I²SDS
The Institute for Integrating Statistics in Decision Sciences

Technical Report TR-2008-7
June 9, 2008

Justified and Unjustified Rank Reversal in Decision Aid Methodologies and Independence of Irrelevant Alternatives

David Zalkind
Department of Decision Sciences
The George Washington University, Washington, DC

Justified and Unjustified Rank Reversal in Decision Aid Methodologies and Independence of Irrelevant Alternatives

David Zalkind
Decision Sciences
George Washington University

Introduction

The notion of the "independence of irrelevant alternatives" is that a decision maker's preference ordering between two alternatives, should remain unchanged if a n additional alternative is made available that is dis-preferred to each of the first two. These dis-preferences presumably should make the new alternative irrelevant to a decision maker. The introduction of an irrelevant alternative is not supposed to change previously determined preference rankings among options that are preferred to it. If indeed preference orderings are changed among these options with the introduction of an irrelevant alternative, it is called *rank reversal*. Rank reversal is supposed to be a bad thing. Advocates of Multi-Attribute Decision Making (MAUT) have criticized the Analytic Hierarchy Process (AHP) as being a methodology that permits rank reversal when irrelevant alternatives are made available to the decision maker (Dyer 1990). Advocates of AHP have criticized MAUT for having the same characteristic (Harker and and Vargas 1990). However, Gass (2005) appears to acknowledge that MAUT with swing weights (MAUTS), does not have rank reversal issues. In response to rank reversal criticisms, advocates of AHP developed an alternative mode of AHP with mathematics changed sufficiently to avoid rank reversal when an irrelevant alternative is introduced, described in Saaty (1990)

Forman and Gass(2001) raise the question "Do we forbid the use of a method that allows rank reversal of alternatives to occur even though there are reasonable real-world situations in which rank reversal can and should occur?"

I argue that there is more than one kind of rank reversal in the presence of irrelevant alternatives. One kind is justifiable. It can occur when there is *appropriate re-consideration or new consideration* of assessments by the decision maker because the alternative provides new information or perspective about the original choices. Decision aid methodologies and decision makers should not be criticized for permitting or exhibiting this behavior. I believe that another kind of rank reversal in the face of an irrelevant (i.e., dis-preferred) alternative is *unjustified*. Its cause may be confusion or distraction on the part of the decision maker. Its cause may be an artifact of the decision aid methodology. If it is the latter, than I believe there is a failing in the decision aid methodology. If it is the former, it may be a failing on the part of the decision aid methodologist to keep the decision maker focused. I provide examples and analyses for each.

One-Dimensional Case

In their seminal work, Luce and Raiffa (1957) first mention the concept of independence of irrelevant alternatives in their axiomatic treatment of utility theory (p. 27). However, they do not fully discuss the concept until presenting their axioms for individual decision making under uncertainty (p. 288). They give the following example illustrating the case where the new alternative conveys information about original choices. This example is now obviously quite old:

> A gentleman wandering in a strange city at dinner time chances upon a modest restaurant which he enters uncertainly. The waiter informs him that there is no menu, but that this evening he may have either broiled salmon at $2.50 or steak at $4.00. In a first-rate restaurant his choice would have been steak, but considering his unknown surroundings and the different prices he elects the salmon. Soon after the waiter returns from the kitchen, apologizes profusely, blaming the uncommunicative chef for omitting to tell him that fried snails and frog's legs are also on the bill of fare at $4.50 each. It so happens that our here detests them both and would always select salmon in preference to either, yet his response is "Splendid, I'll change my order to steak." Clearly, this violates the seeming plausible axiom 6. [**Axiom 6**: Adding new acts to a d.p.u. u. [decision problem under uncertainty], each of which is weakly dominated by or is equivalent to some old act, has no effect on the optimality or non-optimality of an old act.]

