units the criticality is transferred from path 4, the nominal critical path, to path 3, which is the 'uncertainty critical path'. We see here an essential and unique element of the uncertainty analysis. The path which is critical with respect to the nominal conditions is *not* necessarily the critical path when sensitivity to uncertainty is considered in the presence of some extra time.

When $t_c = 30$ path 3 becomes maximally sensitive, and the project robustness is entirely determined by this path: $\widehat{\alpha} = \widehat{\alpha}_3 = 0.31$.

The transfer of path-criticality from path 4 to path 3 is an example of 'reversal of preferences' resulting from intersection of robustness curves, as discussed in section 3.1.8.

We note that in all four cases considered in table 3.2, the range of path-robustness values is quite large. For instance, at $t_c = 28$, the ratio of the most to the least robust path is $\widehat{\alpha}_5/\widehat{\alpha}_3 = 7.2$. It is worth noting that the computations for this example take a fraction of a second on a standard personal computer. Even much larger networks are readily analyzed. Further discussion of this example is found in [33].

3.2.7 Portfolio Investment

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

$$R(q, u) = \sum_{i=1}^{N} q_i u_i = q^T u \tag{3.94}$$

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

Uncertainty model. For a given investment scenario we know the anticipated future values of the securities, $\widetilde{u}_1, \ldots, \widetilde{u}_N$, which we combine in a nominal vector $\widetilde{u}$. Furthermore, we typically have information indicating the relative degree of variability of the securities and the propensity for correlated or anti-correlated variation. Specifically, we will often know an historical covariance matrix for the values of the securities. Let W denote the inverse of the covariance matrix, which is real, symmetric and positive definite. While historical values are often a poor indication of future behavior, we can use this information to formulate an ellipsoid-bound info-gap model for the uncertain future variation of the security values.

The ellipsoid-bound info-gap model for uncertain variation of the actual security-value vector u around the nominal value vector $\widetilde{u}$ is:

$$\mathcal{U}(\alpha, \widetilde{u}) = \left\{ u = \widetilde{u} + v : \ v^T W v \le \alpha^2 \right\}, \quad \alpha \ge 0 \tag{3.95}$$

Robustness function. The decision vector q is chosen to guarantee that the change in the portfolio worth, $R(q,u)$, is no less than a minimum critical reward $r_{\rm c}$, sometimes called a minimum attractive rate of return (MARR) [73]. The robustness of the portfolio investment q for critical reward $r_{\rm c}$ is the greatest value of the uncertainty parameter α such that any vector u in $\mathcal{U}(\alpha, \widetilde{u})$ results in a net worth $R(q,u)$ which is no less than $r_{\rm c}$. This is precisely the robustness in eq.(3.4) on p.40. The least reward up to uncertainty α is readily found to be:

$$\min_{u \in \mathcal{U}(\alpha, \widetilde{u})} q^T u = q^T \widetilde{u} - \alpha \sqrt{q^T W^{-1} q} \tag{3.96}$$

Equating this minimum reward to the critical value $r_{\rm c}$ and solving for the uncertainty parameter α results in the robustness:

$$\widehat{\alpha}(q, r_{\rm c}) = \frac{q^T \widetilde{u} - r_{\rm c}}{\sqrt{q^T W^{-1} q}} \tag{3.97}$$

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

Robust-satisficing investment. The robust-satisficing strategy for choosing the investment portfolio is to select q to maximize $\widehat{\alpha}(q, r_{\rm c})$. The choice of the portfolio is subject to many different possible constraints. For instance, some securities may be accessible only if some other securities are purchased as well. Or, the quantity of an security sold may be limited by the decision maker's holdings. In addition, overall budgetary constraints may limit the total purchasing power. For simplicity, we will consider only the last constraint. Let $\mathcal{Q}$, the set of feasible investment vectors, be the set of all q-vectors which exactly meet the budget, Q:

$$\sum_{i=1}^{N} q_i = Q \tag{3.98}$$

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

$$q^T \mathbf{1} = Q \tag{3.99}$$

The robust-satisficing investment for critical reward $r_{\rm c}$ is the vector q which maximizes $\widehat{\alpha}(q, r_{\rm c})$, as in eq.(3.20) on p.45.

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

$$\widetilde{u} = u_o \mathbf{1} \tag{3.100}$$

where u_o is a known constant. With this simplification, the robustness in eq.(3.97) becomes:

$$\widehat{\alpha}(q, r_c) = \frac{u_o Q - r_c}{\sqrt{q^T W^{-1} q}} \qquad (3.101)$$

where we have employed the budget constraint of eq.(3.99).

