Space Systems, Policy, and Architecture Research Consortium
(SSPARC)
"Beta plus" Revised Draft 5/26/04

"SSPARC BOOK" MATERIAL
for Lecture 4

Prepared by:

Hugh McManus
Metis Design

Adam Ross
MIT

Draft for evaluation only. Do not distribute.

TABLE OF CONTENTS

DEDICATION AND NOTE ON SOURCES

This document is an excerpt of a future book or hyper-book on the MATE-CON method. It is provided for class use as a draft. Suggestions for improvement are welcome, as are warnings of errors or omissions. The notes below apply to the entire work in progress; the work or excepts of it should not be reproduced in any form without these notes.

This document is dedicated to the memory of Joyce Warmkessel, a colleague, mentor, and friend to many in the SSPARC and LAI communities. Many of the core ideas behind this work were originally expressed and developed by her, and she was a key mentor and facilitator to the development of all of this work.

The content of this document was developed by the SSPARC consortium. The primary compilers and codifiers of the MATE-CON method were Lt. Nathan Dillar and Adam Ross, in Master's thesis entitled, respectively, "Utilizing Multiple Attribute Tradespace Exploration with Concurrent Design for Creating Aerospace Systems Requirement," [1] and "Multi-Attribute Tradespace Exploration with Concurrent Design as a Value-Centric Framework for Space System Architecture and Design."[2] Major contributors of the original concepts within the method, and/or complimentary methods and tools, include our SSPARC faculty and staff colleagues Elisabeth Paté-Cornell of Stanford University, Joel Sercel and Fred Cullick of Cal Tech, and Amar Gupta of MIT, post-doctoral researcher Bill Kaliardos, and graduate students Jimmy Benjamin, Jason Derleth, Bobak Ferdowsi, Dave Ferris, Russ Garber, Andre Girerd, Seth Guikema, Cyrus Jilla, Chris Roberts, Satwik Seshasai, Nirav Shah, Todd Shuman, Tim Spaulding, Dave Stagney, Dan Thunnissen, Myles Walton, Annalisa Wiegel, and Brandon Wood, along with their advisors and committees. Many other students, staff, and undergraduate researchers also contributed. Bill Borer, Kevin Ray, and John Ballenthin of the Air Force Research Laboratory, Steve Wall of NASA JPL, and Pete Hendrickson of the Department of Defense aided with the development of the method and the development of the case studies. SSPARC research work has been supported by an active group of industry practitioners, through both an Industrial Advisory Board (IAB) and on-site implementation activities.

The text of this manual is built on SSPARC research and member documents. Much of its contents are excerpts, modifications, or paraphrases of published or unpublished work done under SSPARC sponsorship. Every effort has been made to correctly attribute all contributions. Word-for-word excerpts are identified with quotes or indented, with citations. Many other excerpts have been edited to varying degrees and are integrated into the text for clarity. Their sources are cited in the text or in endnotes. Any omissions or errors of attribution should be brought to the authors' immediate attention for correction.

4. UTILITY THEORY

The concept of utility is used to map the attributes of a design to the preferences of the stakeholders.[3] The attributes are the things that the stakeholders care about. Utilities capture how much they desire various values of the attributes in way that can be quantified. A single attribute utility is a normalized measure of preference for various values of an attribute. A multi-attribute utility combines single attribute utilities into a combined metric that can be used to rank user preferences for any set of possible values of the attributes.

This section will provide an overview of concepts, at a level that all members of a MATE-CON team should understand. Much of the heavy pedagogical lifting will be done by the key references, Richard de Neufville's *Applied Systems Analysis: Engineering Planning and Technology Management*,[4] and Ralph Keeney and Howard Raiffa's *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*.[5] Chapters 18-21 of de Neufville are strongly suggested reading, covering the same ground as this section in considerably more detail, and with explicit procedures for collecting user preferences. Keeney and Raiffa state the formal theory behind the method in great detail. It is important that at least some members of the MATE-CON team understand the theory in some depth to avoid methodological errors.

The purpose of using utility theory in MATE-CON is to make better decisions. Both the theory and its typical application have weaknesses that we will not understate. However, utility theory provides a better way at getting at user preferences and needs than most other techniques, and can be quantitatively coupled to other tools and models. The result is more general and versatile than either fixed requirements, as might be used in a traditional analysis, or a rigid "objective function" as might be used in a multi-objective optimization. As we will see, a utility analysis can be reduced to either of these if desired, but this is best done after the user preferences for a wide range of possible concepts is explored.

A caveat is from the preface of Keeney and Raiffa is of particular relevance here:

> The theory of decision analysis is designed to help the individual make a choice among a set of *prespecified* alternatives. Of course, decision analysts do admit that an insightful generation of alternatives is of paramount importance, and they also take note of the often-overlooked fact that good analysis of a set of existing alternatives may be suggestive of ways to augment the set of alternatives. But this is a side point not suitable for development in a preface.

In two sentences, they frame the role of utility theory in the MATE-CON process, and in a third, decline to pursue other aspects of MATE-CON. The utility theory, elegant and complete as it is, will not be useful if the tradespace (containing the prespecified alternatives) is not correctly scoped and assessed. Also, the output of the utility theory will not be the final word on the issue. A stakeholder presented with the results of a tradespace study may well alter *both* his or her view of what alternatives might be interesting (changing the bounds of the tradespace) *and* what his or her true needs are (changing the utility functions). The running examples contain cases of both types of user-driven updating.

4.1.Single Attribute Utilities

Single attribute utilities map single attributes onto user needs. Here, we will explore typical forms of single attributes utilities and provide some practical examples. Note for the purpose of this discussion we are assuming that the attributes and their acceptable ranges have been previously defined.

Proto-Utilities: Functional Requirements

Consider the traditional method of specifying stakeholder needs, the requirement. If only certain values of an attribute result in a useful system, and others do not, one can state this need as a firm requirement. The attribute is a function or output of the proposed system that the user is interested in, so this would be a functional requirement. Figure 4-1 shows some forms of requirements.

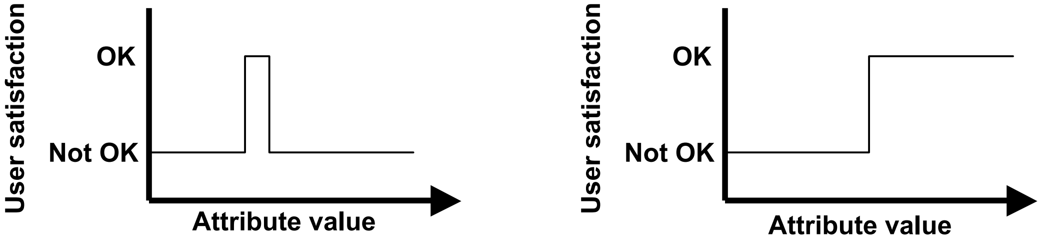

a) Required value (range)

b) Must exceed

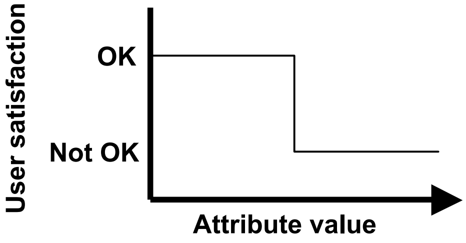

c) Must not exceed

Figure 4-1 Types of Requirements

If the stakeholder really does have a firm need for specific values of a proposed attribute, it is not a good attribute. It should instead be treated as a constraint on the tradespace. Typically, however, the stakeholder needs are not as absolute as Figure 4-1 might suggest. Figure 4-2 shows some more realistic expressions of user needs, superimposed on the artificially binary requirements. Often, failing to meet a requirement does not invalidate the system, although it may displease the user. Also common is the fact that exceeding the requirement would provide additional benefit to the user. Unfortunately, processes based on meeting static requirements often require tedious negotiations if requirements are not met, even if the harm to the user is small, and do not reward "extra" performance at all, even if the benefit to the user is potentially great.

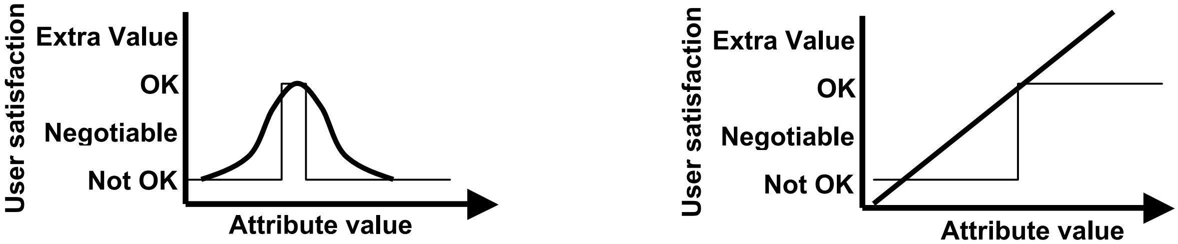

a) Soft target

b) More is better

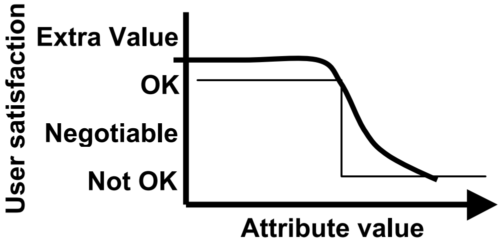

c) Soft threshold

Figure 4-2. True user needs behind requirements

Utilities

Rather than the qualitative measures of user satisfaction shown in Figure 4-2, utilities define a quantitative measure. Somewhat arbitrarily (but usefully) a dimensionless scale from zero to one is used. For any value of an attribute x_i, we define a utility

$$U_i = U_i(x_i) \tag{4-1}$$

We also define the utility to be zero at the lowest (or least desirable) acceptable level of x_i, x_{i*}, and one at the highest (most desirable) level of x_i, x_i^*.

$$\begin{aligned} U_i(x_{i*}) &= 0 \\ U_i(x_i^*) &= 1 \end{aligned} \tag{4-2}$$

Zero represents the lowest possible level of user satisfaction that the user might still consider; it is the bottom end of the negotiable range. This is somewhat non-intuitive, as one is tempted to think zero should mean "no utility", an unacceptable result, but is computationally convenient, and an established convention. Negative values of utility are by definition excluded. One represents full user satisfaction. In some cases (e.g. the "more is better" case from Figure 4-2) this limit may need to be chosen arbitrarily. Values greater than one ("bonus points") create mathematical difficulties and should not be used.

