FORUM

Robustness Analysis

by

Jonathan Rosenhead

Department of Operational Research

London School of Economics & Political Science

London, England

Robustness Analysis is a way of supporting decision making when there is radical uncertainty about the future. It addresses the seeming paradox – how can we be rational in taking decisions today if the most important fact that we know about future conditions is that they are unknowable? It resolves the paradox by assessing initial decisions in terms of the attractive future options that they keep open.

While writing this article an academic colleague on another continent wrote proposing to purchase an air ticket to London for a project meeting scheduled for the summer of 2003 – ie 1 year ahead. The reason: that his local currency (in which the project funds are kept) is inflating fast, so that the ticket will cost much more if purchased later. This can stand in as a simplified example of the dilemmas that life presents us with. Should he purchase an inflexible ticket now? Undoubtedly the cheapest option, but also the one with most exposure to uncertainty. Should he buy now, but choose a ticket with some form of flexibility for subsequent change of flight? Should he delay, and buy an inflexible ticket at a later time when our project uncertainties are less, but when the price will be higher.

This decision needs to be taken now. Not taking a decision is also a decision, the decision not to purchase. But consider the uncertainties. The project may move faster or slower than expected, rendering the planned August meeting untimely. An internal dispute may undermine the project, or the sponsors may pull the plug. Either of us may be subject to illness, or family demands, or to competing time priorities for those particular weeks from other equally valid commitments. And we do not know how fast or slow the future rate of inflation will be.

To resolve this particular problem will not, I am sure, require a massive analytical apparatus. But it does illustrate, in the small, the uncertainty-related issues that can bedevil a wide range of decisions – decisions confronted by individuals, businesses, public agencies, voluntary associations, governments. Many of these decisions are of an order of complexity that does merit serious analytic attention.

The example also points up the organising principle that most sensible people would use, intuitively, when confronted by such dilemmas: namely, to explore the future options for action that are left open by the alternative choices available to them now. What is curious is that OR/MS has almost entirely neglected this concept of flexibility, remaining largely fixated on optimisation or methods derived from it.

Flexibility is not the only criterion that is relevant, but it should be among those that are employed. Employing a number of criteria without a predefined rule for combining them may seem sloppy and incomplete to those who think that the task of analysis is to decide the issue. If, however, we adopt the more modest and practicable aim of providing those who have the problem with structured information relevant to their decision, this difficulty evaporates. And indeed it makes especial sense that this information should, as far as possible, make intuitive sense to those who must use it.

In this article I will first argue for the wide prevalence of uncertainty in strategic decision making situations (and hence the potential relevance of robustness). I will then introduce the basic principles of robustness, how to specify a problem for robustness analysis, and the calculation of the robustness score. A commentary suggests how robustness can be appropriately applied, and there is an indication of the range of practical applications. Finally, robustness analysis is distinguished from a number of other nearby methods.

Prevalence of uncertainty

How widespread are the decision situations for which uncertainty is crucial? I am tempted to wonder whether that question even needs asking in a world so evidently turbulent in its arrangements. I advance in evidence (just a sample, and all at the macro level) the collapse of the Soviet Union (who predicted it?), the dot.com bubble, the

Twin Towers, the 2002 bear market. Clearly there are many decision situations in which uncertainty does not pay a key role. This is particularly true of repetitive operational decisions. For these the rate of change of underlying conditions is usually small compared with the cycle-time of activities, and inherent variability can be accommodated by probabilistic analysis. However these saving graces are not usually available in the case of non-routine, more strategic decision situations.

Businesses are subject to turbulence in the market place, to variations in regulatory regimes, to new technologies threatening established markets, to the unpredictable results of R&D or of mineral resource exploration. Public service providers are vulnerable to the vagaries of governmental funding, to changing expectations of their clientele, to the organisational impacts of new technologies. Nation states may experience violent threats of novel kinds or sources, the impact of decisions by transnational corporations, the erosion of sovereignty to supranational organisations, forceful demands for regional autonomy. Grassroots organisations are hit by the backwash of the decisions of all these more powerful actors. And so on. This is not an attempt at an exhaustive categorisation of the ways in which uncertainty permeates our decision environment. Rather it is an attempt to convince you that uncertainty is significant in particular in the more formative decisions that social organisations confront.

Principles of robustness analysis

Robustness analysis is applicable when

i) uncertainty is a factor that obstructs confident decision – which has been discussed above; and

ii) decisions must be or can be staged. - that is, the commitments made at the first point of decision do not necessarily define completely the future state of the system. There will be one or more future opportunities to modify or further define it.

The first element ensures that uncertainty matters. The second ensures that there is something that we can do about it.