In their discussion of this situation, Luce and Raiffa write (p. 288):

> …can we really argue that he is acting unreasonably? He, like most of us, has concluded from previous experience that only "good" restaurants are likely to serve snails and frog's legs, and so the risk of bad steak is lessened in his eyes.
>
> This illustrates the important assumption implicit in axiom 6, namely that adding new acts to a d.p.u.u. *does not alter one's a priori information to*

which is the true state of nature. In what follows, we shall suppose that this proviso is satisfied. In practice this means that, if a problem is first formulated so that the availability of certain acts influences the plausibility of certain states of nature, then it must be reformulated by redefining the states of nature so that the interaction is eliminated.

Axiom 6 can be strengthened to the following form of the principle of the independence of irrelevant alternatives:

Axiom 7. If an act is non-optimal for a d.p.u.u., it cannot be made optimal by adding new acts to the problem.

Luce and Raiffa's contention that the problem may have to be reformulated in light of new information is a crucial point. I believe that failure to do allow the decision maker the opportunity to do so in the face of changed circumstances would be inexcusable if there were a decision aid methodologist involved in assisting with the analysis. Thus, the newly available dis-preferred alternative is not "irrelevant" at all in the common meaning of the word. It may lead to a reformulation of the problem. Furthermore, that reformulation may result in a rank reversal among the previously considered alternatives. I consider this to be *justifiable rank reversal*.

This process of re-thinking should not be considered to violate any axioms or desirable criteria for a decision aid methodology. I believe it inappropriate to criticize a decision aid methodology by creating examples that assume that there will not be new structuring

or new assessments for a decision problem in the face of new circumstances that may provide information or perspective to the decision maker.

We illustrate the Luce and Raiffa scenario with a pair of decision trees. We do not assume any particular methodology for measuring or scaling "values" of alternatives or outcomes, but it is assumed that there is sufficient cardinal content in the numerical representations of the values that it is mathematically reasonable to calculate weighted sums of the values, e.g., expected values. For simplicity, we will collapse the frog's legs and fried snails options into a single option labeled Flfs. "Before":

MODEL 1

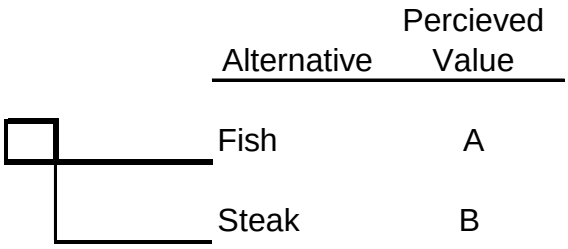

where A > B .

The apparent "After" picture:

MODEL 2

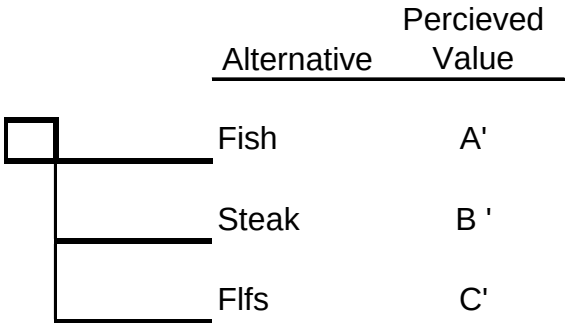

Where A, B, A', and B' are all > C', but A > B while A' < B' . Consternation at this rank reversal for Fish and Steak is understandable, if this is the only picture that we are allowed. However, permitting reformulation and new assessments leads us to another model that appears to be a better representation of the situation, simplified to ease exposition:

MODEL 3

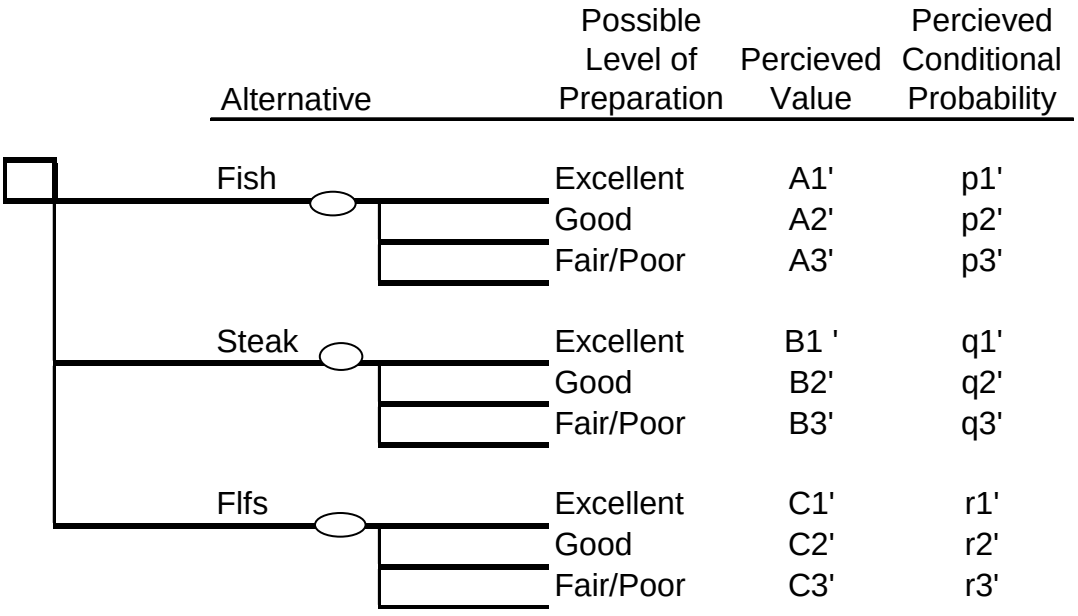

Based on the scenario description, we have ample reason to believe that A1', A2', A3', B1', B2', and B3' are each greater than each of C1', C2', and C3'. (I imagine that the decision maker's preferences could be C1' > A3', but with a trivial subjective probability for each of these outcomes.)

Model 3, which supersedes Model 2 leads us to the realization that Model 1 could be superseded by Model 4.

MODEL 4

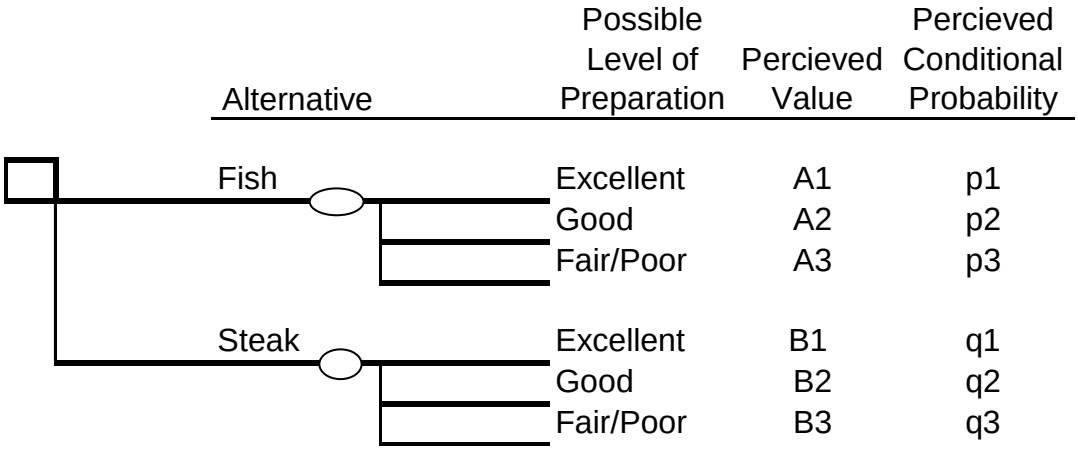

Now that "Before" is represented in Model 4, there is no question that in a Bayesian world updates to the probabilities p1, p2, p3, q1, q2, and q3 should be permitted when new information is made available. These updates are, respectively, p1', p2', p3', q1', q2', and q3'. There may be a reasonable argument that there is no basis for the decision maker to change original *value* assessments, as opposed to probability assessments. That is, we might expect that A1' = A1, A2' = A2, A3' = A3, B1' = B1, B2' = B2, and B3' = B3. However, it should not be a surprise nor a contradiction of any axioms or criteria for a decision aid methodology that rank reversal between the alternatives Fish and Steak may occur when the probabilities are reassessed. Discussion of the possible nebulous nature of the probabilities in the mind of the decision maker is beyond the scope of this paper.