We will concern ourselves with the range of values of attributes that produce utilities between zero and one (the gray box in Figure 4-3). In practice, some designs will be evaluated and found to have values of attributes that do not fall in this range. Values of the attributes that cannot score a utility of zero are not acceptable to the user and the corresponding designs should be excluded from the tradespace. Values of the attributes greater than (or less than, depending on which is the desirable direction) those necessary to score a utility of one need to be handled carefully. There are circumstances in which such designs should be excluded from the tradespace (e.g. if the "excess" attribute is not desirable); more typically this represents excess capacity which has no additional utility to the user but is not a bad thing. In such cases, attributes that are "too good" are simply assigned a utility of one.

The utility scale is dimensionless, and the position of zero is arbitrary. Therefore, care must be taken in interpreting utilities. In some cases, this metric can be given units (e.g. normalized cost per billable minute of a broadband telecom system), in others (e.g. usefulness, to scientists, of scientific data) it can only be used as a relative metric. In the latter case, interpretation of these metrics is somewhat dangerous—a higher metric is better than a lower one, but a utility of 0.5 may not be "half as good" as a 1.0, nor a 0.99 "only 1% off" from a 1.0. Such interpretations usually require returning to the individual metrics, or the decision makers.[6]

The formal requirements for a utility to exist are given in the section below. These requirement and the properties of the utility scale are explored in more detail in de Neufville, Chapter 18.

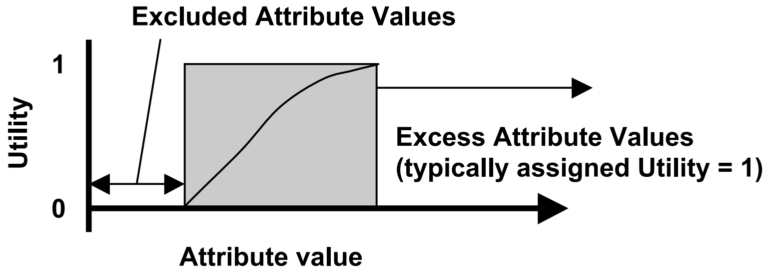

Figure 4-3 Utility as function of a single attribute

Figure 4-4 shows several possible forms of a utility function. The simplest is a linear relation between the attributes and utilities. This form of relationship is common. Also common is the next form, showing diminishing returns for higher attribute values. Threshold, or "S-curve" utilities are also common. The final example shows a non-monotonic function, where the best value is not at one of the extremes. These present difficulties for methods used to both collect the utility functions and to combine them in multi-attribute forms, and so should be avoided if at all possible. Often an attribute with a non-monotonic utility function can be redefined or decomposed into one or two attributes with monotonic utilities.

\

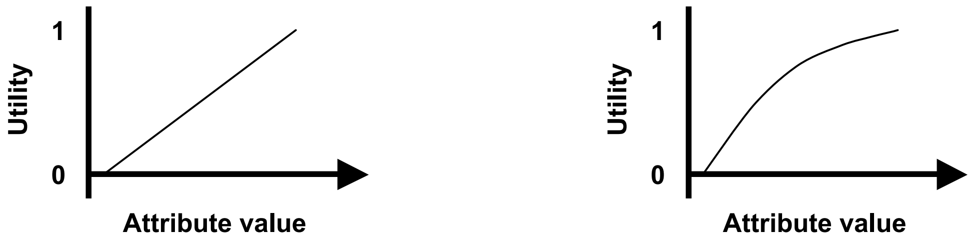

a) linear **b) diminishing returns**

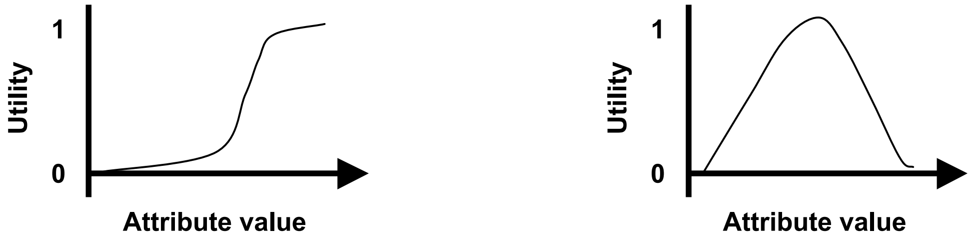

c) threshold **d) non-monotonic**

Figure 4-4 Typical forms of utility functions

Requirements for Single Attribute Utilities

From de Neufville,[4] for a single attribute utility to exist, the attribute must have value to the user in the following senses:

- The user must have a preference for a given value of the attribute over other values of the attribute
- The preferences must be transitive; i.e. if the user prefers A to B and B to C, the user must prefer A to C
- The preferences must be monotonic, i.e. if A is greater than B, and the user prefers A to B, then the user must always prefer a greater value over a lesser.

These criteria simply imply that the function U_i exist for all x_i of interest, and be monotonic.

The utility function is (and must be) an "ordered metric scale," which has the following properties:

- Utilities have meaning only compared to other utilities; they have no absolute value
- The units of utility have constant relative meaning but no absolute meaning.

Mathematically, this implies that the preferences captured by a utility function will not change if the function undergoes a linear transformation, i.e. if

$$U_i' = aU_i + b \tag{4-3}$$

then decisions made using U_i' will be identical to those made using U_i. The analogy to a non-Kelvin temperature scale is exact. The first property implies that "zero degrees" has no physical meaning-it depends entirely on the unit system. The second implies that differences between temperatures DO have physical meaning; it takes the same amount of energy to heat water from 10°C to 20°C as from 30°C to 40°C; and this statement would be equally true in the Fahrenheit system although the numbers would be different. In utility terms, zero utility is defined arbitrarily above (as the lowest utility of interest). The users preference for an attribute with a utility of 0.2 over one with a utility of 0.1 should be just as strong as his or her preference for an attribute of a utility of 0.4 over one with a utility of 0.3 (both are 0.1 apart) and this statement about preferences would still be the same if all the utilities were multiplied by 10.

Additional criteria exist for utilities to be measurable using standard techniques; these are covered in the next section.

Determining Single Attribute Utilities

Often, especially for exploratory studies, simply specifying the form of the utility function with the stakeholders' help and involvement is sufficient. The Spacetug example uses this technique, assuming either diminishing returns or linear utility functions in delta-V, a diminishing return for equipment capability, and a threshold for response time. In cases were the utilities cannot be determined precisely, varying the utilities as part of the tradespace exploration (as was done in the Spacetug example) is appropriate.

If the stakeholder needs are known or can be determined in detail, the Lottery Equivalent Probability (LEP) method for accurately extracting utility functions from participating stakeholders is recommended. This method is described fully in de Neufville, Chapter 19, which is recommended reading. An extended example is found in an excerpt from Ross,[2] covering the X-TOS project; an even more thorough example from the B-TOS system can be found in an excerpt from the B-TOS report.[7] The X-TOS project used software developed by Seshasai.[8] The X-TOS procedure and results are summarized in the example section below.

The LEP method uses questions framed in terms of decisions between two uncertain outcomes ("lotteries") to tease out user utilities in a way that is as free from biases as humanly possible. It is NOT a technique for dealing with true technical uncertainties in the results. These are dealt with later using separate techniques. *The probabilities are simply an artifice for extracting the basic utility curves for the attributes.* This artifice is reinforced in the interviews by having he choices involved be clearly imaginary (see the example). The user is asked to suspend disbelief and answer the questions as accurately as possible.

Use of LEP does, however, place some additional restrictions on the forms of the utilities:

- The attributes must have meaning in the presence of uncertainty, i.e. the statement "there is a 10% probability that the system will have attribute x_i" must be meaningful
- The user must have preference under uncertainty, i.e. must prefer a higher probability of a desirable result over a lower probability
- The user's preference must be linear with probability, *at least within the bounds of the problem stated in the LEP interview.*

The last condition creates some controversy, as users often have non-linear preferences for probabilistic events. Done correctly, the LEP method avoids this problem by:

- Creating an imaginary scenario where the user has no reason to have a true non-linear preference with probability
- Avoiding questions involving very small or very large probabilities (including certainties) which will often invoke either mis-estimations or innate biases in the users.

The LEP method is correctly viewed not as an absolutely accurate method for extracting user utilities, but rather the best available method. A discussion of this point is contained in de Neufville.[4] It has been the practical experience of the SSPARC team that any methodological errors or biases in the LEP method itself are below the level of other sources of "noise" in the measured utilities, discussed next.

The greatest practical difficulty in using the method is finding a representative of the user community that can go through the LEP process successfully. The user must be knowledgeable enough to understand the questions and provide intelligent answers. The user must also be capable of maintaining the detachment from his or her own biases and/or technical problems necessary to go through the interview process. In practice, this requirement has eliminated about half of potential interview subjects for either psychological reasons (e.g. failure to suspend disbelief) or knowledge bias (e.g. too close to some aspect of the problem). The ideal user has been described as a "proxy user"; someone with full knowledge of the needs of the user community and familiarity with the technologies required to solve it, but who does not have a personal stake in either specific user needs or specific technical solutions.

