

MULTIPLE OBJECTIVE DECISION MAKING

Kent D. Wall and Cameron A. MacKenzie *

Defense Resources Management Institute, Naval Postgraduate School

April 20, 2015

1 INTRODUCTION

Imagine the the following decision problem for the fictitious country of Drmecia, a small nation with many security problems and limited funds for addressing these problems:

The Army of Drmecia operates a ground based air-search early warning radar base located near the capital city, Sloat. The radar continues to exhibit low availability because of reliability and maintainability problems. This exposes the country to a surprise air attack from its enemy, Madland. The Army obtained procurement funds two years ago and installed another radar of the same type as a back-up, or stand-by, radar. Unfortunately, the second radar unit also has exhibited low availability and too often *both* radars are off-line. The commander of Sloat Radar Base is now requesting funds for a third radar, but the Army Chief of Staff is very concerned. He knows that the defense budget is under increasing pressure and funds spent on another radar may have to come from some other Army project. In addition, two new radar models are now available for procurement at higher cost than the existing radars. It may be better to purchase one of these than another radar of the existing model.

The Chief of Staff knows he must have a strong case if he is to go to the Minister of Defense and ask for additional funds. He wants to know which radar is best. He wants to know what he is getting for the extra money spent. He wants to know what is the most cost-effective course of action.

The director of analysis for the Army of Drmecia knows that making a cost-effective decision requires two things: (1) a cost analysis and (2) an effectiveness analysis. The analysis of cost is a familiar problem and he has a staff working on it. The effectiveness part of the problem, however, is a concern. What is effectiveness in this situation? How can effectiveness be defined and quantified? How can the cost analysis be integrated with effectiveness analysis so that a cost-effective solution can be found?

The goal of this chapter is to develop a method to think about and quantify effectiveness in the public sector, specifically defense. Effectiveness can best be measured in the public sector by developing a framework for solving decision problems with multiple objectives. The framework will provide you with a *practical tool for quantitative investigation* of all factors that may influence a decision, and you will be able to determine *why* one alternative is more effective than others. This analytical ability is very important because many real-life decision problems involve more than a single issue of concern. This holds true for personal-life decisions, private sector business decisions, and public sector government resource allocation decisions.

Examples of personal-life decision problems with multiple objectives are plentiful: selecting a new automobile; choosing from among several employment offers; or deciding between surgery or medication to correct a serious medical problem. In the private sector, maximizing profit is often the sole objective for a business, but other objectives may be considered, such as maximizing market share, maximizing share price performance, and minimizing environmental damage.

Public sector decision making almost always involves multiple objectives. State and local government budget decisions are evaluated, at least, in terms of their impacts on education programs, transportation infrastructure,

*This is an author's manuscript of a book chapter published in *Military Cost-Benefit Analysis: Theory and Practice*, eds F. Melese, A. Richter, and B. Solomon, New York: Routledge, 2015, pp. 197-236.

public safety, and social welfare. The situation is more complex at the Federal level because national defense issues enter the picture. For example, consider the ways in which the U. S. Department of Defense evaluates budget proposals. Top-level decision makers consider the effects of a budget proposal on (1) the existing force structure, (2) the speed of modernization, (3) the state of readiness, and (4) the level of risk, among other factors. Other national defense objectives include the ability to deter aggression, project force around the world, and defeat the enemy in combat no matter the location and his military capability. This multidimensional character permeates all levels of decision making within a Planning, Programming, Budgeting and Execution System (PPBES). Acquisition decisions, training and doctrine policy changes, and base reorganization and consolidation choices all have multiple objectives.

Government decisions in general, and defense resource allocation decisions in particular, have an added evaluation challenge. Outcomes are difficult, if not impossible, to represent in monetary terms. First, benefit cannot be expressed in terms of profit. Unlike the private sector, the public sector is not profit motivated and this single monetary measure of benefit is not relevant. Second, market mechanisms often do not exist for “pricing out” the many benefits derived from public sector decisions. Thus, it is not possible to convert all the benefits into monetary terms and conduct a cost-benefit analysis. In national defense, benefits are often characterized in terms like deterrence, enhanced security, and increased combat capability. No markets exist that generate a price per unit of deterrence or a unit increase in national military security.

Solving these decision problems requires a structured systematic approach that aids in discovering all the relevant objectives and makes it easy to work with objectives expressed in many different units of measure. It also must allow the decision maker to account for the relative importance of the objectives and the importance of marginal changes within individual objective functions. Following such an approach can allow decision makers to perform cost-effectiveness analysis to compare among different programs and alternatives.

