $f_i(q^*)$, which minimizes the mean-squared discrepancy between measurement and prediction, has zero robustness to modelling errors at its nominal prediction-fidelity, r_c^*.

Second, this is actually true of *any* model, $f_i(q')$. The value of its mean-squared error is precisely r_c' which, as in fig. 4, has zero robustness.

Third, the robustness curves of alternative linear models can cross, as in fig. 4. This shows that a sub-optimal model such as $f_i(q')$ can be more robust to model-uncertainty than the mean-squared optimal model $f_i(q^*)$, when these models are compared at the same aspiration for fidelity between model and measurement, r_c° in the figure.

6 Liability and Robustness

The examples in the previous sections have illustrated, $\widehat{\alpha}(\widehat{q}, \pi_c)$, the robustness of the most robust action, $\widehat{q}(\pi_c)$, which can be selected in an attempt to achieve reward-aspiration π_c. A fundamental theorem of info-gap theory, observed in the previous sections, asserts that this curve decreases monotonically: greater aspiration, π_c, is associated with lower robustness to uncertainty, $\widehat{\alpha}(\widehat{q}, \pi_c)$. One can view the (absolute value of the) slope of this curve as the robustness-price of reward-aspiration. One implication of this, which we will explore in a tentative and preliminary manner in this section, is that info-gap robustness analysis can be used to formulate and evaluate legislation which establishes liability.[3] We begin with a preliminary, intuitive, and hence imprecise discussion of four ways in which this can be realized. We then consider three examples.

6.1 Preliminary Discussion

1. Increased liability induces an agent to move up the robustness curve, to greater $\widehat{\alpha}$ at lower π_c, fig. 5. The robustness curve, $\widehat{\alpha}$-vs.-π_c, is an operating curve on which the decision maker is free to choose an operating point by examining the trade-offs involved. As the agent's liability becomes greater, there is an increased incentive to prefer greater robustness to uncertainty because the cost of failure is enhanced. In other words, as liability is imposed, decision makers will tend to move up their robustness curves. Of course, the psychology of human motivation is sufficiently complex to cause precisely the reverse impact in some situations: greater liability may drive a decision maker to highly ambitious and drastic measures in an attempt to cover accumulated obligations. The transition between these responses to increased liability is evidently related, at least in part, to the agent's perceptions of what constitutes "the long run" and whether or not "the long run" is even part of the agent's universe of discourse.

[3]The basic idea for this section was suggested by Dr. Doug Noonan, School of Public Policy, Georgia Institute of Technology.

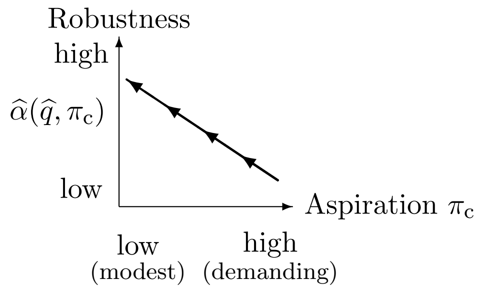

Figure 5: Increased liability induces an agent to move up the robustness curve.

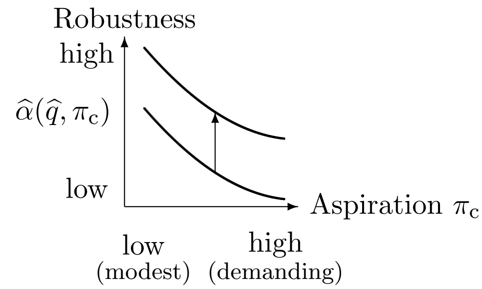

Figure 6: Increased liability induces an agent to prefer a high robustness curve over a low robustness curve.

2. Increased liability induces an agent to prefer a high robustness curve over a low robustness curve, fig. 6. In some situations the decision maker faces two different types of options, each with its own robustness curve. For instance, the decision maker may need to choose between two different projects, only one of which can be launched. Imposing liability on the agent will tend to induce a preference for the more robust option, even though other considerations, not represented by the robustness curve, may mitigate against the more robust alternative.