Even given a proxy user or users, a fair amount of measurement noise can be expected. Spaulding[9] did a MATE analysis of a space-based radar. As part of his study, he studied the issue of measurement error in utilities by having the same proxy user do an LEP interview process (using the MIST software) on two separate occasions, and also sketch by hand the utility curves, again on separate occasions and without knowledge of the MIST results. A typical result is shown in Figure 4-5. The utility of an attribute of the space-based radar (coverage in square miles) is shown. The two LEP interviews, labeled MIST 1 and MIST 2, show similar but not identical results from identical interviews. The hand-drawn results show similar noise, and are also consistently somewhat different from the LEP results. The difference between the LEP interview utilities and the hand-drawn ones is typical. In this case, both show diminishing returns, but the hand drawn one is "gentler;" the user seems unwilling to draw the sharp corner that the LEP interview teases out.

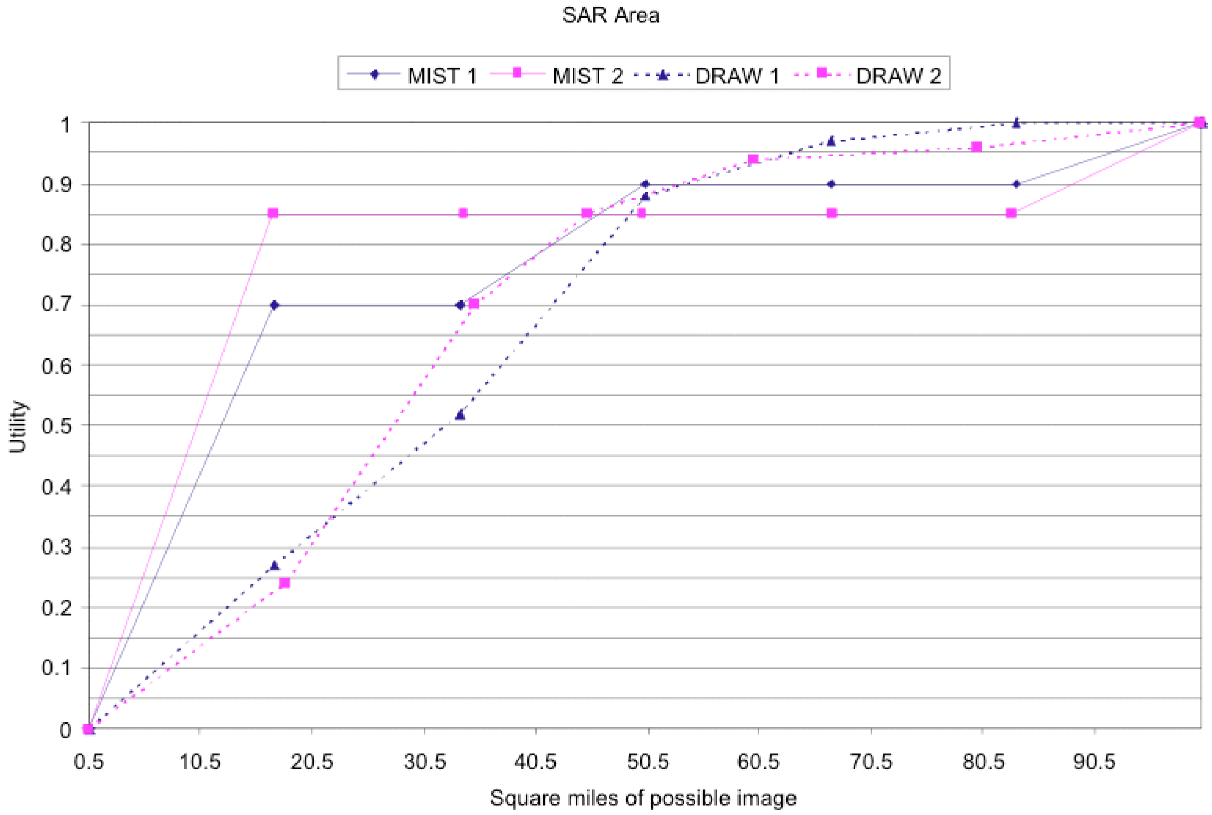

Figure 4-5. Single attribute utilities extracted from the same user (from Ref. 9).

A Note on Risk Aversion

Most texts on utility theory, including de Neufville, include a discussion of risk aversion or risk-seeking. These discussions often lead to needless confusion, due mostly to the historical basis of the terminology. True risk aversion, which can be either a psychological effect or a rational response to a set of specific circumstances, involves valuing a chance of a good outcome at less than its expected value. Risk-seeking or risk-prone behavior (as de Neufville points out, this does not mean recklessness!) involves valuing a chance of a good outcome at more than its expected value. These are both illustrated in Figure 4-6 below. Risk neutral preferences are linear with probability—they express the expected value of the outcome. They are sometimes described as "rational" preferences, although this terminology is also misleading, as risk prone or risk adverse preferences may be entirely rational under some circumstances.

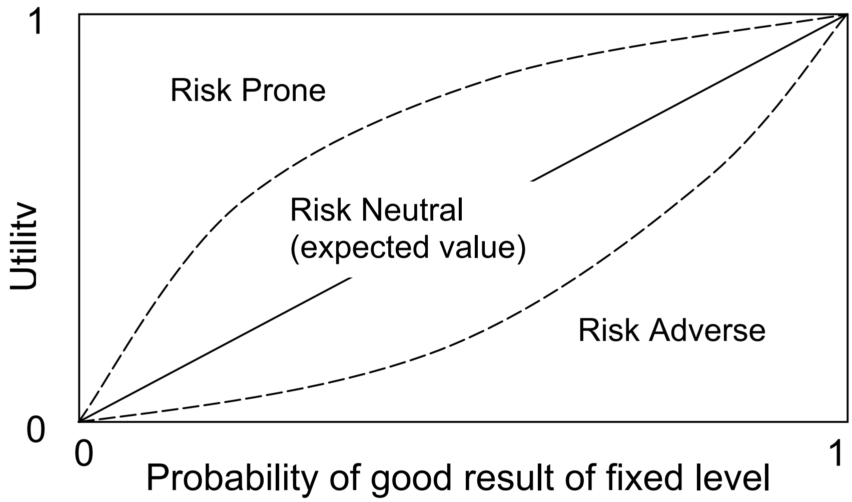

Figure 4-6. Risk adverse, risk prone, and risk neutral behaviors

Due to historical use of terminology, the terms risk-adverse and risk-prone are often used to describe the shapes of SAU curves that do not explicitly include a probabilistic component. In Figure 4-7, SAUs are plotted that show diminishing, linear, and increasing utilities with performance level. *If one assumes that increasing performance level implies increasing risk of failur*e, then a diminishing returns curve (which favors, under this assumption, the more-certain attainment of a lower performance level) could be described as expressing risk-adverse behavior. Likewise, a preference for high performance (with the implication of high risk) could be described as risk seeking.

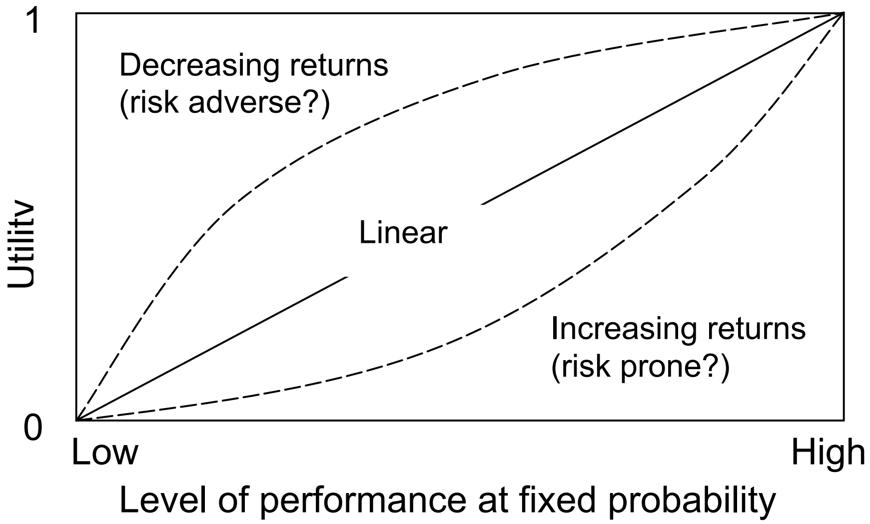

Figure 4-7 SAU versus performance with confusing risk terminology included

For the purposes of the method, as used here, *the above assumption should not be made*. The utilities are intended to express the users preference for various design choices that are assumed to be available in the tradespace. The SAU interviews are expressed in terms of probabilities, but require (and are designed to invoke) a risk-neutral preference for the purely hypothetical choices presented to the user. The issue of actual risks of the product failing to achieve its desired performance (as well as many other sources of risk and uncertainty) will be covered in Chapter 9.

4.2.X-TOS Single Attribute Utilities[10]

The X-TOS project had a well-defined mission (ionospheric science) and customer (the community of scientists developing and improving drag models for LEO satellites). The team carrying out the study had access to a representative of that group who was willing and able to serve as a proxy user. This made the project and ideal case for use of the Lottery Equivalent Probability method. The proxy user, Kevin Ray of the Air Force Research Laboratory at Hanscom AFB, was asked a series of hypothetical questions concerning possible designs for the system, and asked to choose between varying alternatives. The process of repeatedly asking these questions was automated by Seshesai,[8] and the actual interviews were carried out by a web-based software tool.

