FORUM

Robustness Analysis: A Bayesian point of view

David Ríos Insua

Decision Engineering Lab,

U. Rey Juan Carlos-DMR Consulting Foundation

Spain

Introduction

The latest issues of the EWG-MCDA Newsletter have presented several views on what is robustness analysis. In this note, I would like to complement them describing what is understood by such term, within the Bayesian arena. Further details may be seen in Ruggeri et al (2004). The Bayesian approach to inference and decision analysis, see French and Rios Insua (2000), essentially suggests:

- Modelling beliefs about a parameter of interest through a prior which, in presence of further information, is updated to the posterior.
- Modelling preferences and risk attitudes about (multicriteria) consequences, with a multiattribute utility function.

- Associate with each alternative its (multiattribute) posterior expected utility.
- Propose the alternative which maximises the posterior expected utility.

As in any quantitative approach, there are many reasons to check the sensitivity of the output (the optimal alternative) with respect to the inputs (model, beliefs and preferences). In addition, since, in this framework, inputs to the analysis encode the DM's judgements, she should wish to explore their implications and possible inconsistencies. The need for sensitivity analysis is further emphasised by the fact that the assessment of beliefs and preferences is a difficult task. This is an especially important point, as her judgements will evolve through the analysis until they are *requisite*. Robust Bayesian analysis guides this process.

The usual practical motivation underlying robust Bayesian analysis is the difficulty in assessing the prior distribution. Consider the simplest case in which it is desired to elicit a prior over a finite set of states Θi, $i=1,\ldots,I$. A common technique to assess a precise $\Pi(\Theta i)= pi$, with the aid of a reference experiment, proceeds as follows: one progressively bounds $\Pi(\Theta i)$ above and below until no further discrimination is possible and then takes the midpoint of the resulting interval as the value of pi. Instead, however, one could directly operate with the obtained constraints $\alpha i <= \Pi(\Theta i) <= \beta i$, acknowledging the cognitive limitations.

The same situation holds when modelling preferences. One might assess the utility of some consequences through, say, the certainty equivalent method, and then fit a utility function. However, in reality, we only end up with upper and lower constraints on such utilities, possibly with qualitative features such as monotonicity and concavity, if preferences are increasing and risk averse. These constraints can often be approximated by an upper and a lower utility function, leading to the consideration of all utility functions that lie between these bounds. If a parametrised utility function is assessed, the constraints are typically placed on the parameters of the utility, say the risk aversion coefficient. Of course, in developing the model for the data itself there is typically great imprecision, and a need for careful study of model robustness.

A final comment concerning the limits of elicitation concerns the situation in which there are several decision makers and/or experts involved in the elicitation. Then it is not even necessarily possible theoretically to obtain a single model, prior, or utility; one might be left with only classes of each, corresponding to differing expert opinions.

Basic concepts

Robust Bayesian analysis provides tools to check the impact of the utility function, the prior and the model on the optimal alternative, and its posterior expected utility. We distinguish three main approaches to Bayesian robustness. We illustrate it considering robustness with respect to changes in the prior, but similar issues are raised when considering likelihoods and utilities. A “guided tour” through these three approaches is presented in Berger et al. (2000) and the references therein.

Informal approach

The first approach is the *informal* one, which considers several priors and compares the quantity of interest (e.g., the posterior mean) under them. The approach is very popular because of its simplicity. Its rationale is that since we shall be dealing with messy computational problems, why not analyse sensitivity by trying only some alternative pairs of utilities and priors? While this is a healthy practice and a good way to start a sensitivity analysis, in general this will not be enough and we should undertake more formal analyses: the limited number of priors chosen might not include some which are compatible with the prior knowledge and could lead to very different values of the quantity.

It is worth mentioning that the consideration of a finite number of utilities links clearly with multi-objective decision making problems.

Global robustness

The most popular approach in Bayesian robustness is called *global sensitivity*. All probability measures compatible with the prior knowledge available are considered and robustness measures are computed as the prior varies in a class. Computations are not always easy since they require the evaluation of suprema and infima of quantities of interest.

The robustness measures provide, in general, a number that, in principle, should be interpreted in the following way:

1. if the measure is “small”, then robustness is achieved and any prior in the class (possibly one computationally simple) can be chosen without relevant effects on the quantity of interest,

2. if the measure is ``large", then new data should be acquired and/or further elicitation narrows the class, recomputing the robustness measure and stopping as in the previous item; o.w. ….

