the lower value, $0.9\lambda^\star$, at all likelihood aspirations L_c. In other words, $\lambda^\star$ is a better robust-satisficing choice than $0.9\lambda^\star$ at any L_c. However, the robustness curves for $\lambda^\star$ and $1.1\lambda^\star$ cross at $(L_\times, \widehat{\alpha}_\times)$, indicating that the MLE is preferable at large L_c and low robustness, while $1.1\lambda^\star$ is preferred elsewhere. The crossing of robustness curves entails the reversal of preference between $\lambda^\star$ and $1.1\lambda^\star$, as discussed in section 3.1.8.

This pattern is repeated for all of 500 random samples: $\widehat{\alpha}(\lambda^\star, L_c)$ dominates $\widehat{\alpha}(0.9\lambda^\star, L_c)$, while $\widehat{\alpha}(\lambda^\star, L_c)$ and $\widehat{\alpha}(1.1\lambda^\star, L_c)$ cross. The coordinates of the intersection of $\widehat{\alpha}(\lambda^\star, L_c)$ and $\widehat{\alpha}(1.1\lambda^\star, L_c)$ are plotted in fig. 3.20 for 500 robustness curves, each generated from a different 20-element random sample. In each case, $\lambda^\star$ is the MLE of that sample. The vertical axis is the robustness at the intersection, $\widehat{\alpha}_\times$, divided by the maximum robustness for that sample, $\alpha_{\max}$, defined in eq.(3.147). The horizontal axis is the likelihood-aspiration at the intersection, $L_\times$, divided by the maximum likelihood for the sample, $L[X, \widetilde{p}(x|\lambda^\star)]$.

The center of the cloud of points in fig. 3.20 is about (0.5, 0.2). What we learn from this is that the robustness curves for $\lambda^\star$ and $1.1\lambda^\star$ typically cross at a likelihood aspiration of about half the best-estimated value, and at a robustness of about 20% of the maximum robustness. We also see that curve-crossing can occur at much higher values of L_c, and that this tends to be at very low robustness. This happens when $L[X, \widetilde{p}(x|1.1\lambda^\star)]$ is only slightly less than $L[X, \widetilde{p}(x|\lambda^\star)]$. Curve-crossing can also occur at much lower L_c and higher robustness, typically because $L[X, \widetilde{p}(x|1.1\lambda^\star)]$ is substantially less than $L[X, \widetilde{p}(x|\lambda^\star)]$.

Note that the data in this example are generated from an exponential distribution, so there is nothing in the data to suggest that the exponential distribution is wrong. But things change, so a model which fits historical data may "become mis-specified ... [which is] the main problem in economic forecasting" [9, p.246]. The motivation for the info-gap model of eq.(3.144) is that, while the *past* has been exponential, the *future* may not be. The robust-satisficing estimate of λ accounts not only for the historical evidence (the sample X) but also for the future uncertainty about the relevant family of distributions.

In short, the curve-crossing shown in fig. 3.19 is typical, and robust-satisficing provides a technique for estimating the parameters of a pdf when the form of the pdf is uncertain.

3.3 Production Volume With Uncertain Costs

Many decisions under uncertainty involve one or another form of cost-benefit analysis. The prototype of such deliberations is the direct calculation of profit: the difference between monetary earnings and costs. We will illustrate some aspects of the robustness and opportuneness functions with a simple economic example of a small firm which needs to choose its production volume despite severely uncertain production costs.

Consider a static production situation in which the manufacturer must choose the quantity q which will be produced, and let us assume this to equal the quantity which will be sold. $p(q)$ is the known price per item, which may depend on the quantity produced, and $c(q)$ is the total cost of producing q items. We will suppose that the production cost $c(q)$ is an uncertain function of the production volume q. The profit $R(q, c)$ depends on both the manufacturer's decision q and the uncertain costs incurred, c:

$$R(q, c) = qp(q) - c(q) \tag{3.150}$$

Classical theory: profit maximization. If $p(q)$ and $c(q)$ were both completely known, then the manufacturer could maximize the profit by choosing q to satisfy:

$$0 = \frac{\mathrm{d}R}{\mathrm{d}q} = p + q\frac{\mathrm{d}p}{\mathrm{d}q} - \frac{\mathrm{d}c}{\mathrm{d}q} \tag{3.151}$$

However in our case, while the price function $p(q)$ is known, the cost function $c(q)$ is highly uncertain, so eq.(3.151) cannot be directly solved to find the optimum production volume.