The interviewing technique of “Lottery Equivalent Probability” centers around constructing plausible scenarios, allowing the user to decide between two alternatives. The challenge in constructing these scenarios is keeping the user focused on the model, instead of a satellite solution they may have in mind. Use of this method requires a sophisticated and patient user, willing to suspend disbelief and go with the somewhat confusing process of the interviews. As the questions are expressed in terms of probabilities of outcomes, it also requires that the user have preferences that are linear with probability. The user needs to be free of both real and psychological reasons to be risk averse or risk taking when answering these questions; see de Neufville, Chapter 18.

Below are scenarios for three of the six X-TOS attributes. The user is asked to choose between the “new” and “old” technology in each case. The questions were asked repeatedly, with different probabilities substituted for the ## placeholders, and different levels of performance substituted for the XX placeholder. When a combination of XX and ## is found for which the user is indifferent to the choice (he or she prefers them equally), the utility of performance level XX is found to be 2*##. The process is expanded in a detailed example first. Then, three of the six questions and resulting single-attribute utility curves are shown below. The others were relatively linear, and so are not shown for brevity. They may be seen in the extract of Ross referenced above.

Detailed Example

One of the attributes of the system was the diversity of latitudes contained in the data set. The acceptable value for this attribute contained the full range of possibilities, from a single latitude (diversity 0) to all latitudes, pole to pole, for a diversity of 180 degrees. For the purpose of the LEP interviews, an alternate method of obtaining diverse data (using a boat) was postulated. This was understood to be fictitious, but the user was willing to suspend disbelief for the purpose of answering the questions.

The questions were posed in the following form:

> A boat-based sensor capable of collecting pertinent data promises to offer a wide diversity of latitudes. However, there is a chance that the boat will never leave port due an ongoing seamen's strike. If you elect to use traditional methods there is a 50% chance that you will get XX degrees of diversity in latitude of your data, or a 50% chance that you will get 0 degrees diversity of latitude in your data. The boat-based sensor offers a ## chance of getting 180 degrees of diversity of your data or a 1-## chance of getting 0 degrees of diversity of your data

In this case, the "traditional method" is a satellite system, for which we are extracting the utilities. Various XX and ## values were substituted, and the user asked to state his preference for either the boat-based sensor or the "traditional" one. The process for choosing the XX and ## values are given in Seshesai.[8]

A single point on the preference curve, that for 30 deg. of diversity, was extracted by asking the above question with the series of XX and ## values in the table below.

Table 1 LEP interview for one point on the diversity of latitude utility curve

Satellite				Boat				
Attribute XX	Probability	Alternate attribute	Probability	Attribute	Probability ##	Alternate attribute	Probability 1-##	User Choice
30	50	0	50	180	45	0	55	Boat
30	50	0	50	180	10	0	90	Sat
30	50	0	50	180	35	0	65	Boat
30	50	0	50	180	20	0	80	Sat
30	50	0	50	180	30	0	70	Boat
30	50	0	50	180	25	0	75	Boat
30	50	0	50	180	22.5	0	77.5	Indifferent*

The table actually shows the full choice offered to the user. At each step above, the user had to pick between a 50% chance of 30 deg. of diversity (the satellite), or ##% chance of 180 degrees of diversity (the boat). In both cases, the alternative was 0 deg. diversity. The first quest ion was easy – a 45% chance of excellent performance vs. a 50% chance of only 30 deg. of diversity. The second was less so, pairing a 10% chance of excellence vs. a 50% chance of 30 deg.; the user chose the latter. The fourth question, with 20% chance of excellence, has also decided in favor of the satellite; all the others were called in favor of the boat. The final question was not actually asked; the method was presumed to have an accuracy of no better than 5%, and the user had already chosen one way for ## = 20% and the other for ## = 25%. Therefore, the user was presumed to be indifferent to a 22.5% chance of excellent performance versus a 50% chance of 30 deg. of diversity.

The determination of the utility was then made, *based on the assumption of linear preference with uncertainty*. The utility of 180 deg. of diversity is one; that of 0 deg., zero. The *linear* expected utility of the boat choice was therefore 22.5% of 1 (0.225). The utility of the satellite choice was, by definition, the same (the user was indifferent to the choice); and the expected value of the satellite choice was 50% of U(30 deg.), so 0.5(U(30 deg.)=0.225. The utility 30 degrees of latitude diversity was therefore determined to be 2(0.225) = 0.45. This gives the first point on Figure 4-8 below.

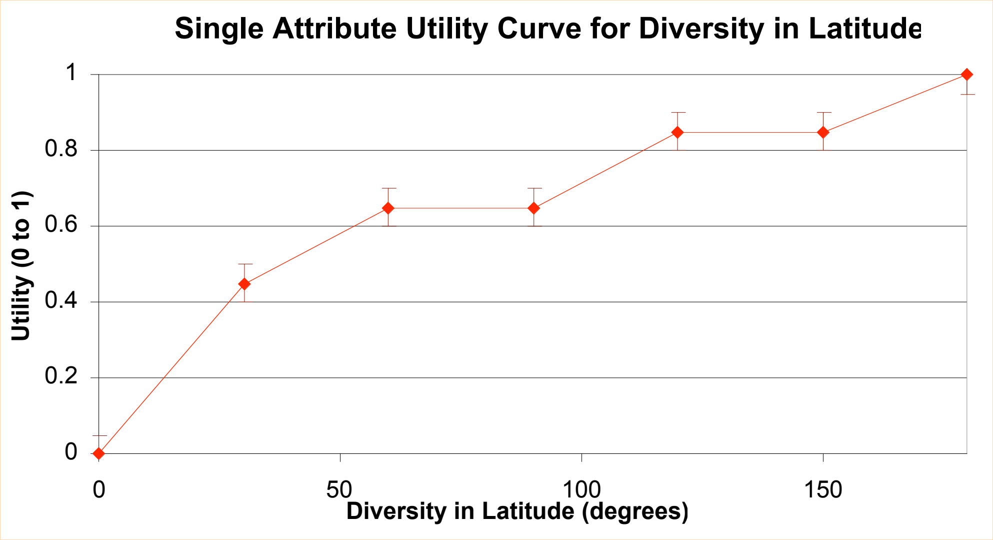

Figure 4-8 Single attribute utility function for Diversity of Latitude

A similar process was used to obtain utility points for all the attributes. In Figure 4-8 we see a very typical case. The utility shows a diminishing return on the diversity of latitude in the data set, with good utility being achieved with modest (60 deg.) diversity.

Time Spent in Equatorial Region

The question was:

> New instruments capable of extracting pertinent data to the AFRL model have been installed on an equatorial ground station. Use of this ground station can get you equatorial data. However, there are many scientific users competing for sole use of these instruments. If you decide not to use this ground station in favor of standard measurement methods, you have a 50% chance of getting XX hours per day of equatorial data or a 50% chance or getting 0 hours per day. Using the new ground station you have a ## chance of getting 24 hours per day or 1-## chance of getting 0 hours per day.

The results, shown in Figure 4-9, are relatively straightforward. The attribute value, bounded on from below by 0 hours per day and from above by 24 hours per day, is monotonically increasing across the range. In this instance, the preference seems to show little sensitivity to the attribute value—the relationship is linear.

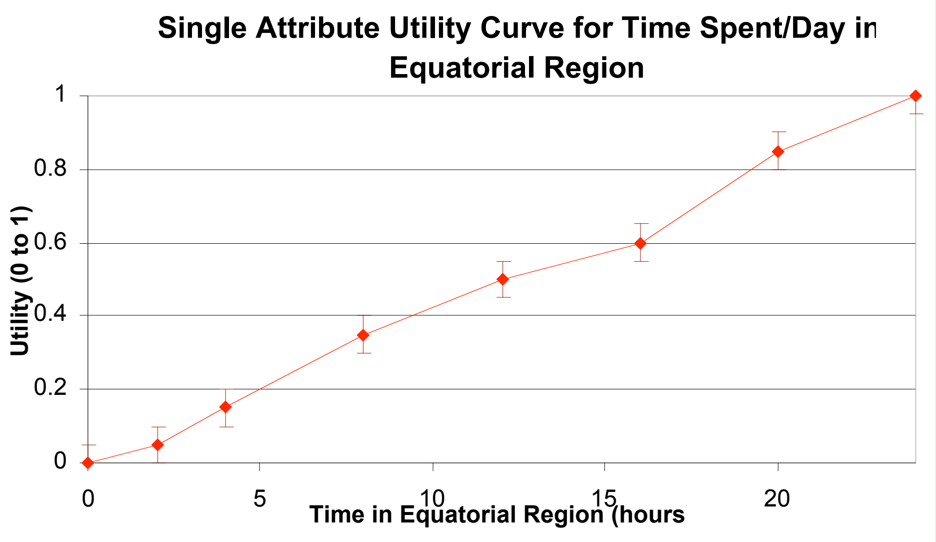

Figure 4-9 Single attribute utility function for Time spent in Equatorial Region

Sample Altitude

The question was:

> A commercial television provider has offered to place a sensor on its geo-synchronous satellite with a lookdown capability to extract pertinent data at 150 kilometers. However, there is a chance that the instrument will become misaligned due to launch vibrations. Your design team has studied the issue and determined that any misalignment will cause the sensor to extract data at 1000 km. You must decide between using this sensor, or traditional methods. The traditional methods will give you a 50% chance of getting data at XX km, or a 50% chance of getting data at 1000 km. The new sensor has a ## chance of extracting data at 150 km or a 1-## chance of extracting data at 1000 km.