3. Liability legislation becomes more effective as the robustness curve becomes less steep, fig. 7. As explained in item 1, increased liability induces the decision maker to move up the robustness curve. When the robustness curve is very steep, a small change in aspiration, π_c, produces a large change in robustness, $\widehat{\alpha}$. The agent's robust-optimal action, $\widehat{q}(\pi_c)$, depends upon π_c, so only small behavioral changes result from increased liability when the robustness curve is steep.

On the other hand, when the slope of the robustness curve is gentle, liability legislation is effective in creating incentives for substantial behavioral changes. A large change in π_c is needed in order to produce even a moderate change in $\widehat{\alpha}$. Hence, since the action, $\widehat{q}(\pi_c)$, depends upon π_c, large behavioral changes result from increased liability as the decision maker is induced to move up the robustness curve.

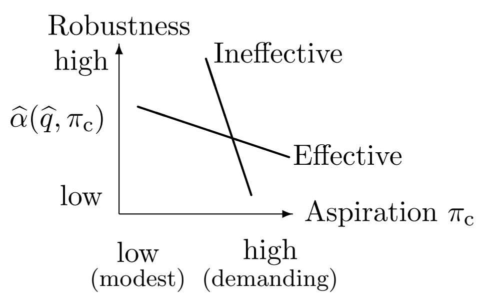

Figure 7: Liability legislation becomes more effective as the robustness curve becomes less steep.

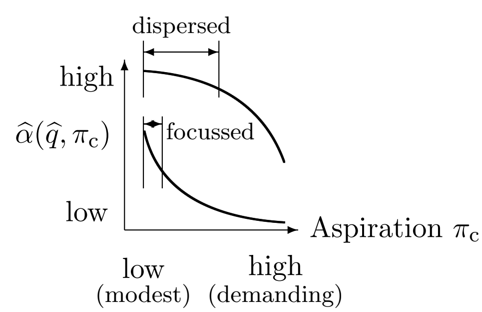

Figure 8: Positive robustness curvature induces focussed behavior; negative curvature induces dispersion.

4. Liability and positive robustness curvature induce focussed behavior; negative curvature induces dispersion, fig. 8. Robustness curves are necessarily monotonically decreasing, so positive curvature, if it is strong enough, will result in a narrow range of aspirations over which the robustness increases rapidly to substantial values, as in the lower curve in fig. 8. Agents' aspirations, and the associated actions, will tend to be focussed in the narrow region of π_c-values which have substantial robustness. On the other hand, a strongly negative robustness curvature, like the upper curve in fig. 8, will have substantial robustness over a large range of aspirations, resulting in dispersion of agents' actions.

This discussion so far has been qualitative and imprecise. We now consider a few simple but quantitative illustrations.

6.2 Example

We return to the oligopolistic example of section 3 to examine the influence of limited liability on decision makers who satisfice profit and maximize robustness. In particular, we will illustrate the phenomenon demonstrated in fig. 5.

First we define the Heaviside function:

$$h(x) = \begin{cases} 0, & x < 0 \\ 1, & x \geq 0 \end{cases} \tag{47}$$

From eq.(3), $\pi_i(x_i, x_{-i})$ denotes the profit function for firm i, where x_i is the firm's output and x_{-i} is the vector of outputs of the other firms. We now introduce limited liability by modifying this as follows:

$$\pi_i^\ell(x_i, x_{-i}) = \pi_i(x_i, x_{-i}) h\left[\pi_i(x_i, x_{-i}) - \pi_\ell\right] + \pi_\ell \left(1 - h\left[\pi_i(x_i, x_{-i}) - \pi_\ell\right]\right) \tag{48}$$

The reward is the original oligopolistic profit $\pi_i(x_i, x_{-i})$ provided that this quantity is no less than π_ℓ; otherwise the reward is π_ℓ. That is, the firm's liability is for providing acceptably large reward, unless the earned reward, $\pi_i(x_i, x_{-i})$, is below the threshold π_ℓ, in which case somebody (like the government) steps in and pays the difference to bring the earnings up to the value π_ℓ.