Satisfice profit, maximize robustness. We now consider a robust-satisficing strategy which both guarantees no less than a specified minimum profit (if possible) and which maximizes the robustness or immunity to uncertainty.

Suppose we have chosen an info-gap model $\mathcal{U}(\alpha, \widetilde{c})$, $\alpha \geq 0$, to match the limited information which we possess about the uncertainty of the production-cost function $c(q)$. Let r_{c} be the lowest profit which the producer tentatively considers to be acceptable. A robust-satisficing strategy for choosing the volume of production is to seek the value of q which guarantees a profit of at least r_{c} and which maximizes the producer's immunity to the unknown variation of the cost function. This "robust rationality" satisfies the producer's need for a minimum profit while also maximizing the avoidance of uncertainty. r_{c} is a parameter which need not be chosen unalterably, but which folds into the firm's decision process. After evaluating the robustness function, $\widehat{\alpha}(q, r_{\mathrm{c}})$, the producer can explore the desirability of different values of critical reward r_{c} by examining the robustness to uncertainty as a function of reward. As in fig. 3.1, $\widehat{\alpha}(q, r_{\mathrm{c}})$ will decrease monotonically as r_{c} increases.

The implementation of this strategy proceeds in two stages. First we find the robustness as a function of q. The robustness is the greatest horizon of uncertainty, α, which is consistent with obtaining a profit no less than r_{c}, which is expressed by eq.(3.4) on p.40.

The second stage is to choose the volume of production to maximize the robustness. This optimal volume, $\widehat{q}_{\mathrm{c}}(r_{\mathrm{c}})$, is defined in eq.(3.20) on p.45.

This approach is quite different from a max-min strategy wherein one maximizes a minimum profit. What one maximizes here is the robustness to uncertainty, making this strategy attractive for avoidance of uncertainty. Of course, one can still gamble since setting a high value for the least acceptable profit r_c will reduce the available robustness $\widehat{\alpha}$. In subsequent examples (and extensively in chapter 4) we will discuss "calibration" of the robustness: while $\widehat{\alpha}$ may be either dimensionless or have problem-specific units, its value can be scaled to give a feeling of whether the robustness is great or small.

Example 1 Uniform-bound info-gap model. Let us suppose that $p(q)$ is known exactly and that we know a nominal or typical value of the production-cost function, $\widetilde{c}(q)$, but we also know that the actual cost function can deviate from $\widetilde{c}(q)$ in some unknown way. This very limited amount of information about the variability of the cost function can be represented by a uniform-bound info-gap model:

$$\mathcal{U}_{\mathrm{u}}(\alpha, \widetilde{c}) = \{c(q) : \ |c(q) - \widetilde{c}(q)| \le \alpha\}, \quad \alpha \ge 0 \tag{3.152}$$

$\mathcal{U}_{\mathrm{u}}(\alpha, \widetilde{c})$ is the set of cost functions $c(q)$ whose deviation from $\widetilde{c}(q)$ is no greater than α. What we *know* about the cost function $c(q)$ is the nominal function $\widetilde{c}(q)$ and that the deviation of $c(q)$ from $\widetilde{c}(q)$ is bounded by α. What we *do not know* is the value of the horizon of uncertainty α, so the info-gap model is an unbounded family of nested convex sets $\mathcal{U}_{\mathrm{u}}(\alpha, \widetilde{c})$, for $\alpha \ge 0$.