2 PROBLEM FORMULATION

A systematic approach to decision making requires developing a model that reflects a decision maker’s preferences and objectives. We assume the decision maker is rational and seeks the most attractive or desirable alternative. Goals describe what the decision maker is trying to achieve, and objectives determine what should be done in order to achieve the decision maker’s goal(s). A model based on the decision maker’s objectives will allow the decision maker to evaluate and compare among different alternatives.

For a decision problem with a single objective, the model for decision making is straight forward. For example, consider the simple problem of choosing the “best” alternative where “best” is interpreted as the greatest range. Not surprisingly, the decision maker should choose the alternative with the greatest range.

Things are never this straight forward when there are multiple objectives because the most desirable or most effective alternative is not obvious. For example, consider the selection of the Sloat Radar as discussed in the introduction. The decision maker knows that range and interoperability are important. Maintenance and reliability are also concerns. How do we define effectiveness in this situation? Do we focus only on range? Do we focus only on interoperability? How can we consider all four objectives at the same time?

2.1 Issues of Concern and Objectives

The issues of concern to the decision maker can always be expressed as objectives. For example, in selecting a radar system, suppose two alternatives are equal in all respects except for range. In this case, the decision maker chooses the radar with more range. This concern with range translates to an objective: the decision maker wants to maximize range. Likewise, suppose that two alternatives are equal in all respects except for required maintenance, and the decision maker selects the alternative with less required maintenance. The concern with required maintenance translates to an objective: the decision maker will want to minimize required maintenance.

The rational decision maker can always be shown to act in a way consistent with a suitably defined set of objectives. Hence, decision problems characterized by many issues of concern are decision problems in which the decision maker attempts to pursue many objectives. In other words, the decision maker confronts a *Multiple Objective Decision Problem*.

The set of objectives is fundamental to the formulation of this type of decision problem. If we have a complete set of objectives that derive from *all* the issues of concern to the decision maker, we have a well formulated problem.

Furthermore, if these objectives are defined in sufficient detail, they tell us how to evaluate the alternatives in terms that have meaning to the decision maker. In the radar example, each alternative can be evaluated in terms of its range, its interoperability, its required maintenance, etc. This knowledge is fundamental to the solution of this type of decision problem. This knowledge tells us how to quantify things or how to measure the attainment of the objectives. We seek to represent the alternatives by a set of numbers that reveal to the decision maker how well each alternative “measures up” in terms of the objectives.

Let us introduce some mathematical notation to help with the formulation, and Table 1 defines all the variables used in this chapter. Let there be M objectives for the decision maker. Let these be defined in enough detail so that we know how to measure the attainment of each objective. Let these measures be denoted x_i where $1 \leq i \leq M$. For example, $x_1 = \text{range}$, $x_2 = \text{interoperability}$, $x_3 = \text{required maintenance}$, etc. Let there be N alternatives indexed by the letter j where $1 \leq j \leq N$. Thus, each alternative can be represented by a set of M numbers:

$$\{x_1(j), x_2(j), x_3(j), \dots, x_M(j)\}.$$

These numbers become the *attributes* for the j^{th} alternative, and each of the M numbers represent how well the j^{th} alternative meets each one of the decision maker’s M objectives. Evaluating an alternative means determining the numerical value of each of the measures. These serve as the raw data upon which the decision is made.

The key to solving the decision problem is to explain how this set of M numbers is viewed in the mind of the decision maker. This is done by constructing a *collective measure of value* that reflects how much the decision maker *prefers* one collection of attributes vis a vis another collection of attributes. For example, how does the decision maker value the set of attributes for the first alternative, $\{x_i(1)\}$, compared to the set of attributes for the second alternative, $\{x_i(2)\}$, where $1 \leq i \leq M$? Is the the set $\{x_i(1)\}$ more “attractive” than the set $\{x_i(2)\}$?

It is challenging for a decision maker to compare alternatives if each alternative has several attributes (i.e., if M is greater than 3 or 4). Developing a function that combines the M attributes into a single number provides a method for the decision maker to compare alternatives based on the attributes that describe each alternative. This single number is called the *Measure of Effectiveness (MOE)*.

2.2 Effectiveness

The function that measures the effectiveness of the j^{th} alternative is called a value function, $v(j)$. This function incorporates the *preferences* of the decision maker and converts each *collection of attributes* as represented by the set $\{x_i(j)\}$ into a number that represents the attractiveness or desirability of the collection in the mind of the decision maker. It measures the extent to which the j^{th} alternative helps the decision maker pursue all the objectives while taking into account the relative importance of each. Effectiveness is a number. It quantifies “how far we go” towards achieving our goals as measured by the value of the objectives. The objectives are not all of equal importance, however, and $v(j)$ also takes into account the relative importance of each objective.