The robustness of firm i's output x_i, given profit-aspiration π_{c}, and uncertainty in the competitions' outputs x_{-i}, is:

$$\widehat{\alpha}^\ell(x_i, \pi_{\mathrm{c}}) = \max\left\{\alpha : \min_{x_{-i} \in \mathcal{U}(\alpha, \widetilde{x}_{-i})} \pi_i^\ell(x_i, x_{-i}) \geq \pi_{\mathrm{c}}\right\} \tag{49}$$

which is the modified version of eq.(6).

Using the ellipsoid-bound info-gap model for uncertainty in firm i's assessment of the outputs of the other $N-1$ firms, eq.(7), we find the robustness function to be:

$$\widehat{\alpha}^\ell(x_i, \pi_{\mathrm{c}}) = \begin{cases} \infty, & \pi_{\mathrm{c}} < \pi_\ell \\ \dfrac{\widetilde{\pi}_i(x_i) - \pi_{\mathrm{c}}}{x_i \sqrt{\mathbf{1}^T W^{-1} \mathbf{1}}}, & \pi_\ell \leq \pi_{\mathrm{c}} \leq \widetilde{\pi}_i(x_i) \\ 0, & \widetilde{\pi}_i(x_i) < \pi_{\mathrm{c}} \end{cases} \tag{50}$$

where $\widetilde{\pi}_i(x_i)$ is the projected profit of firm i based on the firm's nominal estimate of the other firms' outputs, eq.(8). If the firm's profit-aspiration is lower than the liability limit, $\pi_{\mathrm{c}} < \pi_\ell$, then the robustness is infinite since the desired reward is guaranteed. At the other extreme, if the aspiration exceeds the nominal reward, $\widetilde{\pi}_i(x_i) < \pi_{\mathrm{c}}$, then the robustness is zero since the aspiration can not be realized even in the absence of uncertainty (assuming $\widetilde{\pi}_i(x_i) > \pi_\ell$). In the intermediate range of aspirations, the robustness decreases as the aspiration increases.

As in the original formulation, the output which maximizes firm i's robustness, given profit-aspiration π_{c}, is:

$$\widehat{x}_i^\ell = \sqrt{\pi_{\mathrm{c}}} \tag{51}$$

provided that the robustness is bounded and positive for this value of π_{c}.

For this robust-optimal output, the nominal assessment of the competitions' outputs becomes $\widetilde{\pi}_{-i}(\widehat{x}_i^\ell) = (1 - c - \widetilde{x}_{-i}^T \mathbf{1})\sqrt{\pi_\mathrm{c}} - \pi_\mathrm{c}$. Thus the info-gap-optimal robustness becomes:

$$\widehat{\alpha}^\ell(\widehat{x}_i^\ell, \pi_\mathrm{c}) = \begin{cases} \infty, & \pi_\mathrm{c} < \pi_\ell \\ \dfrac{1 - c - \widetilde{x}_{-i}^T\mathbf{1} - 2\sqrt{\pi_\mathrm{c}}}{\sqrt{\mathbf{1}^T W^{-1}\mathbf{1}}}, & \pi_\ell \le \pi_\mathrm{c} \le \left(\dfrac{1 - c - \widetilde{x}_{-i}^T\mathbf{1}}{2}\right)^2 \\ 0, & \left(\dfrac{1 - c - \widetilde{x}_{-i}^T\mathbf{1}}{2}\right)^2 < \pi_\mathrm{c} \end{cases} \tag{52}$$

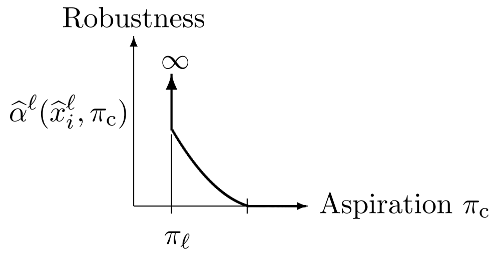

Figure 9: Optimal robustness curve, eq.(52), showing discontinuous slope.