The minimum profit, up to uncertainty α, in the evaluation of the robustness in eq.(3.4) is readily seen to occur for the greatest cost function allowed by the info-gap model at horizon of uncertainty α, which is simply $\widetilde{c}(q) + \alpha$:

$$\min_{c(q) \in \mathcal{U}_{\mathrm{u}}(\alpha, \widetilde{c})} R(q, c) = qp(q) - [\widetilde{c}(q) + \alpha] \tag{3.153}$$

For fixed production volume q, the robustness is the greatest value of the uncertainty parameter α for which this minimum is at least r_c. Equating the minimum profit to the critical value r_c and solving for α results in the robustness function:

$$qp(q) - [\widetilde{c}(q) + \alpha] = r_c \quad \Longrightarrow \quad \widehat{\alpha}_{\mathrm{u}}(q, r_c) = qp(q) - \widetilde{c}(q) - r_c \tag{3.154}$$

unless this expression for $\widehat{\alpha}_{\mathrm{u}}(q, r_c)$ is negative, in which case $\widehat{\alpha}_{\mathrm{u}}(q, r_c) = 0$. The robust-satisficing production volume, defined in eq.(3.20), is found from:

$$0 = \frac{\mathrm{d}\widehat{\alpha}_{\mathrm{u}}(q, r_c)}{\mathrm{d}q} = p(q) + q\frac{\mathrm{d}p(q)}{\mathrm{d}q} - \frac{\mathrm{d}\widetilde{c}(q)}{\mathrm{d}q} \tag{3.155}$$

which is precisely the classical full-information solution, eq.(3.151), if the nominal cost function, $\widetilde{c}(q)$, is employed in the classical theory. In other words, the classical profit maximization with precise information leads to the same production decision as the robust-satisficing strategy with this particular representation of uncertainty in the cost function. Furthermore, the robust choice of q is in fact independent of the producer's lowest acceptable profit, r_c. However, the info-gap robustness will equal zero for the classical best estimate of the profit. Hence the firm will evaluate its decision differently in the classical and in the robust-satisficing analyses.

Example 2 Envelope-bound info-gap model. The equivalence of the classical and the robust-satisficing decision strategies in example 1 results from the structure of the uniform-bound info-gap model of cost-function uncertainty. Suppose, as before, that the demand function $p(q)$ is known exactly, and that the nominal production-cost function $\widetilde{c}(q)$ is known, and that the shape of the envelope of uncertain variation of the actual cost function changes in a known way with the quantity produced. For instance, we may have information about the relative accuracy of the estimated cost, $\widetilde{c}(q)$, at different values of q. The envelope-bound info-gap model is suitable to this somewhat different prior information. Instead of eq.(3.152) we use:

$$\mathcal{U}_e(\alpha, \widetilde{c}) = \{c(q) : \ |c(q) - \widetilde{c}(q)| \le \alpha\psi(q)\}, \quad \alpha \ge 0 \tag{3.156}$$

where $\widetilde{c}(q)$ and $\psi(q)$ are known, and again we have an unbounded family of nested convex sets of possible production-cost functions. $\psi(q)$ determines the *shape* of the envelope within which $c(q)$ varies in an unknown way, while the horizon of uncertainty α (whose value is unknown) determines the *size* of the envelope. Now, instead of eq.(3.153), the least possible profit up to uncertainty α becomes:

$$\min_{c(q)\in\mathcal{U}_e(\alpha,\widetilde{c})} R(q,c) = qp(q) - [\widetilde{c}(q) + \alpha\psi(q)] \tag{3.157}$$

And, instead of eq.(3.154), the robustness of production volume q becomes:

$$\widehat{\alpha}_e(q, r_c) = \begin{cases} \dfrac{R[q,\widetilde{c}(q)] - r_c}{\psi(q)} & \text{if } R[q,\widetilde{c}(q)] \ge r_c \\ 0 & \text{else} \end{cases} \tag{3.158}$$

where $R[q, \widetilde{c}(q)]$ is the estimated profit function, eq.(3.150).

Seeking the most robust production volume, as in eq.(3.155), we will (usually) obtain the classical profit-maximization result *only if* the uncertainty-envelope function $\psi(q)$ is constant (so that $\mathcal{U}_e$ reverts to the uniform-bound info-gap model $\mathcal{U}_u$).