A simple statement of the robustness criterion is that, other things being equal, an initial commitment should be preferred if the proportion of desirable future situations that can still be reached once that decision has been implemented is high. Put still more simply, it is a good thing to keep your options open.

That is the intuitively sensible proposition that underlies robustness analysis. Further specification is needed however to transform it into a systematic methodology that can be applied with some consistency. What counts as a desirable future situation? How do we count them? How do we identify which of them are kept open?

Specifying a problem situation for robustness analysis

The first set of elements which must be specified are

- a set of alternative initial commitments to be considered
- (normally) a set of 'futures' representative of possible environments of the system
- a set of relevant possible configurations of the system which the decisions will modify.

A commitment may be an allocation of resource in a particular decision domain, or it may comprise an integrated package of such allocations. Commitments may be those which appear logically possible, or those proposed by stakeholders with some influence over decision making. The futures, similarly, may be generated by systematic or more clearly subjective processes, or a mixture of the two. The configurations may be relevant in the sense that they are plausible extensions of the directions set by particular initial commitments; or that they can be expected to perform well in one or more of the identified futures; or that they have been proposed as a longer term goal by partisans within the management process.

It is evident that these three elements can be inter-dependent. Configurations may be generated by thinking about futures; the extrapolation of commitments may lead to possible configurations; and so on. Specification is often best achieved in interactive mode with those who are faced with the need to decide. That is, the analysis is carried out by and under the control of the relevant management group, with the assistance of one or more consultants. This and other features place robustness analysis within the family of Problem Structuring Methods (see Rosenhead and Mingers, 2001).

The three elements above need to be complemented by information of the following types:

- assessments of the compatibility of each commitment-configuration pair
- evaluation of the performance of each configuration in each future.

The former, a zero-one assessment, is needed in order to examine the extent to which options are maintained by particular commitments. The latter is also carried out on a zero-one basis. Is the predicted performance acceptable or not?

In cases where configurations consist, in effect, of an aggregation of the available commitments, compatibility can be directly established. In other cases there will be a degree of subjectivity in the assessment. Likewise, for performance evaluation, it may sometimes be possible to agree a set of multi-dimensional performance measures each with their acceptance thresholds, and to build a model to predict the values of the measures for any combination of configuration and future. In such cases the performance evaluation can be automated. Otherwise it may require discussion among those with relevant experiential knowledge to establish which performances are ‘good enough’.

If these two stages need to rely extensively on elicitation rather than on computation, there is a clear danger of combinatorial escalation rendering the process infeasible. Groups are not good at rapid and repeated but thoughtful evaluations of the kind that are required. There is therefor a strong argument for keeping the dimensions of the problem formulation as small as possible; and it may be necessary for the group to delegate the first attempt at one or both of these stages to one of its members, working with a consultant.

Analysing for robustness

Once these processes of elicitation and evaluation have been carried out, it is possible to gain a picture of the pattern of flexibility which any commitment offers, interpreting flexibility to be the future opportunity to take decisions towards desired goals. The robustness of a commitment is the ratio of the number of acceptably performing configurations with which that commitment is compatible, to the total number of acceptably performing configurations.

Clearly this limits robustness scores to the range (0, 1). A robustness score of zero indicates that no acceptable options are kept open, while a robustness of unity means that they all are.

Each commitment now has a robustness score for each future, since a configuration’s performance will vary across future contexts. Commitments can thus be assessed for the spread of flexibility they offer both within and across futures. This process will rarely identify a dominant commitment, but it will usually eliminate non-contenders, and focus discussion on just a small number of relatively attractive alternatives. It may also concentrate attention on those futures which are most crucial to the choice between these alternatives – raising the question of whether the decision-making group can exert selective influence on what future does (or does not) materialise.

Some comments

It may be noted that this procedure depends on identifying alternative futures which the system under consideration may confront. It is a fair criticism that since the future is infinitely devious, we cannot know that any of our identified futures will capture the key aspects of the future that actually happens. Evidently the elicitation process should endeavour to reduce this risk, for example by selecting a broad range of contrasting possible future environments. However the approach does not, cannot, require that this eventual future is actually identified with certainty.

Consider an initial commitment which is the first step to an ‘optimum’ solution in a single predicted future. It will maintain flexibility at best only by accident. By contrast, a robust commitment will maintain flexibility over a wide range of conceivable futures. The value of this in a future which may be outside the range of those considered cannot be rigorously demonstrated. However it is at least highly plausible that this diversity of options is more likely to include routes to one or more future configurations that will perform acceptably in the eventual future context.