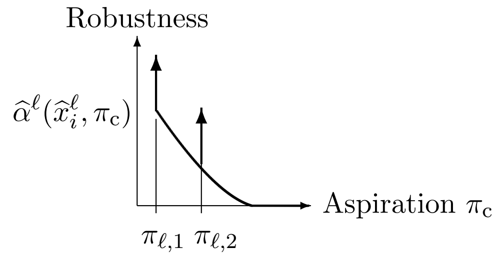

Figure 10: Two optimal robustness curves, eq.(52), for different liability limits.

The robustness curve in eq.(52) has two slope-discontinuities, as shown schematically in fig. 9. We now examine this figure, using the idea behind fig. 5, to explain the impact of limited liability. The effect of the limited liability on low-aspiration firms, those for whom $\pi_\mathrm{c} < \pi_\ell$, is to raise their aspirations to the liability limit, π_ℓ. Likewise, agents whose aspirations are only slightly above π_ℓ will tend to move their aspirations downward to π_ℓ to attain the unbounded robustness offered by the liability limitation. In the latter case the firm's robust-optimal output, $\widehat{x}_i^\ell(\pi_\mathrm{c})$ in eq.(51), depends upon the aspiration, which illustrates the behavioral impact of limited liability.

This example can also be interpreted in terms of the effect of a *change* in liability, as illustrated in fig. 10. Consider a firm whose liability is initially $\pi_{\ell,1}$. If the firm's aspiration, π_c, is greater than $\pi_{\ell,1}$ then there is an incentive to reduce the aspiration from π_c to $\pi_{\ell,1}$, as we have already explained. However, if the liability-limit is increased to $\pi_{\ell,2}$, meaning that the agent has less liability, the agent now faces an incentive to increase the aspiration to at least $\pi_{\ell,2}$.

6.3 Example

We continue by further modifying the oligopolistic example of section 6.2. We now introduce a 'balanced liability' in the form of a tax or insurance premium. Let $\pi_i(x_i, x_{-i})$ denote the original profit function from eq.(3), and modify $\pi_i^\ell(x_i, x_{-i})$ eq.(48) to include not only the liability-limitation π_ℓ, but also a proportional tax of τ dollars on each increment of profit of size π_ℓ:

$$\pi_i^L(x_i, x_{-i}) = \left(1 - \frac{\tau}{\pi_\ell}\right)\pi_i(x_i, x_{-i}) h\left[\pi_i(x_i, x_{-i}) - \pi_\ell\right] + \pi_\ell\left(1 - h\left[\pi_i(x_i, x_{-i}) - \pi_\ell\right]\right) \tag{53}$$

We will assume that $\tau < \pi_\ell$. If $\tau < 0$ then τ represents a subsidy rather than a tax.

Proceeding as in the derivation of eq.(50), we find the robustness function to be:

$$\widehat{\alpha}^L(x_i, \pi_\mathrm{c}) = \begin{cases} \infty, & \pi_\mathrm{c} < \pi_\ell \\ \dfrac{\widetilde{\pi}_i(x_i) - \pi_\mathrm{c}\left(1 - \frac{\tau}{\pi_\ell}\right)^{-1}}{x_i\sqrt{\mathbf{1}^T W^{-1}\mathbf{1}}}, & \pi_\ell \le \pi_\mathrm{c} \le \widetilde{\pi}_i(x_i) \\ 0, & \widetilde{\pi}_i(x_i) < \pi_\mathrm{c} \end{cases} \tag{54}$$

Comparing with eq.(50), we see that a positive tax reduces the robustness in the range of positive but finite robustness:

$$\widehat{\alpha}^L(x_i, \pi_c) \le \widehat{\alpha}^\ell(x_i, \pi_c) \tag{55}$$

(strict inequality at positive tax, $\tau > 0$; equality in the absence of tax, $\tau = 0$). A subsidy, $\tau < 0$, would reverse the direction of the inequality. One way to understand relation (55) is in terms of an *effective aspiration:*

$$\pi_c^L = \frac{\pi_c}{1 - \frac{\tau}{\pi_\ell}} \tag{56}$$

A positive tax, $0 < \tau \le \pi_\ell$, makes the effective aspiration π_c^L greater than π_c. Since the robustness decreases as the aspiration increases, we obtain the robustness-reduction of relation (55). Conversely, a subsidy, $\tau < 0$, reduces the effective aspiration and enhances the robustness.