The robustness of the decision can be calibrated by noting from eq.(3.156) that $c(q)$, $\widetilde{c}(q)$ and α all have monetary units ($\psi(q)$ is a dimensionless envelope function). $\widehat{\alpha}_e(q, r_c)$ is the greatest uncertainty in the cost function $\widetilde{c}(q)$ which is consistent with minimum profit r_c. The ratio $\widehat{\alpha}_e(q, r_c)/\widetilde{c}(q)$ is a dimensionless expression of the greatest acceptable fractional uncertainty in $\widetilde{c}(q)$. If $\widehat{\alpha}_e(q, r_c)/\widetilde{c}(q) \gg 1$ then large fractional variation of $c(q)$ with respect to the nominal function does not jeopardize attainment of the critical profit r_c, and the production volume q is very robust. On the other hand, if $\widehat{\alpha}_e(q, r_c)/\widetilde{c}(q) \ll 1$, then even small fractional deviation of $c(q)$ from its nominal value entails shortfall below the critical profit, so this value of q is extremely fragile to uncertainty.

To illustrate the robustness function let the price per item be constant, $p(q) = p_0$, which means that the market price is unaffected by the manufacturer's production volume, as would be expected when a small firm faces a large competitive market. Also, let the envelope function be linear: $\psi(q) = q$, meaning that the uncertainty in the production cost is proportional to the production volume. Finally, let the nominal cost function be: $\widetilde{c}(q) = \widetilde{c}_0\sqrt{q}$. The robustness function in eq.(3.158) becomes:

$$\widehat{\alpha}_e(q, r_c) = \begin{cases} p_0 - \dfrac{\widetilde{c}_0}{\sqrt{q}} - \dfrac{r_c}{q} & \text{if } qp_0 - \widetilde{c}_0\sqrt{q} \geq r_c \\ 0 & \text{else} \end{cases} \tag{3.159}$$

The robustness, $\widehat{\alpha}_e(q, r_c)$, increases monotonically with increasing production volume q, so the manufacturer will be inclined to maximize production. However, the marginal increase in robustness decreases steadily with increasing q.

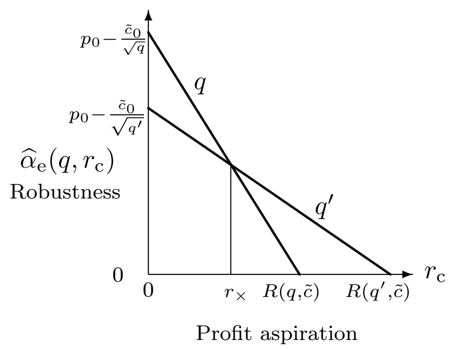

Figure 3.21: Robustness functions vs. profit aspiration for two different production volumes, eq.(3.159).

Robustness curves for two different production volumes, q and $q' > q$, are shown in fig. 3.21, where $\sqrt{q'} + \sqrt{q} > \widetilde{c}_0/p_0$ implying that $R(q, \widetilde{c}) < R(q', \widetilde{c})$. The robustness decreases as the critical profit, r_c, increases, illustrating the trade-off asserted in eq.(3.23). Furthermore, the robustness vanishes precisely when the aspired profit r_c equals the anticipated profit as stated in eq.(3.25). Finally, these robustness curves cross, indicating the potential for preference reversal. The nominal profit at production volume q' exceeds the nominal profit at volume q, suggesting that q' might be preferred over q. However, the slope of the robustness curve is more negative for q than for q', so the curves cross at $r_\times$. This means that q will be preferred over q' at profit aspiration less than $r_\times$.

Information and robustness. One would generally expect that more information would enable the choice of a more robust strategy. We can give this intuition a precise statement, and also understand when to expect surprises. The extent of our information about the uncertain elements of the analysis is represented by an info-gap model. The relation "more information" in one formulation than another is expressed by set-inclusion of the corresponding info-gap models: $\mathcal{U}_1(\alpha, \widetilde{c}) \subset \mathcal{U}_2(\alpha, \widetilde{c})$. Model '1' is more informative than model '2' because the former constrains the uncertain quantity more tightly than the latter. In $\mathcal{U}_1(\alpha, \widetilde{c})$ the information gap is smaller than in $\mathcal{U}_2(\alpha, \widetilde{c})$. It then results from eq.(3.4) that the corresponding robustnesses are ranked in reverse order: $\widehat{\alpha}_1(q, r_c) \geq \widehat{\alpha}_2(q, r_c)$. (We will explore this ranking further in chapter 7.)