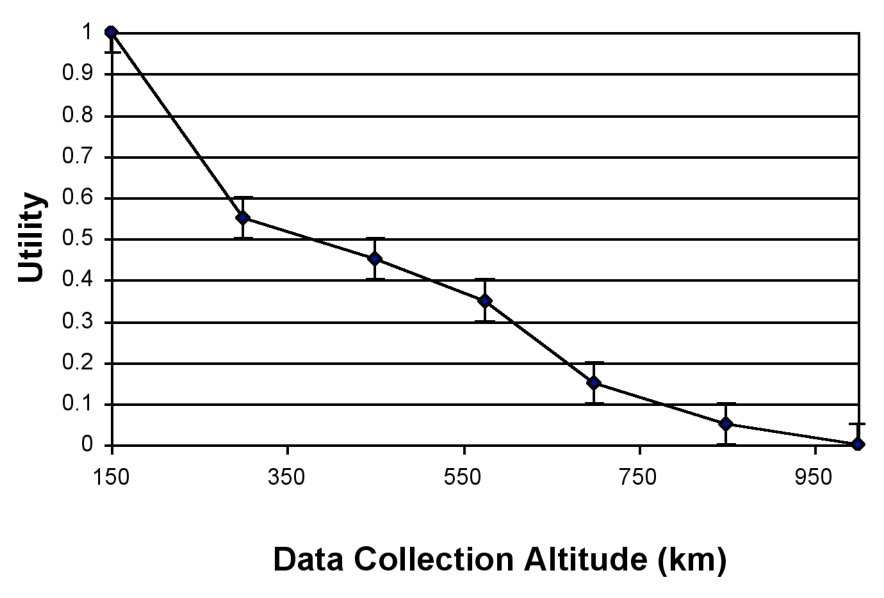

Figure 4-10 Single attribute utility for Data Collection Altitude

This curve shows a strong preference for lower altitude data. The utility drops quickly, and higher altitudes are on a "tail" where the utilities are very low. This is a "threshold" behavior, with the threshold hard against the low end of the range.

4.3.Spacetug Single Attribute Utilities[11]

The Spacetug case presented the opposite situation from the X-TOS project. No users were available for this hypothetical system, and many potential users could be imagined.

In the absence of real users from which to collect more sophisticated functions, it was decided that a simple function that could be explored parametrically was most appropriate. The three attributes are assigned single-attribute utilities. These are dimensionless metrics of user satisfaction from zero (minimal user need satisfied) to one (fully satisfied user). The utilities were chosen *a priori*, with potential users in mind, but without real users available to interview. They utilities were explored as one of the parameters in the tradespace exploration.

The delta-V utility is shown in Figure 4-11. Delta-V is a continuous attribute calculated for each system considered. Utility is assumed to increase linearly with delta-V, with diminishing returns above the levels necessary to do Low Earth Orbit (LEO) to GEO transfers. Variations on this utility are shown in Figure 4-12, which show respectively the utilities of a GEO-centric user (large steps in utility for achieving GEO and GEO round-trip capabilities) and a delta-V-hungry user (continued linear utility for very high delta-V). The manipulator mass (capability) attribute has discrete values, assumed to correspond to increasing utility as shown in Figure 4-13. The response time of a real system would be a complex function of many factors; at the level of the current analysis it is reduced to a binary attribute, valued at one for high impulse systems, and zero for low impulse ones.

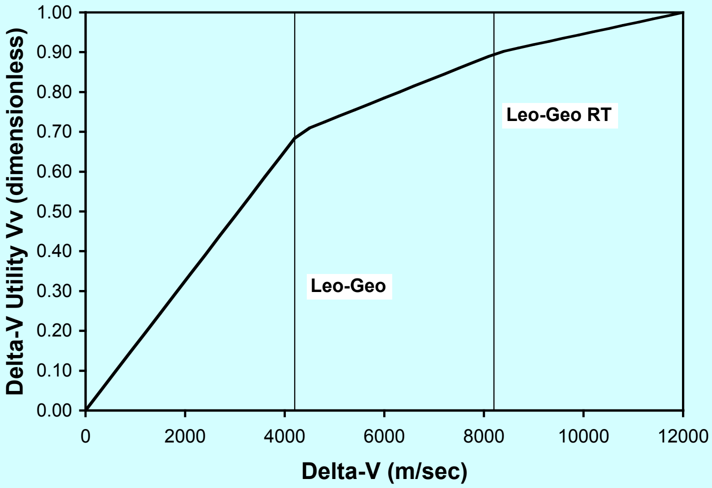

Figure 4-11 Nominal single attribute utility for Delta-V

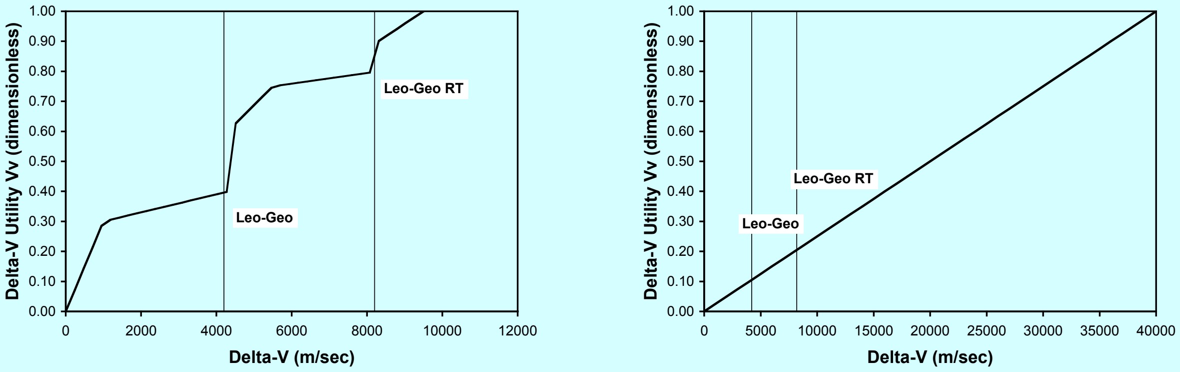

Figure 4-12 Alternate delta-V utilities for GEO-centric user and delta-V hungry user

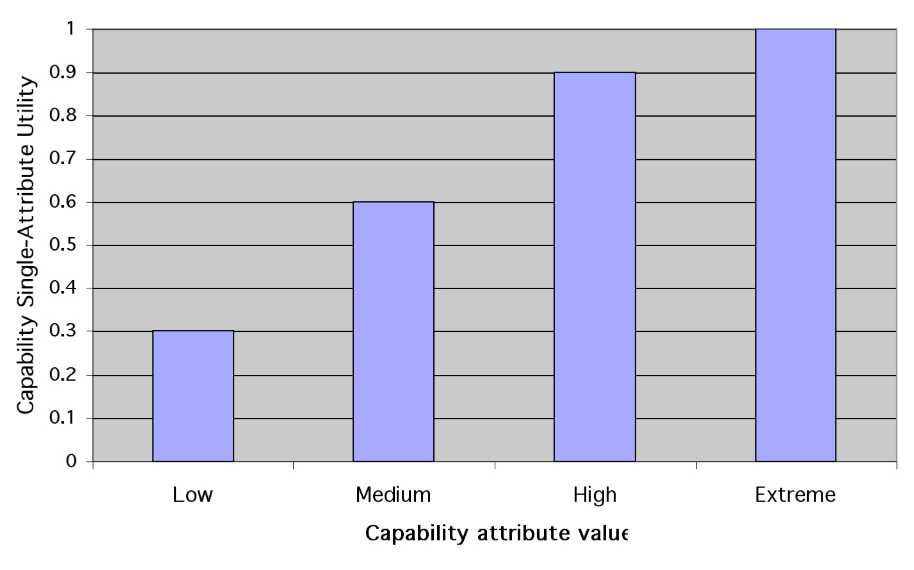

Figure 4-13 Single attribute utility for grappling and observation equipment capability

4.4.Multi-Attribute Utilities

We now have n attributes x_i with known utilities $U_i(x_i)$. The question is now can we combine those utilities in an overall, multi-attribute utility U. In this section, we will explore some of the practical possibilities, while leaving the theory to the key references.

Additive Utility Function (the weighted sum)

The simplest multi-attribute utility function is the weighted sum:

$$U = \sum_{i=1}^{n} k_i U_i \tag{4-4}$$

where k_i is a scalar weight for utility i. It can be found by asking the user the combined utility of a design with attribute i set to its best single-attribute-utility value $x_{i(\max)}$, and all the other attributes set to their worst value $x_{j(\min)}$ for $j \neq i$.

$$k_i = U\left(x_{1(\min)}, x_{2(\min)}, \ldots, x_{i(\max)}, \ldots, x_{n(\min)}\right) \tag{4-5}$$

For U to be a properly normalized utility, these coefficients are under the restriction that

$$\sum_{i=1}^{n} k_i = 1 \tag{4-6}$$

This function is only valid under some very strict limits. The key one is that the single attribute utilities be completely independent of each other in the sense that the utility of one attribute is not affected in any way by the value of other attributes. Formally, an additive multi-attribute function can be used if and only if Eq. 4-4 holds in all cases (see Keeney and Raifa <pg>).[5] This can be true in practice, although it is very difficult to prove rigorously for anything other than very small values of n. Nevertheless, this function is a reasonable choice if there is reason to believe that the single attribute utilities are independent and simplicity and ease of understanding and manipulation are important.

Simple Multiplicative Utility Function

Another simple function is

$$U = \prod_{i=1}^{n} U_i \tag{4-7}$$

This function implies a high degree of interaction between the utilities, of a simple type: the user requires *all* of the single attribute utilities to have a high value for the combined utility to be high. For a combined utility to approach one, all the individual utilities must clearly approach one. Conversely, if *any* of the individual utilities is low (approaches zero) so will the combined utility. This function represents a demanding user who wants all attributes of the system to excel and/or a demanding system function that requires high performance of all of its parts.