The output which maximizes the firm's robustness at aspiration π_c, in analogy to eq.(51), is:

$$\widehat{x}_i^\ell = \sqrt{\frac{\pi_c}{1 - \frac{\tau}{\pi_\ell}}} \tag{57}$$

A positive tax, $0 < \tau < \pi_\ell$, forces the firm to increase its output; a subsidy, $\tau < 0$, reduces output. The resulting info-gap-optimal robustness, the analog of eq.(52), is:

$$\widehat{\alpha}^L(\widehat{x}_i^L, \pi_c) = \begin{cases} \infty, & \pi_c < \pi_\ell \\ \dfrac{1 - c - \widetilde{x}_{-i}^T \mathbf{1} - 2\sqrt{\frac{\pi_c}{1-\frac{\tau}{\pi_\ell}}}}{\sqrt{\mathbf{1}^T W^{-1} \mathbf{1}}}, & \pi_\ell \le \pi_c \le \left(\dfrac{1 - c - \widetilde{x}_{-i}^T \mathbf{1}}{2}\right)^2 \left(1 - \frac{\tau}{\pi_\ell}\right) \\ 0, & \left(\dfrac{1 - c - \widetilde{x}_{-i}^T \mathbf{1}}{2}\right)^2 \left(1 - \frac{\tau}{\pi_\ell}\right) < \pi_c \end{cases} \tag{58}$$

The imposition of the liability-tax τ reduces the optimal robustness:

$$\widehat{\alpha}^L(\widehat{x}_i^L, \pi_c) \le \widehat{\alpha}^\ell(\widehat{x}_i^\ell, \pi_c) \tag{59}$$

with strict inequality at positive tax ($\tau > 0$). As before, a subsidy reverses the direction of the inequality.

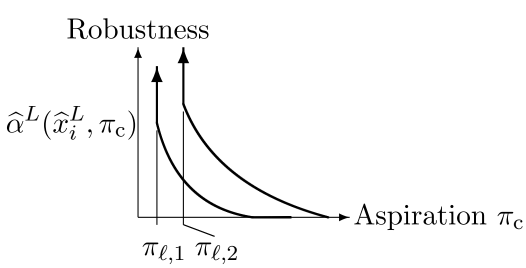

Figure 11: Two optimal robustness curves, eq.(58), for different liability limits.

The effect on the optimal robustness curve $\widehat{\alpha}^L(\widehat{x}_i^L, \pi_c)$, of a change of liability-limit π_ℓ, is shown in fig. 11. Comparing this with fig. 10 we see the effect of balancing the liability-limitation π_ℓ with a positive tax τ. While the 'balanced' robustness is less than the tax-free case, eq.(59), we see in fig. 11 that an increase in the liability-limit π_ℓ (a decrease in the firm's liability) causes enhanced robustness (at fixed aspiration π_c) throughout the finite non-zero range. The curvature of the robustness curve may also change.

6.4 Example

We consider a different example of the influence of limited liability on the info-gap robustness function. Using the Heaviside function, eq.(47), we define the reward for action vector q with uncertain vector u and limited liability parameter π_ℓ:

$$R(q,u) = u^T q\, h\left(u^T q - \pi_\ell\right) + \pi_\ell \left[1 - h\left(u^T q - \pi_\ell\right)\right] \tag{60}$$

The reward is $u^T q$ provided that this quantity is no less than π_ℓ; otherwise the reward is π_ℓ. That is, the decision maker's liability is for providing acceptably large reward, unless the earned reward, $u^T q$, is below the threshold π_ℓ, in which case somebody (like the government) steps in and pays the difference to bring the earnings up to the value π_ℓ.