The decision maker desires to maximize effectiveness by choosing the alternative with the highest $v(j)$ subject of course to cost considerations. The fundamental problem in formulating the decision problem is defining $v(j)$ based on the set of attributes $\{x_1(j), x_2(j), x_3(j), \dots, x_M(j)\}$ and then integrating cost with the MOE.

Formulating a decision problem with multiple objectives requires four pieces of information:

1. We need a list of the relevant objectives.
2. We need to know how to value the measures associated with each individual objective.
3. We need to know the relative importance of these objectives.
4. We need to know the relative importance between the cost and effectiveness of each alternative.

We proceed to address each of these needs in the rest of the chapter.

3 DISCOVERING THE RELEVANT OBJECTIVES

We must know what matters in a decision problem or the consequences that a decision maker considers when thinking about each solution alternative. These are the issues of concern and are represented by objectives. We cannot judge alternatives without knowing the objectives of the decision maker.

Table 1: Variable definitions

Variable	Definition
a	Assessed parameter in the exponential value function
b_1	Assessed parameter for squared term in the exponential value function
b_2	Assessed parameter for cubic term in the exponential value function
$c(j)$	Life-cycle cost of the j^{th} alternative
c_{ideal}	Dollar amount at which the value of cost equals 1
$c_{\text{too much}}$	Dollar amount at which cost is too high or value of cost equals 0
i	A single objective or attribute
j	A single alternative
$k = \frac{1}{c_{\text{too much}}}$	Reciprocal of too high of cost
w_i	Importance weight for the i^{th} attribute
w_A	Importance weight for availability
w_C	Importance weight for complexity
w_E	Importance weight for electronic counter-counter measures (ECCM)
w_I	Importance weight for interoperability
w_L	Importance weight for cognitive load
w_P	Importance weight for performance
w_R	Importance weight for range
w_U	Importance weight for ease-of-use
$v(j)$	Measure of effectiveness for the j^{th} alternative
$v_c(c(j))$	Value function for cost of the j^{th} alternative
$v_i(x_i)$	Value function for the i^{th} attribute
x_i	Measurement for the i^{th} attribute
$x_i(j)$	Measurement for the i^{th} attribute for the j^{th} alternative
x_{ideal}	Ideal measurement for an attribute
x_{max}	Maximum measurement for an attribute
x_{min}	Minimum measurement for an attribute
$x_{\text{too little}}$	Too little for an attribute for more-is-better case
$x_{\text{too much}}$	Too much for an attribute for less-is-better case
z	Swing in attribute that contributes least to an objective
z_i	Difference between x_i and either x_{max} or x_{min}
K	Normalization constant for exponential value function
K_1	Normalization constant for quadratic exponential value function
K_2	Normalization constant for cubic exponential value function
M	Number of objectives or attributes
N	Number of alternatives
$V(j)$	Payoff function for cost-effectiveness of the j^{th} alternative
V^* and V^{**}	Payoff values of cost-effectiveness
W_C	Importance weight for cost
W_E	Importance weight for effectiveness

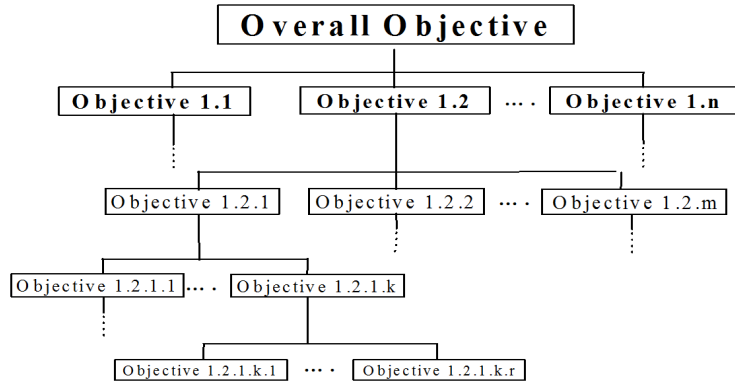


Figure 1: Generic hierarchy

Discovering all the relevant objectives is the first step in solving the decision problem. It is helpful to employ a graphical construction called an hierarchy or “tree structure.” (Think of a tree: a trunk with a few main branches that have many more smaller branches that have even more smaller branches—now turn that picture upside down and you have a picture of an hierarchy.) Figure 1 depicts a hierarchy with several levels. The hierarchy begins at the top most level with a single over-arching objective that captures in its definition all that the decision maker is trying to do. For multiple objective problems in public policy, the overall objective is to *maximize effectiveness*. The objectives in the first level below the overall objective tell us how we define the top level hierarchy. We say the overall objective is refined or defined in more detail by the objectives listed on the next level down. The lower level provides more detail as to what is meant by the objective at the next higher level.