For instance, suppose that the envelope function $\psi(q)$ in eq.(3.156) satisfies: $0 \leq \psi(q) < 1$. Then the envelope-bound info-gap model is more informative than the uniform-bound model, $\mathcal{U}_e(\alpha, \widetilde{c}) \subset \mathcal{U}_u(\alpha, \widetilde{c})$, and the corresponding robustnesses, eqs.(3.154) and (3.158), are ranked as $\widehat{\alpha}_e(q, r_c) > \widehat{\alpha}_u(q, r_c)$. No robust strategy based on the uniform-bound model will yield a robustness greater than that obtainable from the envelope-bound model (with the same values of q and r_c).

Surprises can occur in the relative robustnesses of different formulations as a result of the incautious assessment of how informative an info-gap model really is. A Fourier ellipsoid-bound info-gap model, such as eq.(2.26) on p.26, is certainly more complicated and intricate than the uniform-bound model, eq.(3.152). However, it is not so clear that the Fourier-bound model is more informative than the uniform-bound model in the strict sense of more tightly constraining the uncertain cost function for all values of $\alpha \geq 0$. In fact, these info-gap models can be chosen so that neither includes the other, though they do intersect. While one model may lead to a more robust decision than the other at some aspiration levels, we are unclear about which is which until we perform the analysis.

Information surrogates. Given two info-gap models, where $\mathcal{U}_1(\alpha, \widetilde{c}) \subset \mathcal{U}_2(\alpha, \widetilde{c})$, we can now simply say that $\mathcal{U}_1(\alpha, \widetilde{c})$ is more informative than $\mathcal{U}_2(\alpha, \widetilde{c})$. However, it is nonetheless still possible that the maximum robustnesses of these two models are the same: $\widehat{\alpha}_2(\widehat{q}_2, r_c) = \widehat{\alpha}_1(\widehat{q}_1, r_c)$. The decision maker will have the same vulnerability to info-gaps with $\mathcal{U}_2(\alpha, \widetilde{c})$ as with $\mathcal{U}_1(\alpha, \widetilde{c})$, so we can use the terminology of Simon [160, p.239] and refer to $\mathcal{U}_2(\alpha, \widetilde{c})$ as an *information surrogate* for $\mathcal{U}_1(\alpha, \widetilde{c})$.

In fact, given any two info-gap models, regardless of whether one is included in the other, if they have the same maximum robustness we can say that either model is an information surrogate for the other, or that the two models have the relation of *information equivalence.* The decision maker's robustness will be identical, whichever information set is used, though different robust-satisficing strategies may be adopted. That is, $\widehat{q}_1$ may differ from $\widehat{q}_2$. (We will explore a related issue in chapter 9 when we ask: by how much can the info-gap models of two decision makers differ and yet they still agree on a course of action.)

We might also say that the uniform-bound info-gap model is an information surrogate for precise knowledge of the cost function in example 1 of this section, since the robust-satisficing strategy based on $\mathcal{U}_{\rm u}$ leads to the same choice of q as the classical profit-maximization strategy (compare eqs.(3.151) and (3.155)). While the decision maker makes the same choice with a robust-satisficing strategy based on $\mathcal{U}_{\rm u}$ as with complete knowledge, this is serendipitous, since the optimization goals are different in these two situations. In one case profit is optimized while in the other case robustness is optimized and profit is satisficed. Such serendipitous information surrogates are not uninteresting or unimportant, but we need not be disappointed when they are lacking, since the decision maker may in fact deliberately prefer the robust-satisficing strategy even though it fits a less informative model.

A note of caution. We have said that $\mathcal{U}_1(\alpha, \widetilde{c})$ is more informative than $\mathcal{U}_2(\alpha, \widetilde{c})$ if $\mathcal{U}_1(\alpha, \widetilde{c}) \subset \mathcal{U}_2(\alpha, \widetilde{c})$, which suggests an absolute scale of information. This inclusion does imply an absolute ordering of extrema such as the robustness functions $\widehat{\alpha}(q, r_{\rm c})$ or $\widehat{\alpha}(\widehat{q}_{\rm c}, r_{\rm c})$. However, we have noted that info-gap models which are not nested one in the other are not commensurable in this way. But even info-gap models which *are* commensurable by inclusion may have very different implications when evaluated on the basis of performance criteria other than simple extrema. In short, the measurement of information is usually not absolute, but rather relative to an end use.