The robustness of action q, with reward-aspiration π_c, is:

$$\widehat{\alpha}(q,\pi_c) = \max\left\{\alpha : \min_{u\in\mathcal{U}(\alpha,\widetilde{u})} R(q,u) \ge \pi_c\right\} \tag{61}$$

Let the uncertainty in u be represented by an ellipsoid-bound info-gap model:

$$\mathcal{U}(\alpha,\widetilde{u}) = \left\{u = \widetilde{u} + v :\ v^T W v \le \alpha^2\right\}, \quad \alpha \ge 0 \tag{62}$$

where $\widetilde{u}$ is the known nominal value of u and W is a known, real, symmetric, positive definite matrix.

For this info-gap model the robustness function becomes:

$$\widehat{\alpha}(q,\pi_c) = \begin{cases} \infty, & \pi_c < \pi_\ell \\ \dfrac{\widetilde{u}^T q - \pi_c}{\sqrt{q^T W^{-1} q}}, & \pi_\ell \le \pi_c \le \widetilde{u}^T q \\ 0, & \widetilde{u}^T q < \pi_c \end{cases} \tag{63}$$

As shown in fig. 12, the robustness curve has two slope-discontinuities. The robustness is infinite when the profit-aspiration π_c is less than the liability level π_ℓ. At the other extreme, the robustness is zero for any aspiration π_c which exceeds the nominal reward $\widetilde{u}^T q$. For intermediate aspirations the robustness is a linearly decreasing function whose slope is independent of the liability level.

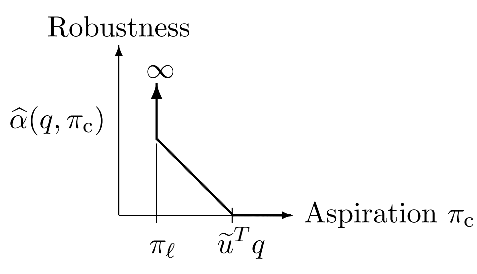

Figure 12: Robustness curve with fixed action q, eq.(63), showing discontinuous slope.

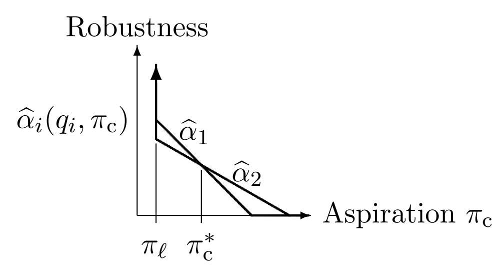

Figure 13: Two robustness curves with fixed actions q_i, eq.(63), for different and mutually exclusive options.

Now let us consider the choice between two mutually exclusive alternatives, each with the limited-liability reward function of eq.(60), but with a different decision vector q_i. The uncertainties also differ, each represented with its own ellipsoid-bound info-gap model, eq.(62), with centerpoint $\widetilde{u}_i$ and shape matrix W_i. The robustness function for option i is $\widehat{\alpha}_i(q_i,\pi_c)$, given by eq.(63) (with q_i, $\widetilde{u}_i$ and W_i).

These robustness functions are shown in fig. 13. They cross when the profit-aspiration equals $\pi_{\rm c}^*$, and they both are infinite for aspiration below the liability limit π_ℓ. The intersection at $\pi_{\rm c}^*$ represents a preference-reversal: alternative '2' is more robust and hence preferred for aspirations above $\pi_{\rm c}^*$, while alternative '1' is preferred below $\pi_{\rm c}^*$. However, if the liability limit π_ℓ is close to the intersection point $\pi_{\rm c}^*$, then this choice-reversal can be eliminated: the aspiration is reduced to π_ℓ and the decision maker is indifferent between the alternatives, both of which have infinite robustness.

A Info-gap Models of Uncertainty

Our quantification of knowledge-deficiency is based on non-probabilistic information-gap models [3]. An info-gap is a disparity between what the decision maker knows and what could be known. The range of possibilities expands as the info-gap grows. An info-gap model is a family of nested sets. Each set corresponds to a particular degree of knowledge-deficiency, according to its level of nesting. Each element in a set represents a possible event. There are no measure functions in an info-gap model.