The definition is not operational unless it is useful in measuring things. We must have a way to develop the hierarchy in enough detail so that we get specific enough in our definition of the objectives that it becomes obvious to us (and everyone else) how “to measure things.” What we need is a method of construction that takes us from the top level down to the lowest level where measurement is obvious. There are several ways of doing this but we will only discuss the “top-down” method and the “bottom-up” method.

3.1 The Top-Down Approach

In the top-down approach you start with the obvious, to maximize effectiveness. This becomes the top-level objective. Because this objective can mean many things to many people and is too vague to be operational, the next step is to seek more detail. You proceed by asking the question: “What do you mean by that?” The answer to this question will allow you to write down a set of sub-objectives, each of which derives from a more detailed interpretation by the decision maker of just what overall effectiveness means to him or her. For example, in the case of Sloat Radar the decision maker may say: “Maximizing availability is part of maximizing overall effectiveness, I’ve got to have a radar that works almost all the time.” He may also say that “I need a high performance radar that works almost all the time, so maximizing radar performance is also part of maximizing overall effectiveness.” Finally, he may say that he needs a radar that is easy to use. Thus minimizing radar complexity for the user is also important.

The top-down approach continues in this way until there is no doubt what the objectives mean because we will be able to measure their value for each alternative. For example, in the first level down from the top we know exactly what we mean by acting so as “to maximize availability.” Availability has a well known precise definition: availability is the probability that a system will work at any given time. It can be measured either directly or by computation using the system’s Mean Time Between Failure (MTBF) and its Mean Time To Repair (MTTR). Each alternative can be evaluated in terms of its availability, and maximizing availability is pursued by seeking the system that exhibits the highest availability.

Performance, however, does not have a precise agreed upon definition that tells us how we can measure it. Here we must ask once again: “What do you mean by that?” The answer to this question will allow us to understand what is meant by maximizing performance. Suppose, for example, that the response to this question is: “High performance means great effective radar range, resistance to electronic counter measures, and high interoperabil-

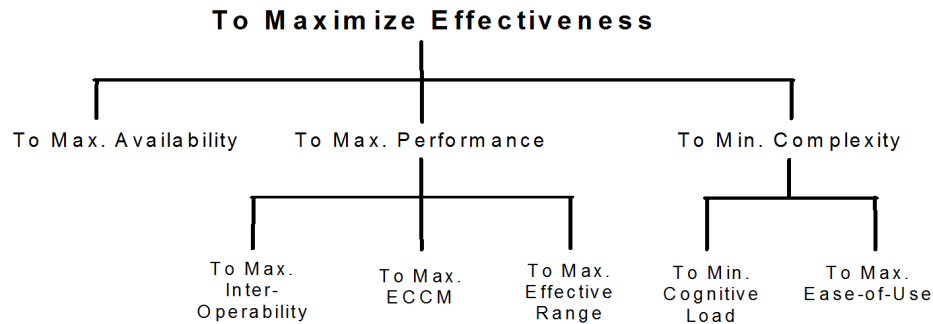


Figure 2: Sloat Radar objectives hierarchy

ity.” We do not need to ask any other questions here because we know how to measure (evaluate) range for each alternative. We know how to evaluate resistance to electronic counter measures: a simple “yes” or “no” answer will do. Either an alternative has electronic counter-counter measures (ECCM) capability or it does not. Finally, we know how to measure interoperability: we can count the number of communication links that can be operated by each alternative (so it can feed target information to the various anti-air forces).

Minimizing user complexity also requires refinement so we need to ask: “What do you mean by that?” The answer in our example may be that to minimize user complexity we need to minimize the cognitive load placed on the radar operator (measured by how many things the operator has to watch, sense, react to, adjust, and refine) and maximize the ease-of-use of the radar by the operator. Each alternative could be evaluated by cognitive psychologists and human-machine interface industrial engineering specialists. In this case each evaluation could result in a simple rating scheme that uses “high,” “medium,” or “low” coding. Therefore both objectives are specific enough to be measurable.