Simple Inverse-Multiplicative Utility Function

A final simple function is

$$1 - U = \prod_{i=1}^{n} \left(1 - U_i\right) \tag{4-8}$$

or

$$U = 1 - \prod_{i=1}^{n} \left(1 - U_i\right) \tag{4-9}$$

This function also implies a high degree of interaction between the utilities, of a simple type: the user requires *any* of the single attribute utilities to have a high value for the combined utility to be high. The combined utility approaches one if any of the individual utilities approach one. Conversely, *all* of the individual utilities must be low (approach zero) for the combined utility to do so. This is an “easy to please” user who will be satisfied by excellence in any one area, and/or a system whose attributes are complementary in the sense that good performance in any one area can make up for poorer performance in others.

The Keeney-Raiffa Multiplicative Utility Function

Keeney and Raiffa derive the following form

$$KU + 1 = \prod_{i=1}^{n} \left(Kk_i U_i + 1\right) \tag{4-10}$$

where the k_i are found from Eq. 4-5, but *without* the restriction of Eq. 4-6, and K is the largest non-zero solution to

$$K + 1 = \prod_{i=1}^{n} \left(Kk_i + 1\right) \tag{4-11}$$

This function allows a *single* interaction between the utilities, expressed by the value of K. This interaction can be understood, intuitively if not strictly rigorously, as spanning the continuum of simple interactions covered by the simpler functions above.

If the k_i values collected using Eq 4-5 tend to be high, such that

$$\sum_{i=1}^{n} k_i > 1 \tag{4-12}$$

the user or system is satisfied with partial solutions. In this case, Eq. 4-11 yields K values less than one. In the limit of even some of the k_i's approaching 1, K approaches –1 and the Keeney and Raiffa function reduces to

$$1 - U = \prod_{i=1}^{n} \left(1 - k_i U_i\right) \tag{4-13}$$

which is a modified form of the inverse multiplicative function. It is identical to the inverse multiplicative function if the k_i's are in fact all 1.

If Eq. 4-6 is satisfied (the k_i's sum to one), there is no meaningful non-zero solution for K, so K=0 and Eq 4-10 (after some manipulation) reduces to Eq. 4-4, the weighted sum.

If the If the k_i values collected using Eq 4-5 tend to be low, such that

$$\sum_{i=1}^{n} k_i < 1 \tag{5.11}$$

the user or system is dissatisfied with partial solutions. In this case, Eq. 4-11 yields K values greater than zero, and frequently quite large. In the limit of even some of the k_i's approaching 0, K approaches +∞ and the Keeney and Raiffa function reduces to Eq. 4-7 , the simple multiplicative function.

Requirements for Keeney-Raiffa Multi-Attribute Utility Function

Formally, for the Keeney-Raiffa MAU function to be valid, the single attribute utility functions are under two additional constraints:

- If a user chooses a *pair* of attributes, consisting of attribute x with value x_1 and attribute y with value y_1, over a second pair x_2 and y_2, that choice will not be affected by the value of a third attribute z.
- The single attribute utility function U_i will be altered by no more than a linear transformation by changes in the values of any other attributes.

The first criteria, referred to as "preference independence," requires not only that a user's preference order in one attribute be independent of the values of the other attributes, but that a choice that trades two attributes against each other be independent of the value of third. To understand the later point, imagine that x_1 is a good (high utility) value, while y_1 is poor, and x_2 is poor while y_2 is good. The choice is then about the relative preferences given to attributes x and y; the condition requires that it not be changed by any value of z. This criteria is difficult to check formally (see de Neufville, Chapter 19).[4] Seshesia[8] includes spot checks to assure that it is not grossly violated in the MIST code.

The second criteria, referred to as "utility independence," is slightly stricter in the mathematical sense than the first, but presents no additional difficulties in practice. It simply requires that the single attribute utilities remain utilities in the mathematical sense (see Eq. 4-3) for all values of the other attributes.

Understanding Keeney-Raiffa Functions

The result is a family of well-behaved MAU functions spanning the range from multiplicative to inverse multiplicative, and including weighted sum, differentiated by the parameter K.

Figure 4-14 plots the total utility U against two single-attribute utilities U_1 and U_2, assuming they are moving in the same direction ($U_1 = U_2$). This illustrates the behavior of the multi-attribute functions as the constituents all improve. The full range of Keeney-Raiffa Functions are shown, with k_1 and k_2 values (assumed equal for this example) ranging from almost one to almost zero. The top line shows the inverse multiplicative function. Total utility rises quickly, reaching 0.75 when the two component utilities have reached only 0.5. The straight line in the middle of the plot shows the additive utility function, with total utility rising in proportion to the constituent utilities. The lowest line shows the multiplicative utility function. Total utility rises slowly at first, reaching a value of only .25 when the two component utilities have reached 0.5. It rises steeply as both constituent utilities approach one. The continuous range of functions between the extremes, as the values of k_1 and K vary, can be seen.

Figure 4-15 plots the total utility U against two single-attribute utilities U_1 and U_2, assuming they are moving in opposite directions ($U_1 = 1 - U_2$). This illustrates the behavior of the multi-attribute utility functions as one attribute is traded off for the other. The same family of functions are shown. The inverse multiplicative function is one if either U_1 or U_2 is one, and has its lowest value (0.75) when they are both 0.5. The weighted sum is a straight line, as the tradeoff between the equally weighted constituents is a wash. The multiplicative function is zero if either U_1 or U_2 is zero, and reaches its *maximum* value (of only 0.25) in the center. Again, a full range of functions between the extremes is available.

MAU functions that involve more than two attributes have the same sort of behavior shown in Figure 4-14 and Figure 4-15, only with more spread between the functions.

The point here is not abstract; it is important that the nature of the trades between the single attributes be understood as part of the tradespace evaluation. The Spacetug study used a weighted sum for simplicity. The X-TOS study used a Keeney-Raiffa function, but the coefficients resulted in a K of 0.07; examination of Figure 4-14 and Figure 4-15 shows this is indistinguishable from a weighted sum. This implies that the X-TOS attributes were fully independent, and can be traded simply by treating the k_i as weights. Interestingly, when the k_i's were changed quite drastically late in the study, K changed to 0.28, implying the MUA was still very close to a weighted sum.

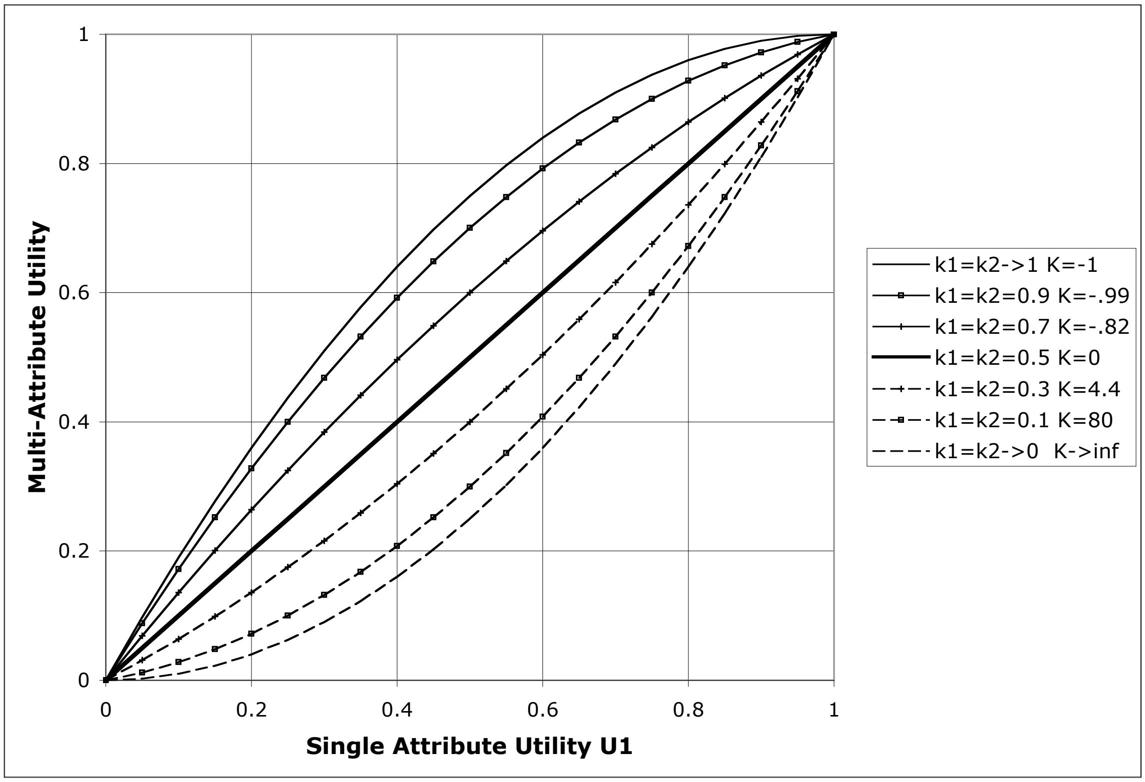

Figure 4-14 Family of valid MAU functions for two attributes moving in the same direction (*U_2=U_1*)