Examining eq.(3.101) we see that the robust-satisficing investment $\widehat{q}_c$, which maximizes $\widehat{\alpha}(q, r_c)$, is the vector which minimizes $q^T W^{-1} q$ subject to the constraint in eq.(3.99). Using Lagrange optimization one readily finds the robust-satisficing investment vector to be:

$$\widehat{q}_c = \frac{Q}{\mathbf{1}^T W \mathbf{1}} W \mathbf{1} \qquad (3.102)$$

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

$$\widehat{q}_{c,i} = \frac{w_{ii}}{\sum_{j=1}^{N} w_{jj}} Q \qquad (3.103)$$

The investment in an security is inversely proportional to its relative propensity for variation. In both cases, eq.(3.102) and (3.103), we see that the robust-satisficing investment in securities with equal nominal values is controlled entirely by the info-gap uncertainty, and is independent of the demanded critical return r_c.

Substituting the robust-satisficing investment $\widehat{q}_c$ of eq.(3.102) into the robustness function of eq.(3.101) we obtain the maximal robustness:

$$\widehat{\alpha}(\widehat{q}_c, r_c) = \frac{(u_o Q - r_c)\sqrt{\mathbf{1}^T W \mathbf{1}}}{Q} \qquad (3.104)$$

(recalling that $\widehat{\alpha}(\widehat{q}_c, r_c) = 0$ if the right-hand side is negative.) Eq.(3.104) shows the trade-off between immunity to uncertainty (large $\widehat{\alpha}$) and reward (large r_c): the decision maker can confidently demand great reward only in exchange for low immunity against failure due to uncertain fluctuations in the security values. The robustness equals zero when the critical reward, r_c, equals the nominal, anticipated, value for the portfolio, $u_o Q$.

Comparing portfolios. Now consider the choice between two different portfolios, each with its own set of securities, its own nominal values and its own ellipsoid-bound info-gap model of uncertainty.[8] Assuming eq.(3.100)

[8]The two portfolios may contain different numbers of securities. This would require re-normalization of the uncertainty parameters of the two info-gap models so that the robustnesses of the two portfolios are commensurable.

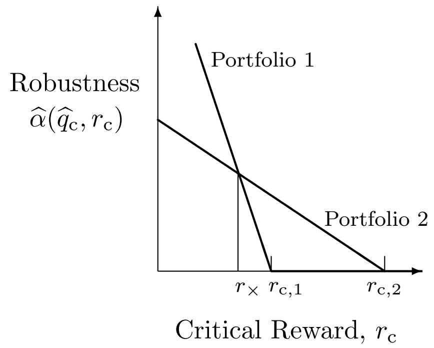

Figure 3.9: Robustness functions for two different portfolio investment alternatives.

holds, separately, for each set of securities, let $u_{o,i}$ be the anticipated value of each security in the ith portfolio, and assume that $u_{o,1} < u_{o,2}$. Likewise, let W_1 and W_2 be the shape matrices for the two uncertainty models, which are historical inverse covariance matrices. The maximum-robustness functions for the two portfolios are each described by eq.(3.104), as shown schematically in fig. 3.9 versus the critical reward r_c.

The anticipated value for the portfolios are $r_{c,i} = u_{o,i}Q$. Based on anticipated returns, portfolio 2 is more valuable than portfolio 1: $r_{c,2} > r_{c,1}$. However, each robustness function equals zero when the critical reward, r_c, equals the anticipated return.

Fig. 3.9 assists the decision maker to assess the relative riskiness of the two portfolios. Both robustness functions vanish for critical rewards in excess of $r_{c,2}$, so neither portfolio is acceptable if rewards this large are needed. For critical rewards between $r_\times$ and $r_{c,2}$, and especially above $r_{c,1}$, the second portfolio is the clear favorite over the first, since the second portfolio has greater robustness. The riskiness of the two portfolios becomes equal when the robustness curves cross, and if values of r_c less than $r_\times$ are acceptable then the first alternative becomes increasingly preferable because it affords greater immunity at the same level of guaranteed return. In this example we see how crossing of robustness curves entails reversal of preference, as discussed in section 3.1.8.

Opportuneness function. We now consider the opportuneness function $\widehat{\beta}(q, r_w)$, which is the least level of uncertainty needed to sustain the possibility of reward as large as r_w, as expressed in eq.(3.5) on p.40. The opportuneness function assesses the immunity to windfall gain r_w, so a small value of $\widehat{\beta}$ — low immunity to windfall — is desirable, unlike the robustness function for which a large value is needed to assure large immunity to failure. Windfalling, upon which the opportuneness function is based, is different from satisficing which underlies the robustness function, though on the surface the mathematics looks quite similar.