The end result of this top-down approach for the Sloat radar example is the objectives hierarchy depicted in Figure 2. The key to the top-down approach is repeated application of the question “What do you mean by that?” to provide more and more detail. You stop refining the structure when the answer to the question defines a quantity that can be measured, quantified, or evaluated. When you have reached this point for each part of the hierarchy then you have completed the process and obtained what you desire—a complete description of what you mean by effectiveness and a way to measure it.

3.2 The Bottom-Up Approach

The bottom-up approach starts where the top-down approach ends: with a collection of very detailed and specific measures. These are structured, or grouped, into a hierarchy in a way that assures we are not forgetting anything important and we are not double counting objectives. The key to this approach is the repeated application of the question: “Why is that there?”

The construction of the list is accomplished in many ways. First, it can be considered a Christmas “wish list.” The decision maker can be asked to list everything he or she would like to have in an alternative. For example, with Sloat Radar the decision maker may respond with the following. “I’d like to have maximum effective range and complete immunity to electronic ‘jamming.’ I’d also like to have it very easy to use so my least technically adept soldier could operate it.” Second, specific measures can be found in the “symptoms” listed in the original problem statement. If the Army is upset with the existing radar availability then it is obvious that the decision maker would also be interested in improving or maximizing availability.

For each item listed the decision maker is asked: “Why is that there?” The response to this question will provide information that aids in grouping the measures. This makes it easy to define higher level objectives for each group. For example, electronic counter measure resistance and range are included so that “when the radar works, it has enough capability to do its job better than any other alternative.” This response may bring to mind performance issues, and we group these two measures under an objective that seeks to maximize performance. Once performance is included as a higher objective it may provoke consideration of other ways one interprets performance, and this may make the decision maker think of high interoperability. The same process applied to

the ease-of-use measure would lead to considering cognitive loading on the operator and a general concern with user complexity.

Finally, this process aids the decision maker in clarifying the higher level objectives. A structure emerges that helps insure nothing is forgotten and nothing is double counted. For example, once the higher objectives of maximizing performance and minimizing complexity are evoked, the decision maker will be able to see what is important: (1) the radar has got to work almost all the time, i.e., maximize availability; (2) when it works it must be the best, i.e., maximize performance; and (3) it should be easy to operate or else all the other stuff is not worth anything, i.e., minimize complexity. The result is a hierarchy of objectives as in the Figure 2.

3.3 Which Approach to Use

Each approach has its advantages and disadvantages. The top-down approach enforces a structure from the very beginning of the exercise. It is, however, often difficult to think in general terms initially. Humans find it easier to focus on specifics, like range and availability. The bottom-up approach is more attractive in this respect, but this will not produce a logical structure without additional work. Producing an extensive wish list will, most likely, produce redundant measures. Oddly enough, the longer the list, the more likely something important will be forgotten. Long lists are harder to critically examine and locate omissions. This is where structure helps.

The best approach is perhaps a combination of the two. First, construct an hierarchy with the top-down approach, and then “reverse” direction—construct a hierarchy using the bottom-up approach. Critical examination of the results provides information for a more complete final hierarchy. The result should be a hierarchy that is:

1. Mutually exclusive (where each objective appears once).
2. Collectively exhaustive (where all important objectives are included).
3. Able to lead to measures that are operational (can actually be used).

The bottom level of the hierarchy composes the M objectives or attributes necessary to build the function for effectiveness. The Sloat Radar hierarchy has $M = 6$ attributes (availability, interoperability, ECCM, range, cognitive load, and ease-of-use), which will be used to compare among the different alternatives.

3.4 The Types of Effectiveness Measures

The effectiveness measures that result from development of the objectives hierarchy can be of three general types: (1) *natural* measures; (2) *constructed* measures; and (3) *proxy* measures.

3.4.1 Natural measures

Natural measures are those that can be easily counted or physically measured. They use scales that are most often in common use. Radar range, interoperability (as measured by the number of communication links provided), and availability (as measured by the probability that the system will be functioning at any given time) are examples. Weight, payload, maximum speed, size, volume, and area are other examples. Whenever possible, we should try to refine our objectives definitions in the hierarchy to obtain this type of measure.

3.4.2 Constructed measures

This type of measure attempts to measure the degree to which an objective is attained. It often results when no commonly used physical or natural measure exists. For example, cognitive load is a constructed measure. It is assessed as high, medium, or low depending on the number of visual and audible signals to which the operator must attend. Suppose indicators of operator cognitive abilities as measured by aptitude scores, training rigor, and education level indicate significant error in operator function if the combination of audio and visual cues is greater than 5. A radar with 5 or more cues is then assessed as having a high cognitive load. Similarly, a combination of cues between 3 and 4 constitutes a medium cognitive load, and a combination of cues less than 3 constitutes a low cognitive load. Constructed measures are the next best thing to natural measures.

