Responses to reviewers' comments on "Boundary linear utility and sensitivity of decisions with imprecise utility trade-off parameters"

Malcolm Farrow and Michael Goldstein

May 20, 2009

Review 1

1. *I would have liked to have access to reference [4], that would have thrown light on Section 4, but unfortunately I could not find it. . . . Perhaps the authors can include a new reference to this workshop or other paper that is already published.*

   The paper has now appeared in print and the reference has been changed accordingly. A copy can be obtained from the following Web address.

   `http://www.mas.ncl.ac.uk/~nmf16/ISIPTA09/`

2. *I suggest that the authors consider improving the readability of the paper, for example introducing the practical example earlier and using it to illustrate the main parts of the paper, as they did in their previous work.*

   We have spread the example throughout the paper to illustrate the various parts.

Review 2

General points

1. *While the article looks carefully written and elaborated in detail, it lacks some didactical impetus to relate to ISIPTA community. It is more on imprecise utility rather than imprecise probabilities (IPs) that come in (at least to my understanding) only as analogies. One may argue that therefore the article misses the community. (. . .) The authors must take a much more severe effort to address the IP community, as I will outline below.*

   The paper is concerned with imprecise utility, more precisely imprecise trade-offs in multi-attribute utility, rather than imprecise probability. However "imprecision in utilities and expected utilities" was listed as a theme in the Call for Papers for the symposium. Furthermore we feel that probability and utility are linked concepts and issues of imprecision in one are closely related to issue of imprecision in the other. In this paper we do not consider imprecision in the evaluation of expectations of utilities but this would be a natural next step for us.

Our responses to a number of other specific points are given below.

2. *On a more fundamental level, I am not sure whether the metrics introduced to demonstrate robustness, are in-line with the IP community's fundamentals: The fact that the λ-volume is equivalent to some derivative of a uniform distribution sounds somewhat "objective Bayesian-type style", the IP community is just happy to have got rid of. Some axiomatic background on the choices of sensitivity measures would be very welcome.*

   The role of the formal decision analysis is to act as a decision support tool which identifies for the decision maker those aspects of the decision choice which are straightforward and those aspects for which the decision choice is delicately poised and which therefore would benefit from more careful consideration. In our analysis, if the form of the feasible region was simple to display graphically, then we would draw graphs showing the way that our decision preferences changed as we moved across the feasible region. As the feasible region is complicated and high dimensional, and we have many decision contrasts to consider, we summarise this graphical display in terms of volumes of the feasible region associated with preference for the competing alternatives. These sub-volumes are complicated to calculate and we may potentially require such comparisons for many decision pairs. Therefore, we exploit the formal equivalence between these volume calculations and the probabilistic calculations that we would make to evaluate the probabilities of certain events given a uniform probability specification over the elements of λ. This allows us to use the local computation properties of the implied graphical model to simplify the volume calculations that we require, and the probability specification is simply used as a device to reduce a complicated expression into its component parts.

   We have modified the wording of section 4.2.

3. *Finally, the article would benefit a lot by making much clearer what results refer to imprecise (multi-attribute) utilities and what comes in specifically from the hierarchical structure. To my taste, it is the latter that could lead (has led) to highly non-trivial, timely insights.*

   This work follows the development of Farow and Goldstein (2006). In that paper, we firstly develop the notion of imprecise utility in general, then determine the implications of our construction for mutually utility independent attributes with imprecision over the tradeoff parameters, and finally we extend these conditions to utility hierarchies with such tradeoff imprecision at each level in the hierarchy, tying the different levels together by our use of the standard scale at each node. All of the results on which we rely in this paper are explicitly based on the final stage in this sequence, and space does not allow us to develop a parallel strand when we remove the hierarchic imprecise structuring that we have imposed, as such a case will, in general, be far more complicated to consider.

See also our response to point 1b below.

Corrections

1. *Section 1, §1:*

   (a) *What are "rules"? Mappings from the stated preference order onto decision space?*

   We have used "rules", "choices" and "alternatives" more or less interchangably. We have added a comment in Section 1 to make clear that a "rule" is just a "choice".

   (b) *Motivate "utility independence"!*

   We have added some sentences in section 2.1.

   (c) *Does "Pareto optimality" refer to competing attributes?*

   We mean Pareto optimality with respect to the set of possible trade-off specifications. We have added a phrase in Section 1 to clarify this.

2. *Section 1, §3: "need for practical ways": this sounds like a somewhat arbitrary simplification that could well be as ad hoc as specifying a precise utility (although that precision is spurious). Please explain why, after having implemented that "practical way", the decision situation is much better modelled than by precise utilities.*

   When utility is imprecise, we do not always have strict dominance relations between competing decisions. Therefore, we need to introduce the notions of almost dominance, by allowing the analyst to identify those differences in utility which would have sufficiently small practical consequence that he would be prepared to discount such differences in the analysis. The decision problem is better modelled in this way than using precise utilities in those cases where, firstly, there are no precise utilities to which we can appeal and, secondly, the analyst is prepared to make such judgements as to the consequences of small differences in outcome.

   Along the same line: the quest for "simplicity" in §4.

   This "simplicity" refers to the choice itself, not the method for choosing it or for examining its robustness. We just meant that we might have secondary criteria, which we have not used in the analysis so far. If we can not choose between two alternatives on the basis of the primary criteria, we might use some secondary criterion, "such as simplicity." To avoid misunderstanding, we have removed this phrase.

3. *[4] is excessively cited. Please make that ref (that is announced as being "in press") somehow available.*

   The paper has now appeared in print and the reference has been changed accordingly. A copy may be obtained from the following Web address.