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

$$\max_{u\in\mathcal{U}(\alpha,\widetilde{u})} q^T u = q^T\widetilde{u} + \alpha\sqrt{q^T W^{-1} q} \tag{3.105}$$

whose similarity to the minimum reward in eq.(3.96) is evident. The opportuneness function is obtained by equating this maximum to the windfall reward $r_{\rm w}$ and solving for the uncertainty parameter α, leading to:

$$\widehat{\beta}(q, r_{\rm w}) = \frac{r_{\rm w} - q^T\widetilde{u}}{\sqrt{q^T W^{-1} q}} \tag{3.106}$$

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

If we impose the budget constraint of eq.(3.99) and if, as in eq.(3.100), we assume that the nominal security-values are all equal, then the opportuneness function becomes:

$$\widehat{\beta}(q, r_{\rm w}) = \frac{r_{\rm w} - u_o Q}{\sqrt{q^T W^{-1} q}} \tag{3.107}$$

which is similar to the robustness function of eq.(3.101).

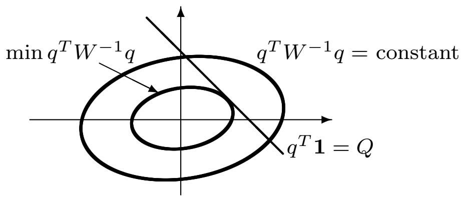

Figure 3.10: Schematic illustration of constrained optimization of $q^T W^{-1} q$.

Because windfalling is different from satisficing, and because opportuneness is different from robustness, we can now see that optimizing $\widehat{\beta}(q, r_{\rm w})$ is very different from optimizing $\widehat{\alpha}(q, r_{\rm c})$. The opportuneness function $\widehat{\beta}(q, r_{\rm w})$ is optimized (minimized) by maximizing $q^T W^{-1} q$, while the robustness is optimized (maximized) by minimizing this same quadratic term. First of all, we obviously cannot do both optimizations simultaneously. Furthermore, the first — optimizing $\widehat{\beta}(q, r_{\rm w})$ — cannot be done at all if only the budget constraint of eq.(3.99) is imposed. There simply is no maximum of $q^T W^{-1} q$ subject to $q^T\mathbf{1} = Q$. This is illustrated in fig. 3.10. No matter how large we make the quadratic term (which defines an ellipsoid) it still intersects the plane defined by the budget constraint. This is unlike the minimization of $q^T W^{-1} q$, which occurs when any further constriction of the ellipsoid would cause it to disconnect from the budget-constraint plane.

In practice of course the budget limitation of eq.(3.99) is not the only constraint. Additional constraints become active as the investment vector q ranges further from the origin: limitations in the supply of securities which can be purchased or constraints on the quantity of holdings which can be sold. Nonetheless, this example demonstrates some of the fundamental differences between windfalling with the opportuneness function and satisficing with the robustness function.

Let us leave the attempt to optimize robustness and opportuneness and note that any improvement in one function is obtained at the expense of deterioration in the other. Comparing the robustness and opportuneness functions in eqs.(3.101) and (3.107) we note that any change in the investment vector q which increases one will increase the other, and likewise any decrease in one function will be accompanied by a decrease in the other. However, "big is better" for $\widehat{\alpha}$ while "big is bad" for $\widehat{\beta}$. These immunity functions are antagonistic in this example: either immunity can be improved only at the expense of the other.

3.2.8 Monetary Policy

Economies experience adverse shocks and surprises which are unpredictable from historical data. Info-gap theory is well suited for the selection of monetary policy to counter these surprises, as we illustrate in this adaptation of Brainard's example [39] which is discussed by Blinder [36, pp.11–12]. We will show that policies which, based on best-estimated models would seem to optimize the outcome, should sometimes be avoided in favor of less aggressive policies, as was suggested by Brainard.

Consider the macroeconomic model:

$$y = Gx + z \tag{3.108}$$

where G and z are both highly uncertain, with best-estimates $\widetilde{G}$ and $\widetilde{z}$, respectively. The central bank wishes to choose x so as to pilot y towards a target value, $y^\star$.