Info-gap theory provides a quantitative model for Knight's concept of "true uncertainty" for which "there is no objective measure of the probability", as opposed to risk which is probabilistically measurable [9, pp.46, 120, 231–232]. Further discussion of the relation between Knight's conception and info-gap theory is found in [3, section 12.5]. Similarly, Shackle's "non-distributional uncertainty variable" bears some similarity to info-gap analysis [16, p.23]. Likewise, Kyburg recognized the possibility of a "decision theory that is based on some non-probabilistic measure of uncertainty." [11, p.1094].

Events are represented as vectors or vector functions u. Knowledge-deficiency is expressed at two levels by info-gap models. For fixed α the set $\mathcal{U}(\alpha, \widetilde{u})$ represents a degree of variability of u around the centerpoint $\widetilde{u}$. 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 ($\widetilde{u}$ 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 imperfection of knowledge: the horizon of variation is unbounded.

Let $\Re$ denote the non-negative real numbers and let Ω be a Banach space in which the uncertain quantities u are defined. An info-gap model $\mathcal{U}(\alpha, \widetilde{u})$ is a map from $\Re \times \Omega$ into the power set of Ω. Info-gap models obey a number of axioms, which may take different forms, depending on the application. The most important axioms are 'nesting' and 'contraction'. *Nesting:* $\mathcal{U}(\alpha, \widetilde{u}) \subseteq \mathcal{U}(\alpha', \widetilde{u})$ if $\alpha \le \alpha'$. *Contraction:* $\mathcal{U}(0, \widetilde{u})$ is the singleton set $\{\widetilde{u}\}$. The range of variability increases as the horizon of uncertainty, α, grows, and the centerpoint or nominal function, $\widetilde{u}$, belongs to the sets at all horizons of uncertainty. For more discussion of these axioms see [2].

B Info-gap Immunity Functions: Robustness and Opportunity

In this section we briefly describe the basic decision functions of info-gap theory, based on [3, section 3.1].

B.1 A First Look

The *robustness function* expresses the greatest level of uncertainty at which failure cannot occur; the *opportunity function* is the least level of uncertainty which entails the possibility of sweeping success. The robustness and opportunity functions address, respectively, the pernicious and propitious facets of uncertainty.

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:

$$\widehat{\alpha}(q) = \max\{\alpha : \text{ minimal requirements are always satisfied}\} \qquad \text{(robustness)} \tag{64}$$

$$\widehat{\beta}(q) = \min\{\alpha : \text{ sweeping success is sometimes enabled}\} \qquad \text{(opportunity)} \tag{65}$$

We can "read" eq. (64) as follows. The robustness $\widehat{\alpha}(q)$ of decision vector q is the greatest value of the uncertainty parameter α for which specified minimal requirements are always satisfied. $\widehat{\alpha}(q)$ expresses robustness — the degree of resistance to uncertainty and immunity against failure — so a large value of $\widehat{\alpha}(q)$ is desirable. Eq. (65) states that the opportunity $\widehat{\beta}(q)$ 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. $\widehat{\beta}(q)$ is the immunity against windfall reward, so a small value of $\widehat{\beta}(q)$ is desirable. A small value of $\widehat{\beta}(q)$ reflects the opportune situation that great reward is possible even in the presence of little ambient uncertainty. The immunity functions $\widehat{\alpha}(q)$ and $\widehat{\beta}(q)$ are complementary and are defined in an anti-symmetric sense. Thus "bigger is better" for $\widehat{\alpha}(q)$ while "big is bad" for $\widehat{\beta}(q)$. The immunity functions — robustness and opportunity — are the basic decision functions in info-gap decision theory.

The robustness function in eq.(64) 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. By selecting an action q according to its robustness $\widehat{\alpha}(q)$, the robustness function underlies a satisficing decision algorithm which optimizes the immunity to pernicious uncertainty.

The opportunity function in eq.(65) 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". When $\widehat{\beta}(q)$ 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.

B.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 $\mathcal{U}(\alpha, \widetilde{u})$, $\alpha \geq 0$. The uncertain u 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 auxilliary variable of no inherent significance which nonetheless influences the overall reward.