Robustness and information-gap uncertainty. The 'robust rationality' which is developed in these examples is an alternative to the classical maximization of profit. This info-gap robust-satisficing analysis is based on a non-probabilistic quantification of uncertainty. The gap between what the manufacturer *does know* (a nominal cost function, in these examples) and what *needs to be known* in order to implement the classical analysis (the true cost function), is represented by an info-gap model. It is in the context of info-gap models that robust-satisficing provides a particularly natural decision algorithm. An info-gap model is an unbounded family of nested uncertainty sets, where the horizon of uncertainty determines both the level of nesting and the magnitude of the info-gap. The robust-satisficing decision procedure in eqs.(3.4) and (3.20) takes immediate advantage of the info-gap model structure to define a strategy which maximizes the decision maker's immunity to the info-gap. While this approach may appeal to some decision makers on its own merits as a criterion of optimality, it may also be a pragmatic necessity in those situations where lack of information prevents implementation of the more classical analysis. Furthermore, as we will show in section 11.4, the robust-satisficing strategy is a better bet than direct best-model optimization in a wide class of problems.

Example 3 Robustness and opportuneness. Let us continue the production volume decision problem with the envelope-bound info-gap model for uncertainty in the cost function, eq.(3.156) in example 2, and evaluate the opportuneness function.

The greatest possible profit, for any cost function up to horizon of uncertainty α, occurs when the cost function falls on the lower envelope:

$$\max_{c(q)\in\mathcal{U}_{\mathrm{e}}(\alpha,\widetilde{c})} R(q,c) = qp(q) - [\widetilde{c}(q) - \alpha\psi(q)] \tag{3.160}$$

Equating this greatest possible profit to the windfall value r_{w} and solving for the uncertainty parameter α yields the opportuneness function:

$$qp(q) - [\widetilde{c}(q) - \alpha\psi(q)] = r_{\mathrm{w}} \quad\Longrightarrow\quad \widehat{\beta}_{\mathrm{e}}(q, r_{\mathrm{w}}) = \frac{r_{\mathrm{w}} + \widetilde{c}(q) - qp(q)}{\psi(q)} \tag{3.161}$$

unless this expression for $\widehat{\beta}_{\mathrm{e}}(q, r_{\mathrm{w}})$ is negative, in which case $\widehat{\beta}_{\mathrm{e}}(q, r_{\mathrm{w}}) = 0$.

The robustness function for this example, eq.(3.158), and the opportuneness function are an anti-symmetric pair, and are related as:

$$\widehat{\beta}_{\mathrm{e}}(q, r_{\mathrm{w}}) = -\widehat{\alpha}_{\mathrm{e}}(q, r_{\mathrm{c}}) + \frac{r_{\mathrm{w}} - r_{\mathrm{c}}}{\psi(q)} \tag{3.162}$$

We can now explore how the robustness and opportuneness functions, $\widehat{\beta}_{\mathrm{e}}(q, r_{\mathrm{w}})$ and $\widehat{\alpha}_{\mathrm{e}}(q, r_{\mathrm{c}})$, vary as the manufacturer's choice of the production volume q is altered. Robustness and opportuneness *may* be sympathetic, but in this example they will sometimes be antagonistic.

Differentiation eq.(3.162) with respect to q results in:

$$\frac{\partial\widehat{\beta}}{\partial q} = -\frac{\partial\widehat{\alpha}}{\partial q} - \frac{r_{\mathrm{w}} - r_{\mathrm{c}}}{\psi^2(q)}\frac{\mathrm{d}\psi}{\mathrm{d}q} \tag{3.163}$$