3. …. if the measure is “large” and the class cannot be modified, then a prior can be chosen in the class but we should be wary of the relevant influence of our choice on the quantity of interest. Alternatively, we may use an ad hoc method such as the *G-minimax*, to select an alternative.

Given a class *G* of prior measures, global sensitivity analysis will usually pay attention to the range of variation of a posterior (or predictive) functional of interest as the prior ranges over the class.

Local robustness

The last approach looks for *local sensitivity* and studies the rate of change in inferences and decisions, using functional analysis differential techniques, such as Frechet derivatives of posterior expected utilities and their norms, total derivatives or Gateaux differentials.

Decision and utility robustness

An important and occasionally controversial issue is the distinction between decision robustness and expected utility robustness. A variety of situations may hold. For instance, when performing sensitivity analysis, it may happen that expected utility changes considerably with virtually no change in the optimal Bayes action, even if the utility is fixed.

Foundations

A number of results show that we may model imprecision in beliefs and preferences through a class of probability distributions and a class of utility functions. These results have two basic implications. First, they provide a qualitative framework for sensitivity analysis, describing under what conditions we may undertake the standard and natural sensitivity analysis approach of perturbing the initial probability-utility assessments, within some reasonable constraints. Second, they point to the basic solution concept of robust approaches, thus indicating the basic computational objective in sensitivity analysis, as long as we are interested in decision analytic problems: that of *non-dominated alternatives*. This corresponds to a Pareto ordering of decision rules, see White (1982), based on inequalities on the posterior expected utility.

As a consequence of this model, the solution concept is the set of non-dominated alternatives. In some cases, non-dominance is a very powerful concept leading to a unique non-dominated alternative. However, in most cases the non-dominated set will be too large to imply a final decision. It may happen that there are several non-dominated alternatives and differences in expected utilities are non-negligible. If such is the case, we should look for additional information that would help us to reduce the classes, and, perhaps, reduce the non-dominated set. Some tools based on functional derivatives to elicit additional information may be seen in Martín and Ríos Insua (1997). Tools based on distance analysis may be seen in Ríos Insua (1990).

Stability Theory

Stability theory provides another unifying, general sensitivity framework, formalising the idea that imprecisions in elicitation of beliefs and preferences should not affect the optimal decision greatly. When *strong stability* holds, careful enough elicitation leads to decisions with expected utility close to the smallest achievable; when *weak stability* holds, at least one stabilised decision will have such property. However, when neither concept of stability applies, even small elicitation errors may lead to disastrous results in terms of large losses in expected utility.

Conclusion

The approach we propose may be summarised as follows: at a given stage of analysis, we elicit information on the DM's beliefs and preferences, and consider the class of all priors and utilities compatible with such information. We approximate the set of non-dominated solutions; if these alternatives do not differ too much in expected utility, we may stop the analysis; otherwise, we need to gather additional information, possibly with the tools outlined above. This would further constrain the class: in this case the set of non-dominated alternatives will be smaller and we could hope that this iterative process would converge until the non-dominated set is small enough to reach a final decision. It is conceivable in this context that at some stage we might not be able to gather additional information yet there remain several non-dominated alternatives with very different expected utilities. In these situations, $L \times G$*-maximin solutions* may be useful as a way of selecting a single robust solution. We associate with each alternative its worst expected utility; we then suggest the alternative with maximum worst expected utility.

Acknowledgements

Supported by grants from MCYT, URJC and DMR Consulting Foundation.

References

1. Berger, J.O., D. Rios Insua, and F. Ruggeri (2000). Bayesian robustness. In D. Rios Insua and F. Ruggeri, eds., *Robust Bayesian Analysis*, Springer-Verlag, New York, USA.

2. French, S. and Rios Insua, D. (2000) *Statistical Decision Theory*, Arnold.

3. Martin, J. and Rios Insua, D. (1997) Local sensitivity analysis in Bayesian Decision Making (with discussion) en Berger et al (eds) *Bayesian Robustness*, 1997, Institute of Mathematical Statistics (con J. Martin), 119-135.

4. Rios Insua, D. (1990). *Sensitivity Analysis in Multiobjective Decision Making*. Springer-Verlag, New York, USA.

5. Rios Insua, D., and F. Ruggeri (2000). *Robust Bayesian Analysis*. Springer-Verlag, New York, USA.

6. Ruggeri, F., Rios Insua, D. and J. Martin (2004) Robust Bayesian Analysis, in Dey (ed) *Handbook of Statistics: Bayesian Analysis*, North Holland. White, D.J. (1982). *Optimality and Eficiency*. Wiley, New York, USA.