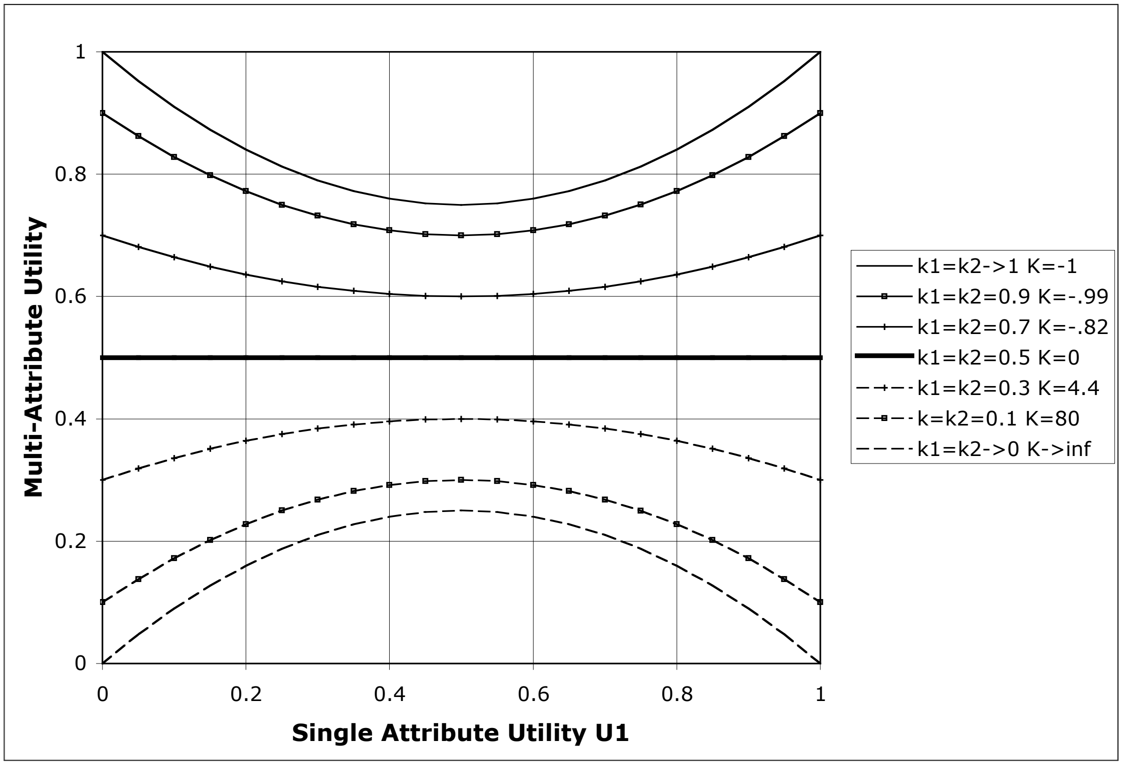

Figure 4-15 Family of MAU functions for two attributes moving in opposite directions (*U_2=1-U_1*)

In contrast to these examples, the B-TOS study[7] used a Keeney-Raiffa function which had a K of –0.998. The MAU was clearly a modified inverse multiplicative function. This greatly complicated the interpretation of the multi-attribute utility. The complication was furthered by the fact that one of the single-attribute utilities had a k very close to 1 and a utility of 1 for all competitive systems. The result was that all competitive systems had U varying from 0.99 to 0.99999. This did not invalidate the MAU function for ranking of the alternatives, but made it difficult to interpret, and impossible to present, the absolute values of the U's.

Alternate Methods

The MAU technique is recommended for typical tradespace analysis situations. It may not be the most appropriate in all situations. On one hand, if the user needs are very well understood and can be reduced to one commodity good, a quantitative measure of this commodity would be more useful than the dimensionless and more abstract utility. In the other extreme, user needs may be obscure and/or may not conform to the restrictions (such as independence of single attribute preferences) required by MAU theory. In these cases, the tradespace may need to be explored with parametric MAU's, or explored in the absence of a single utility function.

The simplest situation is when the performance of a system can be reduced to a single desired quantity that accurately expresses what the users want. In analyzing broadband communication systems, Gumbert *et al*.[12] proposed product of such systems was a commodity – billable minutes of communication time. Thus, the sole measure of value is the cost per billable minute. Shaw et al[13] used this metric to do an early tradespace-type analysis. The generalization of this idea is to measure Cost Per Function (CPF). To have an ascending measure of value, one could simply invert CPF and measure Function Per Cost.

A single, quantifiable metric such as CPF makes comparisons between a broad range of different systems easy. This approach is limited, however, to situations where the function desired can be expressed simply and quantitatively as a single value. This tends to happen when the function is well understood and can be thought of in commodity terms (e.g. communications services). It presupposes (possibly incorrectly) that other aspects of the service (e.g. reliability, timeliness, etc) are either not important, or can be incorporated into the CPF in some way. It also presupposes that the market's demand for the commodity is understood.

Slightly more complex is a constructed metric that looks like a single performance metric. An analysis of the Terrestrial Planet Finder (TPF) system uses a CPF metric, cost/image, but the image is actually and average of a number of different types of images, with a fixed (assumed!) mix, and no measure of the relative value of different image types.[14] Again, the advantage is ease of comparison; the clear weakness is a set of assumptions needed to be made to allow an "average" function to be calculated. These assumptions may bury a host of issues.

In the other extreme, the user's desires may be unknown, and/or not correspond to the requirement of MAU theory. If the user is unknown, or very uncertain of his or her needs, the MAU can be defined and explored parametrically. This was done in the SPACETUG case, by postulating a series of potential users, differentiated by the weights they placed on the SAU's, and in some cases by the shapes of the SAU's. In this case the MAU theory is not used as a formalism for capturing user preferences, but rather as a way to organize thoughts about the preferences of potential users.

If the user has preferences, but they are not independent, this does not invalidate the tradespace exploration concept, but it does require some creativity for determining user preferences across many possibilities without requiring the user to rank preferences exhaustively across thousands of choices. A very interesting approach is presented by Belegundu *et al.*[15] A "design by shopping" paradigm is proposed, where the user is presented with many possible designs, represented as a multi-dimensional data set, projected using advanced tools. In our context, the data set could consist of the single attribute utilities; the user could decide where in the tradespace of possible single attribute utilities their best overall utility could be found without creating a reduced MAU. As we will see in the next chapter, Tradespace Exploration, it is wise to explore this tradespace even if an MAU is created. The MAU provides some guidance into which designs may provide the best user utility, which provides guidance for more detailed explorations.

Multiple Stakeholders

If the wishes of multiple stakeholders clash, there is no formalism for finding a "best" solution. Indeed, it has been demonstrated that there is no simple optimal solution to this class of problems.[16,17] Scott and Antonsson discuss reasons that this may not be as severe a problem in engineering problems as in social ones.[18] De Neufville has a relatively light coverage of what to do in this situation in his Chapter 21.

Ultimately, conflicting stakeholders need to negotiate solutions. If every stakeholder has an understanding of the tradespace for their own utilities they enter the negotiation with more useful knowledge. Understanding which trades are truly unavoidable (e.g. the trades on the Pareto front) as opposed to design changes that can be made to please all parties focuses negotiation on the real issues. Reevaluating the tradespace for a variety of stakeholder utilities (as was done in the Spacetug example) is straightforward and may be a major contributor to negotiations between stakeholders looking for a mutually acceptable solution. Alternately, in a competitive situation, tradespace knowledge may be a major advantage. In either case, tradespace understanding is a powerful tool for moving the human process of multi-stakeholder decision-making forward.

4.5.X-TOS Multi-attribute Utilities

The X-TOS study used a Keeney-Raiffa MUA function. The coefficients k_i were determined using the relation in Eq. 4-5. They were collected by the MIST software at the same time as the single attribute utilities. They are shown in the figure below.

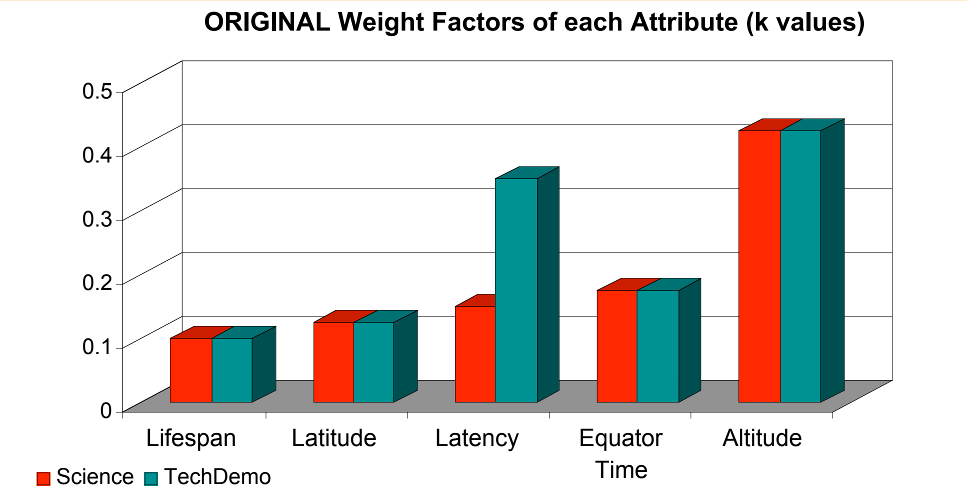

Figure 4-16 Original k values for the X-TOS MAU function

This resulted in a K value of 0.07 for the Science mission. After reviewing the data from the first interview, Kevin Ray realized that he had not put enough emphasis on the spacecraft lifetime. The ability to capture many atmospheric cycles (such as day/night, monthly, yearly, and solar cycles) is actually quite important for a successful mission. This retrospection changed the weight factor of the data lifespan attribute from 0.1 (lowest) to 0.3 (second highest). In a similar manner, Mr. Ray realized that there was very little importance on latency for a science mission. He reduced this weight factor from 0.15 (3rd highest) to 0.1 (lowest). Lastly, Mr. Ray altered the shape of the data lifespan utility curve, which resulted in a somewhat linear relationship between utility and data lifespan. The utility of a 2-year mission was decreased from 0.35 to 0.3, and the utility of a 4-year mission was increased from 0.35 to 0.44. The resulting K value was 0.28.