3.4.3 Proxy measures

These measures are like natural measures in that they usually are countable or physically measurable. The difference is that they *do not directly* measure the objective of concern. In one sense we could consider radar range to be a proxy for measuring how much an alternative helps to maximize warning time to react better to an airborne threat. The actual reaction time cannot be directly measured because we would need to know the exact attack speed of the threatening bombers. We do know, however, that the farther out we can detect an attack, the more reaction time we would have. So maximizing range is a proxy measure for maximizing time to react.

4 DECISION MAKER PREFERENCES

After building the objectives hierarchy, we know what is important and have created a list of the relevant objectives. Each alternative can be evaluated using the list of individual effectiveness measures for the different attributes, but these measures involve a variety of incommensurable units. We must develop a way to convert all these disparate measures to a common unit of measure. This common unit of measure represents value in the mind of the decision maker.

Calculating the MOE for the j^{th} alternative as represented by the function $v(j)$ requires two types of preference information: (1) information that expresses preference for more or less of a single attribute and (2) information that expresses relative importance between attributes. The first represents preference information *within* a single attribute, while the second represents preference information *across* different attributes.

The first is important because of the way a decision maker values marginal changes in a single effectiveness measure. For example, in the Sloat Radar case, a decision maker may value the increment in range from 200 km to 400 km more than twice as much as the same increment from 400 km to 600 km. The same increment in range (of 200 km) from 600 km to 800 km may be valued very little. This valuation is quite reasonable. More range is preferred to less range, but there is a point beyond which additional range has little additional value to the decision maker. We already will have “enough” range for the purposes of early warning against air attack. Such information is very important for another reason. It allows us to construct an *individual attribute value function* that provides a scaling function to convert the natural units of measurement into units of value on a scale of 0 – 1. This has the advantage of removing any problem caused by the units of measurement for the individual attribute. For example, measuring radar range in terms of meters, kilometers, or thousands of kilometers can influence the numerical analysis and ultimately the answer. We need a process that is *independent of the units of measurement*. This is one of the important things we achieve with an individual attribute value function. Such a function also makes the process independent of the range of values of the individual effectiveness measures. This preference information will provide the answer to the question: “How much is enough?”

The second type of information is important because a decision maker values some individual effectiveness measures of attributes more than others. For example, in the Sloat Radar problem, the decision maker may feel it much more important to increase availability than to increase range. Such information implies that a decision maker would be willing to obtain more availability at the expense of less range. This preference information will provide the answer to the question: “How important is it?”

4.1 How Much Is Enough (of an Individual Attribute)?

Once we know all the relevant effectiveness measures, we need to describe how a decision maker values marginal changes in each measure. This is done by constructing a function that converts the nominal measurement scale into a 0 – 1 value scale where 0 is least preferred and 1 is most preferred. There are several ways of doing this but only one method is presented here. It is a three step procedure.

First, divide the nominal scale into “chunks,” or intervals, over which the decision maker can express preference information. For example, consider radar range. The decision maker may find it difficult to express preferences for each additional kilometer of range but may find it easier to do so if we consider range in 100 kilometer chunks.

Second, ask the decision maker to tell you, using a scale of 0 – 10, how valuable is the first 100 km of radar range. Then ask how valuable is the next 100 km of range (the second increment from 100 km to 200 km). Repeat this process for each additional 100 km of range until you reach the upper limit of interest. Suppose this exercise gives the following sequence when applied to range measurements between 0 km and 1000 km: 10, 10, 8, 7, 5, 3, 1, 0.5, 0.1, 0.05. These are the *increases in marginal effectiveness* to the decision maker.