We have no probabilistic model for the error in the estimates $\widetilde{G}$ and $\widetilde{z}$, and what we can say is that the fractional error in these estimates is unknown. That is, true (or truer) values G and z deviate from the estimated values $\widetilde{G}$ and $\widetilde{z}$ by no more than a fraction α. However, the horizon of uncertainty α is unknown. An info-gap model for this uncertainty is the following unbounded family of nested sets of G and z values:

$$\mathcal{U}(\alpha, \widetilde{G}, \widetilde{z}) = \left\{ G, z : \left| \frac{G - \widetilde{G}}{\widetilde{G}} \right| \le \alpha, \ \left| \frac{z - \widetilde{z}}{\widetilde{z}} \right| \le \alpha \right\}, \quad \alpha \ge 0 \tag{3.109}$$

At any horizon of uncertainty, α, the estimates $\widetilde{G}$ and $\widetilde{z}$ may err fractionally by as much as α. However, the value of α is not known. Thus an info-gap model does not allow a 'worst case' analysis: these is no known worst case since the horizon of error is unknown. We are deep in the domain of Knightian uncertainty.

The performance function is the squared difference between the desired value $y^\star$ and the realized value y:

$$f(x, G, z) = [y(x, G, z) - y^\star]^2 \tag{3.110}$$

In the spirit of Simon's bounded rationality and the concept of satisficing, we desire the output error, $f(x, G, z)$, to be no greater than the critical value E_c^2:

$$f(x, G, z) \le E_c^2 \tag{3.111}$$

E_c^2 can be chosen to be small or large to express demanding or modest performance aspirations.

The robustness of policy choice x is the greatest fractional error in the estimates $\widetilde{G}$ and $\widetilde{z}$, up to which every realization G and z results in acceptable squared error. Formally, the robustness of decision x with aspiration E_c is:

$$\widehat{\alpha}(x, E_c) = \max \left\{ \alpha : \left(\max_{G, z \in \mathcal{U}(\alpha, \widetilde{G}, \widetilde{z})} f(x, G, z) \right) \le E_c^2 \right\} \tag{3.112}$$

Large robustness $\widehat{\alpha}(x, E_c)$ implies that policy choice x is immune to error in the estimated model while satisficing the outcome-error at E_c. Low robustness implies that outcome-error as small as E_c cannot be confidently expected with choice x.

Let $\widetilde{y}(x) = \widetilde{G}x + \widetilde{z}$ denote the best estimate of the outcome, given choice x, and let us consider values of x for which $\widetilde{y}(x) \le y^\star$. We will assume that $\widetilde{G} > 0$ and $\widetilde{z} > 0$. The robustness of choice x is found to be:

$$\widehat{\alpha}(x, E_c) = \begin{cases} \dfrac{E_c - [y^\star - \widetilde{y}(x)]}{\widetilde{y}(x)} & \text{if } E_c \ge y^\star - \widetilde{y}(x) \\ 0 & \text{else} \end{cases} \tag{3.113}$$

Outcome-error no greater than E_c is guaranteed with policy choice x if the horizon of uncertainty is no larger than $\widehat{\alpha}(x, E_c)$.

As illustrated in fig. 3.11, the robustness gets worse ($\widehat{\alpha}$ decreases) as the aspired output error improves (E_c gets smaller). That is, robustness trades-off against performance.

Furthermore we see in eq.(3.113) and fig. 3.11 that the robustness vanishes when the aspiration E_c equals or is less than the best-estimate of the output error, $y^\star - \widetilde{y}(x)$. This is true for any choice of x. We can have little confidence in attaining fidelity as good as the best-estimated fidelity; only poorer fidelity has positive robustness. Since this is true for any x, it is also true for the choice of x which minimizes the estimated error, $f(x, \widetilde{G}, \widetilde{z})$.

This is beginning to sound like Brainard's conclusion that policies which optimize the outcome should sometimes be avoided, but there is more.

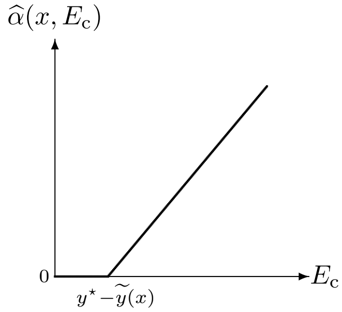

Figure 3.11: Schematic illustration of the robustness function $\widehat{\alpha}(x, E_c)$ in eq.(3.113).

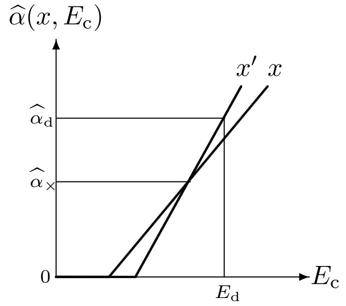

Figure 3.12: Comparison of two policy choices: reversal of preferences.

Let us now consider two policy alternatives, x and x', where:

$$\widetilde{y}(x) > \widetilde{y}(x') > 0 \tag{3.114}$$

Note that, if $\widetilde{G} > 0$, then eq.(3.114) implies that $x' < x$. That is, x' is a less aggressive intervention that x.