In any case the principle advantage of robustness analysis lies more in its process than in its product. It does not offer a simple decision rule – “calculate the highest robustness score, and select the commitment that provides it”. Rather it provides a language in which the logic of option maintenance can be worked through. Furthermore this language is accessible also to those without developed quantitative skills. It therefor opens up for systematic dialogue with and between those who must accept responsibility for any decision, an uncertainty-based discourse that optimisation-oriented methods do not provoke.

Applications of robustness

Practical uses of robustness analysis have included

- brewery location
- chemical plant expansion
- hospital location
- regional health planning
- oil field development

- personal educational and career planning
References to most of these will be found in Rosenhead (2001a, 2001b).

Relationship to other approaches

It may be helpful to compare and contrast this notion of robustness analysis with other related approaches.

Statistical robustness

The term ‘robustness’ is used in statistics to refer to a desirable characteristic of statistical procedures. One says that a procedure is robust against some departure from the assumptions of the model when the procedure continues to work well even when, to a greater or lesser extent, the assumptions do not hold. Such assumptions, often adopted for ease of computation, might be that an underlying distribution is Normal, or that the observations have constant variance. In the case of statistical hypothesis testing, which approaches most nearly the decision-focussed approach adopted in this article, a robust test avoids the difficulty of a decision (here between two hypotheses) resting in too unstable a fashion on a particular assumption.
Bayesians give the term a rather more specific meaning. A Bayesian application is robust if the posterior distribution for an unknown parameter is not unduly affected by the choice either of the prior distribution or of the form of model taken to be generating the data.
In either approach uncertainty, though limited to knowledge about whether the specific assumptions do in fact hold, clearly lies behind the need for this concept. It does not of course purport to address other types of uncertainty, or sequentiality of decisions.

Sensitivity analysis

Sensitivity analysis is a systematic procedure used to explore how an optimal solution responds to changes in inputs – which are typically either known values which might vary in the future, or parameters whose values are open to question. Thus the analysis is based round a prior assumption that optimisation is centre stage, with uncertainty viewed as a potentially disruptive factor. The analysis aims at discovering how sensitive the ‘optimal’ solution is to changes in crucial factors. An insensitive solution is an advantage and, to add to linguistic confusion, is sometimes termed ‘robust’.

Robustness analysis (after Roy)

This use of the term ‘robustness analysis’ entered the literature some 13 years after it was first introduced in the sense employed in this article. As with sensitivity analysis this approach seeks to incorporate the real world experience of uncertainty into the understanding of mathematically derived results. It differs from sensitivity analysis in two ways. The first difference is that it aims to handle not only optimisation but a range of other computational results – eg that a certain solution is feasible, or that it is near optimal. The second is that its perspective is virtually the mirror image of that of sensitivity analysis. This is to identify the domain of points in the solution space for which a particular result continues to hold. Uncertainty, however, remains attached to parameter values, rather than to the swathe of intangible uncertainties that may be resistant to credible quantification. And as with sensitivity analysis, the idea of exploiting sequentiality to achieve flexibility is absent.
The purpose of this comparison is not to criticise these formulations, but by distinguishing robustness analysis (in the sense of this article) from them to clarify its characteristics. Each of them performs functions which robustness analysis does not attempt, and vice versa.

In conclusion

For a more extended introduction to robustness analysis, see Rosenhead (2001a, 2001b). Fuller references are available there.
This summary has been couched largely in terms of the practicalities of decision-making. A more polemical

case, but no less legitimate, could be advanced in the language of sustainable development. To quote Russ Ackoff (1988)

"The freedom to decide, to make choices, is for me the most important freedom people of any age can have. But this freedom is empty without alternatives from which to choose. To deprive future generations of options is a deprivation of their rights."

REFERENCES

R.L. Ackoff 'The Future is now', Systems Practice 1(1), 1988, pp. 7-9.

J.Rosenhead (2001a) 'Robustness analysis: keeping your options open'. In J. Rosenhead and J. Mingers (eds.) Rational Analysis for a Problematic World Revisited: problem structuring methods for complexity, uncertainty and conflict, Wiley, Chichester, 2001, pp. 181-207.

J.Rosenhead (2001b) 'Robustness to the first degree'. In J. Rosenhead and J. Mingers (eds.) Rational Analysis for a Problematic World Revisited: problem structuring methods for complexity, uncertainty and conflict, Wiley, Chichester, 2001, pp. 209-223.

J. Rosenhead and J. Mingers (eds.) Rational Analysis for a Problematic World Revisited: problem structuring methods for complexity, uncertainty and conflict, Wiley, Chichester, 2001.

ACKNOWLEDGMENTS

My thanks are due to John Howard and Larry Phillips for advice on statistical concepts touched on in this article.

EWG-MCDA Newsletter, Fall 2002