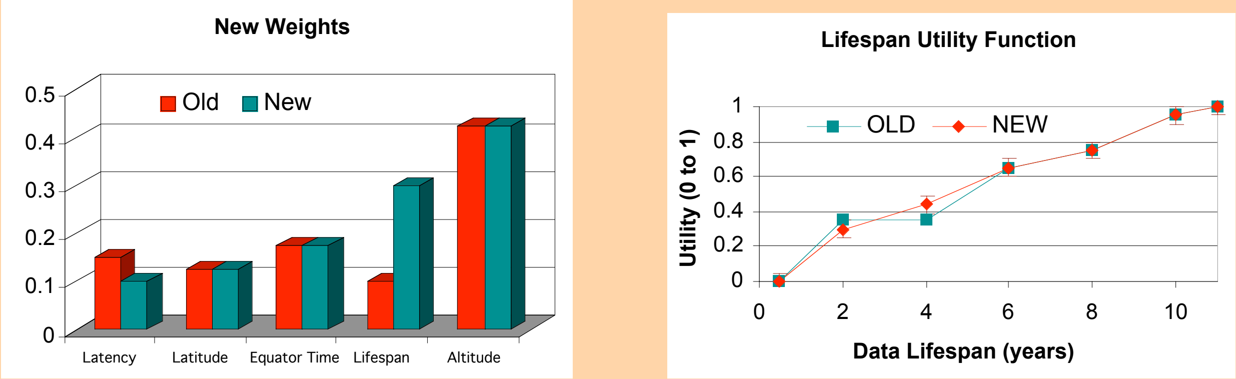

Figure 4-17: New weights and Lifespan utility function

4.6.Spacetug Multi-Attribute Utilities

The space tug study used weights determined for a number of hypothetical users. These are summarized in Table 2.

Table 2 Utility weightings for various hypothetical Spacetug Users

Attribute	Nominal Weights	Capability Stressed User	Response Time Stressed User
Delta-V	0.6	0.3	0.2
Capability	0.3	0.6	0.2
Response Time	0.1	0.1	0.6

These were used in a simple weighted sum MAU. The tradespace was evaluated separately for each user. Note that there were two more users considered, the GEO-centric and Delta-V hungry users, who used the nominal weights but had modified single-attribute utilities for delta-V (see Figure 4-12 above). The tradespace was evaluated for the needs of all five potential users.

4.7.Concluding Thoughts

MAU is a formal method for extracting user preferences which is very useful for tradespace exploration. However, it must be remembered that the point is not to perform a perfect extraction of the user's needs, which may be very imperfectly formed in any case. The point is to gain insight into, and quantify (even if imperfectly), the users preferences, in order to proceed with the exploration of the tradespace. No choices are fixed, correctly or incorrectly, at this stage. We will close with two quotes. From Otto and Antonsson,[19] a defense against critics that might point out the difficulty of proving that a utility approach is mathematically rigorous:

> The defense against this criticism is that one must give the designer credit for some intelligence. The designer must determine what is optimal. Any method should allow designers to iteratively determine this by choosing which trade-off strategy appears most appropriate, and to allow the designer to modify any and all of these initial choices. ... Design problems *are* commonly solved in an iterative manner, not usually with a single formalization and subsequent optimization. In an iterative design process, a designer makes determinations without complete understanding, thus enabling the designer to (ultimately) form a more complete understanding.

From Thurston,[20] some thoughts on the correct use of MAU:

> Utility analysis cannot be the only analytic tool employed in design. It cannot contribute much to the creative or configuration phase, except to free the designer to think in terms of function rather than form. It cannot tell the designer which raw material options are available, nor the beam cross-section required to bear a particular load. Neither can it fully resolve the problem of de- fining the optimal group decision, one which has long plagued economists. Like many useful analytic tools, it can be used naively or incorrectly, and there are special cases that yield inconsistent or nonsense results.
>
> However, design theory and methodology is an arena worthy of endeavor because traditional design processes sometimes take too long, result in products that are too costly, are difficult to manufacture, are of poor quality, don't satisfy customer needs, impact the environment adversely, and provide design teams with only ad hoc methods for communicating and resolving conflicting preferences. Utility analysis can help remedy these problems by quickly focusing creative and analytic efforts on decisions that affect important design functions, by identifying the best tradeoffs—particularly under uncertainty!—and by disaggregating design team decision problems into subproblems on which consensus can be reached. So, while decision theory by itself does not constitute a theory of design, integrating it throughout several design phases, including configuration and analysis, can improve both the process and the product of design.

4.8. Problems

1. de Neufville problem 18.4.

2. de Neufville problem 20.6

3. A user has the following single attribute utility functions for attributes A, B and C:

Attribute A	Utility		Attribute B	Utility		Attribute C	Utility
50	0.0		1	0.0		5	1.0
100	0.2		2	0.3		7	0.9
150	0.4		3	0.6		8	0.7
200	0.6		5	0.8		9	0.3
250	0.8		7	0.9		12	0.1
300	1.0		10	1.0		15	0.0

The user states a design with attributes (300,1,15) would have multiattribute utility 0.9, a design with attributes (50,10,15) would have utility 0.5, and a design with attributes (50,1,5) would have utility 0.3. The user is presented with the following design choices:

Design	Attribute A	Attribute B	Attribute C
Alpha	150	10	10
Beta	300	2	7
Gamma	275	7	9
Delta	225	5	7.5
Epsilon	100	10	5
Phi	150	3	8

Rank these designs using:

a) Lexographic ordering
b) A weighted sum, “normalized” (by multiplying by a constant factor) so that the multiattribute utility scales from 0 to 1
c) A multiplicative function
d) A modified inverse multiplicative function
e) a Keeney-Raiffa function

Comment. What ranking is most likely closest to the user’s true desires?

NOTES AND REFERENCES

---

[1] Diller, N. P., “Utilizing Multiple Attribute Tradespace Exploration with Concurrent Design for Creating Aerospace Systems Requirements,” Master of Science Thesis in Aeronautics and Astronautics, Massachusetts Institute of Technology, June 2002.

[2] Ross, A. M., “Multi-Attribute Tradespace Exploration with Concurrent Design as a Value-Centric Framework for Space System Architecture and Design,” Master of Science Thesis in Aeronautics and Astronautics, Massachusetts Institute of Technology, June 2003.

[3] Most of the time, the utility of the key stakeholders, or d*ecision makers*, is what we desire. See Chapter 3, section 4.3, for further discussion of this issue.

[4] de Neufville, Richard, *Applied Systems Analysis: Engineering Planning and Technology Management*, McGraw-Hill, New York, 1990.

[5] Keeney, Ralph L., and Raiffa, Howard, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs,* Cambridge University Press, Cambridge, UK, 1993.

[6] Text of this paragraph taken verbatim from Hugh L. McManus, Daniel E. Hastings, and Joyce M. Warmkessel, “New Methods for Rapid Architecture Selection and Conceptual Design,” *Journal of Spacecraft and Rockets,* January 2004.

[7] “B-TOS Architecture Study,” 16.89 Space Systems Engineering Final Report, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, May 2001.

[8] Seshasai, Satwiksai, “Knowledge Based Approach to Facilitate Engineering Design,” Masters Thesis in Electrical Engineering, Massachusetts Institute of Technology, May 2002.

[9] Spaulding, Timothy J., “Tools for Evolutionary Acquisition: A Study of Multi-Attributes Tradespace Exploartion (MATE) Applied to the Space Based Radar (SBR), Master of Science Thesis in Aeronautics and Astronautics, Massachusetts Institute of Technology, June 2003, see page 42 and Appendix A.

[10] Most text and graphics taken from “X-TOS: 16.89 Space Systems Engineering Final Design Report,” Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, May 2002.

[11] Some text modified from McManus, H. L. and Schuman, T. E., “Understanding the Orbital Transfer Vehicle Trade Space,” AIAA Paper 2003-6370, Sept. 2003.

[12] Gumbert, C. C., Violet, M. D., Hastings, D. E., Hollister, W. M., and Lovel, R. R., “Cost per Billable Minute Metric for Comparing Satellite Systems, Journal of Spacecraft and Rockets, Vol. 34, No. 6, 1997, pp. 837-846.

[13] Shaw, G. M., Miller, D. W., and Hastings, D. E., “Development of the Quantitative Generalized Information Network Analysis (GINA) Methodology for Satellite Systems,” *Journal of Spacecraft and Rocket*s, Vol. 38, No. 2, 2001, pp. 257-269.

[14] Jilla thesis and paper.

---

[15] Belegundu, A. D., Halberg, E., Yukish, M. A., and Simpson T. W., “Attribute-Based Multidisciplinary Optimization of Undersea Vehicles,” AIAA Paper 2000-4865, 8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Long Beach, CA.

[16] Arrow, K. J., *Social Choice and Individual Values*, 1st ed., J. Wiley, New York NY, 1951.

[17] Arrow, K. J., and Raynaud, H., *Social Choice and Multicriterion Decision-Making*, The MIT Press, Cambridge, MA,1986.

[18] Scott, M. J. and Antonsson, E. K., "Arrow's Theorem and Engineering Design Decision Making," Research in Engineering Design, Vol. 11, No. 4, pp. 218-28.

[19] Otto, K. N., and Antonsson, E. K., “The Method of Imprecision Compared to Utility Theory for Design Selection Problems.” <full ref?>

[20] Thurston, D. L., "Real and Misconceived Limitations to Decision Based Design with Utility Analysis," Journal of Mechanical Design, Vol. 123, June 2001, pp. 176-82.