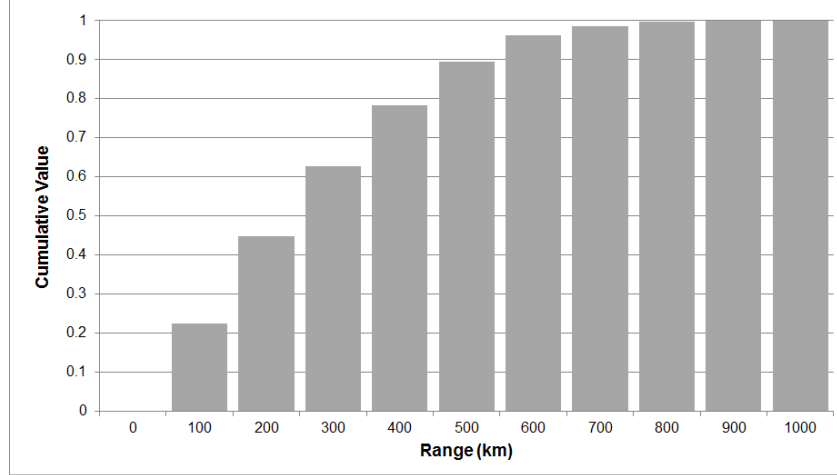


Figure 3: Radar range value function

Third, use this marginal information to create a cumulative value function scaled to give value scores between 0 and 1. The radar range example gives a total cumulative value (the sum of the marginal increments) equal to $10 + 10 + 8 + 7 + 5 + 3 + 1 + 0.5 + 0.1 + 0.05 = 44.65$. Divide all marginal increments by this sum: $10/44.65, 10/44.65, 8/44.65, 7/44.65$, etc. We obtain a sequence of smaller numbers that add to one. The cumulative value, calculated by adding the marginal values, is what we use to determine the value of a measure. For example, the value of 300 km range is equal to the value derived from increasing the range from 0 to 100 km plus the additional value derived from increasing the range from 100 km to 200 km plus the additional value derived from increasing the range from 200 km to 300 km. The cumulative value of these normalized marginal increments is presented graphically in Figure 3.

This method achieves two important things. First, if we repeat this process for each attribute, we have a common unit of measure standardized to a 0 – 1 interval. Second, and most importantly, we now have a way of valuing radar range. A range of 700 km has a value of 0.985 and a range of 800 km has a value of approximately 0.997, which informs us that the decision maker prefers 800 km to 700 km but that this increase is not highly valued (only an increase in value of 0.012). On the other hand, increasing the range from 200 km (with a value of 0.448) to 300 km (with a value of 0.628) is much more valuable to the decision maker. The increase in value here is 0.18. Thus we know the decision maker prefers the 100 km increment from 200 to 300 km more than the going from 700 to 800 km.

The analysis of effectiveness is facilitated by converting this graphical information to algebraic form so we can obtain a model of preferences that allows us to quantify these preferences for every possible numerical measure for an attribute. We construct a value function for an individual attribute where x_i represents the numerical value for the i^{th} attribute and $v_i(x_i)$ is the value function for the i^{th} attribute. A useful algebraic form for capturing the important characteristics of Figure 3 is the *exponential function*:

$$v_i(x_i) = \frac{1 - e^{-a[x_i - x_{\min}]}}{K} \quad (1)$$

where $K = 1 - e^{-a[x_{\max} - x_{\min}]}$ is a normalization constant. This function requires three parameters: x_{\min} , x_{\max} , and a . The first two are known already because the decision maker has specified the range of interest over which x_i will vary: $[x_{\min} \leq x_i \leq x_{\max}]$. The parameter a can be obtained using least-squares regression using the Solver “add-in” tool in Excel. The parameter can also be estimated graphically by adjusting a until the curve generated by the function $v_i(x_i)$ corresponds to the points representing the decision maker’s preferences in the cumulative value function as ascertained by the three-step method. More flexible but more complicated forms of this exponential function can be used (see the Appendix). The purpose of these algebraic forms is to match the original cumulative value function given to us by the decision maker as closely as possible.

Calculating a numerical value for a in $v_i(x_i)$ that corresponds to the information for the range of the radar, as depicted in Figure 3, is done in Excel using the Solver add-in tool. Figure 4 presents $v_i(x_i)$ corresponding to the




Figure 4: Fitted exponential value function for radar range




Figure 5: Radio weight value function

data in Figure 3 when we use an exponential form with $a = 0.0033$, $x_{\min} = 0$, and $x_{\max} = 1000$.

The preceding example has illustrated a case where “more is better.” Exactly the same approach is used when “less is better.” For example, consider a man-portable field radio. Here less weight is preferred to more weight. Now the decision maker is asked a series of questions relating to marginal value but we “start from the other end” of the measurement scale—we start with the heaviest weight and work backwards to the lightest weight. For example, suppose it is decided that field radios weighing more than 20 kg are useless because they weigh too much for a typical soldier to carry. Imagine we use marginal changes in weight of 2 kg. We begin by asking the decision maker to give us a number (on a scale of 0 – 10) expressing the value of a 2 kg reduction (from a weight of 20 kg to a weight of 18 kg). Then we ask what value is associated with a further 2 kg reduction in weight (to 16 kg). We repeat the process until we get an answer for the value attached to the last 2 kg of weight. Suppose we obtain the sequence: 10, 9, 8, 5, 3, 1, 0.4, 0.1, 0.0. These sum to 36.5 so we normalize the decrements by this total. The result is presented in Figure 5.