To reiterate, the Luce and Raiffa scenario is an example of *justifiable rank reversal* when an irrelevant alternative is made available. It can be recognized as such through revising the initial implicit or explicit decision model.

As a counterpoint to the justifiable rank reversal in the Luce and Raiffa scenario, consider this futuristic example:

A flying saucer lands in a field where a decision maker happens to be walking. A space alien emerges and tries to communicate with the decision maker. They have no common language and the meaning of gestures may not be particularly clear to either one, but the decision maker discerns that the space alien is offering him a gift. The decision maker is to choose one of two boxes that look essentially identical to him. The contents of each box are not well understood by the decision maker, but there has been enough communication so that she has acquired some idea (right or wrong) of what might be in each box. She decides that she prefers one of the boxes (and its contents) and indicates this to the space alien. Just as the space alien is about to hand her the box, a second space alien appears in the doorway of the flying saucer, holding another box which appears to be identical to the first two. The second space alien "speaks" to the first one and the first space alien then makes it known to the decision maker that this third box (and its contents) is another possible choice. The decision maker knows nothing about the contents of the third box. The decision maker does not select the third box but now switches her initial choice from one box to the other of the initial two.

I believe that there is a crucial distinction between the Luce and Raiffa example and space alien example. The objects among which the decision maker is deciding are *perceptually invariant* (PI) in the context of the problem. She knows little (with any

confidence) about the initial two boxes and there is no basis for a change in perception when the third box is presented. (Note: We make no inference with regard to preference for the third box and the box that was originally available but not selected. This situation is weaker than the earlier one, when there was dis-preference for the new alternative relative to the original alternatives.)

Decision tree models for the space alien scenario would be identical to Models 1 and 2 except for the labeling. There is no reason to believe that there are the counterparts to Models 3 and 4. Reversal of rank between the original two alternatives is inappropriate when an irrelevant alternative is presented since there are no reasonable inferences that the decision maker can make with regard to the original boxes. If the rank reversal were to occur in this scenario, the responsibility would be with the decision maker, not with a decision aid methodologist. We can distinguish between the Luce and Raiffa scenario and the space alien scenario by recognizing that in the former the initial alternatives had characteristics that were initially not explicitly recognized, but turned out to be relevant. In the latter scenario, there do not appear to be such characteristics.

There is something else in the description of the Luce and Raiffa scenario that shows that it may have been mis-modeled, even by our newer representations. The text shows that the decision maker was concerned about *price*. Perhaps the appropriate model and methodology should be multidimensional (e.g, multi-attribute, multi-objective, etc.) If the availability of new information, derived from a newly available mis-preferred alternative, should lead to going from one dimension to more in the analysis, than a rank

reversal that occurs as a result of incorporating multiple dimensions should not be viewed as unjustified rank reversal (unless there is direct evidence that not enough has changed with regard to assessments to warrant the reversal). As an example, suppose the decision maker is making a choice based on which object is "bigger". Here is what he initially perceives:

MODEL 5

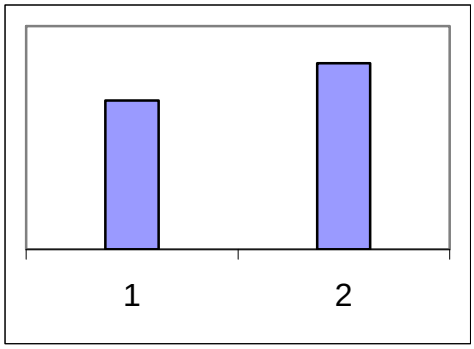

There is no optical illusion here. The decision maker's preferred alternative should be "2". Now, another alternative becomes available. The new object is shorter than both of the original two objects. However, after seeing the new object, it occurs to the decision maker that there is an attribute that he had not previously considered: width. Now, this is what he perceives:

MODEL 6

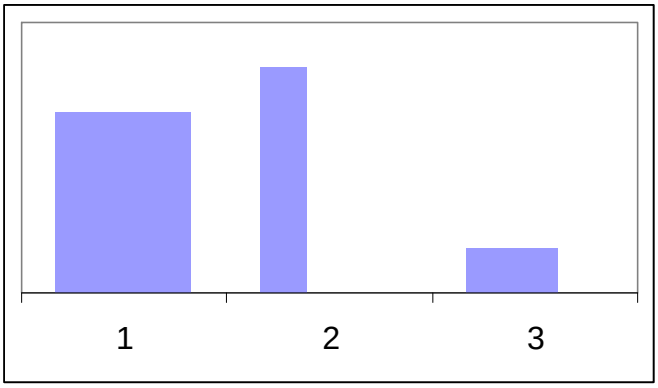

In light of this new view, he may choose to change his original notion of "bigger" that just reflected *height*. Suppose that his new notion of "bigger" reflects both height and width, now that he recognizes that there is such a thing as width and that widths may vary among objects. The notion of "bigger" becomes *half perimeter* rather than *height*, assessed by adding the height and the width of each object. Object 3 is dis-preferred to both objects 1 and 2, no matter which notion of "bigger" is used and thus could be considered irrelevant. However, the perspective inspired by the availability of object 3 leads to a new assessment model and which results in a rank reversal between objects 1 and 2. The second dimension was always present, but was not recognized or perceived as relevant by the decision maker (or by the decision aid methodologist who constructed the assessment model, if any). This scenario represents *justified rank reversal*. New information about the original options has been assessed and incorporated into the basis for assessing preferences.

Multi-dimensional case

We have already illustrated how there can be justified rank reversal when an additional dimension is used in the analysis, when its relevance is newly recognized after the introduction of a dis-preferred alternative. In that illustration, we did not consider relative importance of the old and new dimensions (or attributes or objectives). We assumed that they had equal importance. In general, relative importance of the dimensions is a crucial issue. We are not concerned here with how the decision maker's

importance weights are ascertained, just that they are and that they are used either implicitly or explicitly by the decision maker or the decision aid methodologist.

In the next example, the decision maker is making a choice between two objects based on capacity, attractiveness, and safety. We posit that she is getting professional help for the decision so that there has been some formal process of assessing her perceptions of relative capacity, attractiveness, and safety for each object. Furthermore, there has been an assessment of the perceived relative importances of these three dimensions. Now suppose that a new object is introduced. The decision maker judges that the new object is inferior in every dimension to each of the two original objects, thus it is evaluated as being inferior overall to each of them. In that sense it is irrelevant. Suppose that in the safety dimension, the degree of perceived inferiority sparks a realization by the decision maker that she wants to give more importance weight to this dimension relative to the other two. Perhaps when this is done the calculation that is made by the decision aid methodologist reverses the previous ranking of the original two objects. Is this rank reversal justified?

I argue that the answer is "yes", as a practical matter. It would be unreasonable to prevent the decision maker from re-thinking previous assessments in the light of new insight sparked by a changed situation. This insight is different from the previous examples. It does not involve newly assessing something that was not considered before. It provides an opportunity, decision aid methodologist willing, to re-think something that had already been settled. Any rank reversal that occurs owing to the irrelevant alternative

in this example is not a weakness of either the decision maker or of a decision aid methodology. I believe this because changes to the existing numbers in the original model (e.g., relative preferences within a dimension, relative importances of the dimensions, or probabilities (if any) *are done at the direction of the decision maker*. They are not an artifact of black box calculations that occurred when the decision maker's assessments for the new object were put in the model with *unchanged* (by the decision maker) assessments based on the original objects. (Of course, we may be faced with the situation where the decision maker wants to change importance assessments without any apparent relevant inspiration. This should also be permitted, although it might distress the decision aid methodologist. For example, it is not unreasonable to have "second thoughts" after "sleeping on it".)