Given a scalar reward function $R(q, u)$, the minimal requirement in eq.(64) is that the reward $R(q, u)$ be no less than a critical value r_{c}. Likewise, the sweeping success in eq.(65) 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}:

$$r_{\text{w}} > r_{\text{c}} \tag{66}$$

The robustness and opportunity functions of eqs.(64) and (65) can now be expressed more explicitly:

$$\widehat{\alpha}(q, r_{\text{c}}) = \max\left\{\alpha : \min_{u \in \mathcal{U}(\alpha, \widetilde{u})} R(q, u) \geq r_{\text{c}}\right\} \tag{67}$$

$$\widehat{\beta}(q, r_{\rm w}) \quad = \quad \min \left\{ \alpha : \max_{u \in \mathcal{U}(\alpha, \widetilde{u})} R(q, u) \geq r_{\rm w} \right\} \tag{68}$$

$\widehat{\alpha}(q, r_{\rm c})$ is the greatest level of uncertainty consistent with guaranteed reward no less than the critical reward $r_{\rm c}$, while $\widehat{\beta}(q, r_{\rm w})$ is the least level of uncertainty which must be accepted in order to facilitate (but not guarantee) windfall as great as $r_{\rm w}$. The complementary or anti-symmetric structure of the immunity functions is evident from eqs.(67) and (68).

$\widehat{\alpha}(q, r_{\rm c})$ is the least upper bound of a set of α values. If this set is empty then we define the robustness to be zero: the immunity to failure vanishes. Similarly, $\widehat{\beta}(q, r_{\rm w})$ is the greatest lower bound of a set of α values. If this set is empty then we define the opportunity function to be infinite: the immunity to windfall is unbounded since windfall cannot occur at any horizon of uncertainty.

C References

1. Kenneth J. Arrow, 1984, *Individual Choice Under Certainty and Uncertainty: The Collected Papers of Kenneth J. Arrow,* Vol. 3, Belknap Press, Cambridge, Massachusetts.
2. Yakov Ben-Haim, 1999, Set-models of information-gap uncertainty: Axioms and an inference scheme, *Journal of the Franklin Institute,* 336: 1093–1117.
3. Yakov Ben-Haim, 2001, *Information-gap Decision Theory: Decisions Under Severe Uncertainty,* Academic Press.
4. Dixon, Huw David, 2001, *Surfing Economics: Essays for the Inquiring Economist,* Palgrave.
5. D. Galai and Z. Wiener, 2002, Government support of investment projects in the private sector: A micro-economic approach, Working Paper, Samuel Neaman Institute for Advanced Research in Science and Technology, Technion, Israel.
6. J. Galbraith, 1973, *Designing Complex Organizations,* Addison-Wesley Publ. Co.
7. Friedrich A. Hayek, *Individualism and Economic Order.* University of Chicago Press, 1948.
8. John Maynard Keynes, 1936, *The General Theory of Employment, Interest, and Money,* Harcourt Brace & World.
9. Frank H. Knight, 1921, *Risk, Uncertainty and Profit.* Houghton Mifflin Co. Re-issued by University of Chicago Press, 1971.
10. Frank H. Knight, 1951, *The Economic Organization,* Harper Torchbook edition, 1965.
11. Henry E. Kyburg, jr., 1990, Getting fancy with probability, *Synthese,* 90: 189–203.
12. R.P. Mack, 1971, *Planning and Uncertainty: Decision Making in Business and Government Administration,* Wiley.
13. James G. March, 1988, *Decisions and Organizations,* Basil Blackwell Ltd.
14. A. Mas-Colell, M.D. Whinston and J.R. Green, 1995, *Microeconomic Theory,* Oxford University Press.
15. Paul A. Samuelson, 1983, *Foundations of Economic Analysis.* Enlarged edition. Harvard University Press.
16. G.L.S. Shackle, 1972, *Epistemics and Economics: A Critique of Economic Doctrines,* Cambridge University Press.
17. Herbert A. Simon, 1997, *An Empirically Based Microeconomics,* Raffaele Mattioli Lectures, Cambridge University Press.