The quantity $(r_{\mathrm{w}} - r_{\mathrm{c}})/\psi^2$ is for all practical purposes positive, since the wildest dream for windfall reward r_{w} is doubtlessly greater than the critical survival level of reward r_{c}. This means that the extrema of $\widehat{\beta}$ and $\widehat{\alpha}$ will not occur at the same value of q unless $\mathrm{d}\psi/\mathrm{d}q$ happens to vanish precisely at an optimal q. That is, we cannot have:

$$\frac{\partial\widehat{\beta}}{\partial q} = 0 \quad\text{and}\quad \frac{\partial\widehat{\alpha}}{\partial q} = 0 \tag{3.164}$$

at the same value of q unless

$$\frac{\mathrm{d}\psi}{\mathrm{d}q} = 0 \tag{3.165}$$

as well. Furthermore, not only can we not simultaneously optimize both $\widehat{\beta}_{\mathrm{e}}(q, r_{\mathrm{w}})$ and $\widehat{\alpha}_{\mathrm{e}}(q, r_{\mathrm{c}})$, eq.(3.163) also implies that there can be some ranges of q-values for which both functions have the same slope. Since "bigger is better" for $\widehat{\alpha}$ while "big is bad" for $\widehat{\beta}$, this means that for some ranges of production volumes one can improve one of these functions only at the expense of degrading the other. In these ranges of q one can trade robustness for opportuneness or vice versa. On the other hand, this antagonism need not be universal. If, for instance, $\mathrm{d}\psi/\mathrm{d}q$ is positive, then a range of q-values over which $\widehat{\alpha}(q, r_{\mathrm{c}})$ increases (which is desirable) will be a range of q-values over which $\widehat{\beta}(q, r_{\mathrm{w}})$ decreases (which is also desirable). We will explore the antagonism and sympathy of the immunity functions further in chapter 5.

3.4 ‡ General Robustness and Opportuneness Functions

For most practical purposes eqs.(3.4) and (3.5) (or eqs.(3.10) and (3.11)) are sufficiently general and powerful assessments of the impact of pernicious and propitious aspects of uncertainty. In some circumstances, however, one needs a more general formulation of robustness and opportuneness. In this section we briefly address this need. These concepts are applied in sections 11.1 and 11.2.

Robustness function. Let q be a vector of actions made (or contemplated) by the decision maker, and let $\mathcal{Q}$ be the set of all allowed or available actions. Let $u_1, \ldots, u_N$ be info-gap variables with info-gap models $\mathcal{U}_1(\alpha, \widetilde{u}_1), \ldots, \mathcal{U}_N(\alpha, \widetilde{u}_N)$, $\alpha \geq 0$. Let $\mathcal{R}_* \left[q, \mathcal{U}_1(\alpha, \widetilde{u}_1), \ldots, \mathcal{U}_N(\alpha, \widetilde{u}_N)\right]$ be a reward algorithm which evaluates action q with respect to the sets $\mathcal{U}_1(\alpha, \widetilde{u}_1), \ldots, \mathcal{U}_N(\alpha, \widetilde{u}_N)$, at fixed uncertainty α, and returns a real-valued assessment of the reward. The reward algorithm $\mathcal{R}_*$ is *monotonically decreasing* in the sets upon which it acts. That is, $\mathcal{R}_*$ has the following property:

$$A_i \subseteq B_i, \ i = 1, \ldots, N, \quad \text{implies} \quad \mathcal{R}_*[q, A_1, \ldots, A_N] \geq \mathcal{R}_*[q, B_1, \ldots, B_N] \tag{3.166}$$

We stress that the reward algorithm $\mathcal{R}_*$ is monotonically decreasing on the uncertainty sets, but we have assumed nothing about how $\mathcal{R}_*$ varies with the decision vector q. The most common motivation for requiring this monotonic decrease of the reward algorithm is that it thereby seeks the worst case on the uncertainty-sets in its domain. We will have more to say shortly about worst cases.

Referring to eq.(3.4) we see that the expression:

$$\min_{u \in \mathcal{U}(\alpha, \widetilde{u})} R(q, u) \tag{3.167}$$

is a special case of:

$$\mathcal{R}_* \left[q, \mathcal{U}_1(\alpha, \widetilde{u}_1), \ldots, \mathcal{U}_N(\alpha, \widetilde{u}_N)\right] \tag{3.168}$$

The minimum in (3.167) decreases as the set $\mathcal{U}(\alpha, \widetilde{u})$ grows, which is precisely the property of monotonic decrease required of the generic reward algorithm in (3.168).