Discussion and Conclusion

To deal with the issue of rank reversal, the concepts of Ideal and Distributive Synthesis modes were developed for AHP (Forman and Selly, 2001). These involve different ways of assigning priority weights to alternatives within each objective. However, the existence of more than one mode in AHP requires recognition and choice on the part of the decision maker with regard to which mode to invoke, depending on whether the decision environment is open (e.g., unlimited resources) or closed (e.g. limited resources). The former mode avoids rank reversal when irrelevant alternatives are introduced. MAUTS does not have a requirement for the decision maker to distinguish between environments. (The implications of the identification of multiple modes are outside the scope of this paper.)

The existence of a prohibition of rank reversal, whether or not the reversal is justified, may be hidden in the underpinnings of a decision aid methodology. With regard to two of their utility theory assumptions, #3 for continuity and #4 for substitutability, Luce and Raiffa (p. 27) state that the combination "is reminiscent of what is known in other work as the assumption of the *independence of irrelevant alternatives*". An example of its hidden presence in a more recent work is in the axioms for expected utility presented in Clemen and Reilly (2001). In this book, the authors discuss some of the controversies relating to the axioms or assumptions and state that "Most of us agree that these assumption are reasonable under almost all circumstances" (p. 573). In the sense of this paper, we need to be aware that there could be justifiable and unjustifiable invocations of combinations of axioms that do not appear to be directly related to the independence of irrelevant alternatives and rank reversal.

(There may be technically justified rank reversal that is not compatible with stable decision making. For example, the decision maker might value "change for the sake of change", leading to changing away from the previous top choice. This could have happened in the space alien scenario. This type of preference consideration suggests instability that might be hard to deal with using any standard decision methodology.)

I believe that

(i) Newly presented alternatives, even if dispreferred to all alternatives considered up to now, may provide a decision maker with information or insights not previously known or considered by the decision maker or an assisting decision aid methodologist.

(ii) The decision maker should be permitted to consider reassessments, based on (i). This may lead to a revised model structure, such as in the Models 1 through 4 examples. If these revisions result in a rank reversal between previously considered alternatives, tht should not be seen as problematic.

A rank reversal among the previously considered alternatives, if it occurs, should be considered to be *justified* when the decision maker gleans some insight from the newly introduced alternative pertaining to the original choices or importance of attributes or objectives. A rank reversal should be considered *unjustified* in the presence of an irrelevant alternative, if the decision maker's previous assessments do not change (after she has been given an opportunity to change them) and the decision aid methodology computes a rank reversal among the original alternatives after data for the dis-preferred alternative is incorporated.

Furthermore, an example of what should *not* be considered to be unjustified rank reversal would be a decision problem where there is no recalculation of importance weights permitted by the decision aid methodologist, even though they would change if the new alternative had been available from the beginning. To call a rank reversal unjustified in these circumstances may be unfair to either the decision maker, who is not permitted to do a reassessment of relative importances, or to the decision aid methodology, which is

restricted in some way from doing a full recalculation in the presence of new data, even though none of the original decision maker assessments may have changed.

Similarly, an example of what *should be* considered to be unjustified rank reversal would be a decision problem such that (i)in the presence of a new dis-preferred alternative, the decision maker is given the opportunity to re-think previous assessments; (ii) None of these assessments change; and (iii) The decision aid methodology, after incorporating the data reflecting the new alternative and fully recalculating, reverses the ranks of two or more original alternatives.

References

Clemen, Robert T. and Terrence Reilly. 2001. Making Hard Decisions with DecisionTools®. Duxbury. United States.

Dyer, James S. 1990. Remarks on the Analytic Hierarchy Process. Management Science, Vol. 36, No. 3. 249 -258.

Forman, E., S. I. Gass. 2001. The analytic hierarchy process: An exposition. *Oper. Res.* 49(4) 469–486

Gass, Saul I. 2005, Model World: The Great Debate – MAUT Versus AHP. Interfaces, Vol. 35, No. 4, pp. 308-312. p 309.

Harker, Patrick T. and Luis G. Vargas. Repley to ""Remarks on the Analytic Hierarchy Process" by J. S. Dyer. 1990. Management Science. Vol. 36, No. 3. 269-273

Luce, R. D., H. Raiffa. 1957. *Games and Decisions.* John Wiley and Sons, New York.