In particular, let x be the policy choice which, based on the best-estimated model, causes the outcome to precisely match the required value: $\widetilde{y}(x) = y^\star$. This choice of x is what would normally be called the optimal policy. Eq.(3.114) means that the estimated fidelity is worse with the less aggressive policy x' than with x:

$$0 = y^\star - \widetilde{y}(x) < y^\star - \widetilde{y}(x') \tag{3.115}$$

The robustness curves for choices x and x' are shown in fig. 3.12. $\widehat{\alpha}(x', E_c)$ intersects the E_c-axis to the right of $\widehat{\alpha}(x, E_c)$ because the best-estimated fidelity of x' is poorer than for x, eq.(3.115). However, eqs.(3.113) and (3.114) imply that the slope of $\widehat{\alpha}(x', E_c)$ is steeper than the slope of $\widehat{\alpha}(x, E_c)$, so these robustness curves cross.

Crossing of robustness curves implies reversal of preference between choices x and x', where x' is less aggressive than x. Let us suppose that the value of robustness, $\widehat{\alpha}_\times$, at which the curves in fig. 3.12 cross is fairly low. If we are quite confident that the estimates $\widetilde{G}$ and $\widetilde{z}$ are accurate, then we don't need much robustness, so $\widehat{\alpha}_\times$ might be enough robustness and we would prefer choice x over choice x'. However, we are considering severe Knightian uncertainty: great error in $\widetilde{G}$ and $\widetilde{z}$ is plausible and we need to choose a policy whose anticipated outcome is both acceptable and reliably achieved. Thus, if outcome-error E_d in fig. 3.12 is good enough fidelity, and if $\widehat{\alpha}_d$ is great enough robustness, then our robust-satisficing preference is for x' over x. If an acceptable combination (E_d, $\widehat{\alpha}_d$) is not found on the

x'-curve, then we need to search for some other choice, x'', whose robustness is adequate at acceptable fidelity. If no such x'' exists, then no acceptable policy choice is available in the current state of knowledge. We either revise our aspirations, or do some data-hunting, or revise our model.

This very simple example has illustrated the info-gap robust-satisficing version of Brainard's dictum, namely, calculate the optimum and then do less, or for what Blinder refers to as "a little stodginess at the central bank" [36, p.12]. Fig. 3.12 shows that the optimal choice, x, is less desirable than the sub-optimal and less aggressive choice x' under severe uncertainty, because the latter more reliably yields acceptable outcomes (if $E_{\rm d}$ is adequate).

While crossing of robustness curves as in this example is very common, it is not universal. It can happen that robustness curves do not cross, in which case the policy-selection stodginess disappears: the optimizing choice will coincide with the robust-satisficing choice. However, the caution remains in assessing what outcome can be considered reliable. Since the trade-off between robustness and outcomes is universal, the robust-satisficing policy maker will not anticipate (or depend upon) the best-estimated outcome because the robustness of this outcome is zero. Rather, by "migrating up" the robustness curve to an acceptable level of robustness, the analyst finds the corresponding outcome which can reliably be anticipated.

3.2.9 Search and Evasion

In a wide class of "tracking" problems an intelligent "hunter" tries to catch an intelligent "prey". Decision-theoretic interest in this problem arises when one side has only fragmentary information about the other side's strategy and intentions. We will use the robustness function to optimize the hunter's tracking strategy in a simple realization of this problem.

The hunter and prey both move along a line, and their positions at time t are denoted $x(t)$ and $u(t)$ respectively. The hunter's initial position is $x(0) = 0$ while the prey is initially at a positive position $u(0)$. The hunter is able to measure the prey's position as it evolves in time, and the hunter's strategy is to move towards the prey with velocity proportional to the distance between them:

$$\frac{{\rm d}x(t)}{{\rm d}t} = q\,[u(t) - x(t)] \tag{3.116}$$

where q, which represents the hunter's "effort", is a positive constant which the hunter must choose before starting the chase.

The hunter has only very limited information about the prey's technique of evasion. The hunter knows that the prey's typical speed is $\widetilde{s}$, and that the prey's speed deviates from $\widetilde{s}$ by no more than a fixed constant, but this constant is not known. A slope-bound info-gap model represents this prior information:

$$\mathcal{U}(\alpha, \widetilde{s}) = \left\{ u(t) : \ \left| \frac{{\rm d}u(t)}{{\rm d}t} - \widetilde{s} \right| \le \alpha \right\}, \quad \alpha \ge 0 \tag{3.117}$$