   `http://www.mas.ncl.ac.uk/~nmf16/ISIPTA09/`

4. *Section 2.1: try to replace "complementary", as this is a fixed terminus in Physics.*

   We use the word "complementary" here in the sense used by Keeney and Raiffa (1976). We have not found a satisfactory concise replacement.

5. *Section 2.1: Eq. (1) in conjunction with Eq. (4) looks confusing. Could one save one of the two Eqs.?*

   We have changed these equations to clarify them but we think that it is necessary to have two equations to make clear the scaling which we are applying.

6. *Section 2.1, last §: Are c_n and C_n meant to denote interval limits? They are ill-defined.*

   They are not interval limits. They are reference outcomes used to define the scaling of the utility function. We have changed the wording to clarify this.

7. *Section 2.2: I did not find out what the authors exactly mean by "imprecise utilities" (or trade-off parameters): just intervals? Or imprecise measures that then lead to the need to utilise expectations of utilities (as noted in §3)? I would appreciate if once the authors clearly wrote what are the constraints that are elicited on the trade-off parameters (or on combinations of u's)?*

   In standard utility theory, the decision maker may state preferences between all combinations of outcomes for the decision problem. In our development of imprecise utility, we suppose that the decision maker can state preferences for some, but not all, outcomes, and, as with imprecise probability, our imprecise utility is defined by obeying all of the constraints implied by the stated preferences. In the mutually utility independent case that we consider, we suppose that the decision maker can make preference statements over all combinations of the individual attributes, and so may create precise utility specifications for each individual attribute, but can only make preference statments for some, but not all, combinations involving more than one of the various attributes. Each such preference statement imposes constraints on the tradeoff parameters which are used to combine the individual attributes into a multi-attribute utility, and the imprecise utility that we refer to is utility defined by obeying all of the constraints implied by the collection of stated preferences. The precise implications of all of the stated constraints are as described in section (2.4)

8. *Section 2.2 :*

   (a) *What do the authors mean by "unconstrained elicitations" (§4)?*
   We mean that, for each additive or binary child node, we state whichever preferences we wish between pairs of utility vectors for the parent nodes. We have changed the wording to clarify this.

   (b) *Add a blank before last sentence.*
   Done.

9. *Section 3.1: Is a "feasible region" the same as a "feasible set" introduced in S2.2?*

   Yes. We have changed "set" to "region."

10. *By the procedure described after Eq5: does d become a function of λ ?*

    The identity of d which maximises $\bar{U}_{d,\lambda}$ does depend on the choice of λ.

11. *S3.1, § "in [3]": This discussion is confusing, as the reader learned in S2.2, that an elicitation results in a certain polyhedronic structure of R – so why do we need now additional sources of justification, why do we want the whole boundary linear U-approach at all? The logical link between S2&3 is poorly explained in the intro of S3. Some relief comes from S3.2, §4, where important equivalences are stated. Basically, sweeping over R is equivalent to sweeping over the λ-simplex, isn't it? So we gain a "natural unit"? What is a Bayes decision?*

    In section 2, we establish the structure which follows directly by imposing the requirement of consistency with all of our stated preferences. This imposes constraints on our preferred decision choice, but need not uniquely determine our preferred choice. Therefore, it is of interest to consider whether there are any additional principles which we may consider sufficiently reasonable that we would like our decision choices to respect them. This is the role of the Boundary Linear Utility, which we define in Section 3. As we say in the text, "In [3], the boundary linear form is motivated by various axiomatic and natural requirements for the combination of group preferences." Space does not allow us to recapitulate all of these axiomatic arguments, but they are given in detail in the reference.

    A Bayes decision is a choice which maximises expected utility. Again, unfortunately, we are limited by space. For this usage, in terms of utility, see, for example, Smith (1988), p33-34.

12. *S3.2, §5 (§ "From this"): Does λ run to "s" or to "r" ? I'd expect the latter. Also, it seems that U_i denotes very different items: compare Eq.2&5!*

    Corrected. It should be r.

13. *Please explain in 1-2 sentences conditional independence graphs; otherwise the reader misses the interesting calculations in S4.2.*

    We have changed the wording so that the required property is now clearer.

14. *Table 1 is missing.*

    It was not missing. It appeared on the next page and therefore after Table 2. We have moved it so that it now appears before Table 2.

15. *Spend a state-of-the-art Figure caption in Fig.2.*

    We have expanded the caption.

16. *The conclusions are rather sparse and should especially be expanded before the last sentence: what are the non-trivial concrete findings on the sensitivity studies?*

    We have expanded the conclusions, refering specifically to the example, within the space available.

17. *Perhaps the last §before S3.2 could be saved.*

    We have shortened this paragraph. We have retained a shorter version of the paragraph because we feel that it provides a link with other work in the area of imprecision in multi-attribute decision problems.

Review 3

1. *Section 3.2:*

    - *2nd paragraph, in (ii), remove first comma (it could be understood that the constraint $\lambda_i > 0$ also applies to (i), in which case (i) would imply (ii) and hence be redundant)*

      Done.

    - *regarding the proof of the multiplicative form: what if some of the λ's are zero?*

      We have altered the wording slightly. At each node there must be at least one vertex with nonzero weight and this can be used as the reference vertex i'_n.

    - *4th paragraph, typo: "$\lambda_1, \ldots, \lambda_s$ " → "$\lambda_1, \ldots, \lambda_r$"?*

      Corrected.

2. *Section 4.2:*

    Last paragraph of left column, first line: brackets missing around superscript n

    Corrected.

3. *Section 4.3:*

    Is λ' the transpose of λ?

    Yes. Note added.

References

Farrow, M. and Goldstein, M., 2009. Almost-Pareto decision sets in imprecise utility hierarchies. *Journal of Statistical Theory and Practice*, 3, 137-155.

Smith, J.Q., 1988. *Decision Analysis: A Bayesian Approach*, London: Chapman and Hall.