As in the case where more is better, the graph of cumulative value can be used to fit an exponential function for conducting analysis of effectiveness. Once again, exponential functions are very useful and can be used to represent this information:

$$v_i(x_i) = \frac{1 - e^{-a[x_{\max} - x_i]}}{K} \quad (2)$$




Figure 6: Fitted exponential value function for radio weight

where K is the same as before. Figure 6 depicts the results of fitting the cumulative value function for radio weight in Excel using the Solver add-in tool where $a = 0.21$, $x_{\min} = 0$, and $x_{\max} = 20$. Once again, closer fit can be obtained by using more terms in the exponent as described in the Appendix.

A linear function can also be used to approximate the exponential function for either the more-is-better case or the less-is-better case. The linear function can use the same cumulative value function as described previously, or just two values can be assessed from the decision maker. The two values necessary are a number corresponding to “not enough performance” and a number corresponding to “good enough performance.” If more is better, the number corresponding to not enough performance is “too little” or $x_{\text{too little}}$. If less is better, the number corresponding to not enough performance is “too much” or $x_{\text{too much}}$. In both cases, the number corresponding to good enough performance is “ideal” or x_{ideal} . For example, in the more-is-better example of radar range, 0 km can represent not enough performance or too little and 600 km can represent good enough performance or the ideal. We choose 600 km rather than 1000 km as the ideal because the assessed value at 600 km is 0.963 and the increase in value from 600 to 1000 km is only 0.037. In the less-is-better example of the radio’s weight, we choose 20 kg as too much and 8 kg as ideal. Eight kilograms is chosen as ideal because the value at 8 kg is 0.986 and the increase in value from 8 to 0 kg is only 0.014. The linear value function is defined for more is better in which $[x_{\text{too little}} \leq x_i \leq x_{\text{ideal}}]$:

$$v_i(x_i) = \frac{x_i - x_{\text{too little}}}{x_{\text{ideal}} - x_{\text{too little}}}, \quad (3a)$$

and for less is better in which $[x_{\text{ideal}} \leq x_i \leq x_{\text{too much}}]$:

$$v_i(x_i) = \frac{x_i - x_{\text{too much}}}{x_{\text{ideal}} - x_{\text{too much}}}. \quad (3b)$$

If $x_i \geq x_{\text{ideal}}$ in the more-is-better case or if $x_i \leq x_{\text{ideal}}$ in the less-is-better case, $v_i(x_i) = 1$. Similarly, $v_i(x_i) = 0$ if $x_i \leq x_{\text{too little}}$ in the more-is-better case or $x_i \geq x_{\text{too much}}$ in the less-is-better case.

4.2 How Important Is It (Relative to the Other Attributes)?

After assessing the preferences of the decision maker for changes in each attribute, we need to know the preferences of the decision maker among different attributes. The model for the MOE uses *importance weights* or *trade-off weights* denoted by w_i , the importance weight for the i^{th} attribute. Each importance weight must satisfy $0 \leq w_i \leq 1$ and together must satisfy $\sum_{i=1}^M w_i = 1$. A larger value of w_i signifies that the effectiveness measure for the i^{th} attribute is more important.

The conditions imposed on the weights will automatically be satisfied if we apply the conditions to *each of the partial hierarchies that appear within the overall objectives hierarchy*. For example, the overall hierarchy in Figure 2 is actually composed of three hierarchies. In the first hierarchy, defined by the top level and the first level

underneath it, maximizing overall effectiveness is supported by maximizing availability, maximizing performance, and minimizing user complexity. The relative importance of each of these three objectives to maximizing overall effectiveness can be expressed with three weights such that $0 \leq w_A, w_P, w_C \leq 1$ and $w_A + w_P + w_C = 1$ (A = availability, P = performance, and C = complexity). In the second hierarchy, maximizing performance is supported by maximizing interoperability, maximizing ECCM capability, and maximizing range. The relative importance of each of these can be expressed with three more weights such that $0 \leq w_I, w_E, w_R \leq 1$ and $w_I + w_E + w_R = 1$ (I = interoperability, E = ECCM, and R = range). Finally, minimizing complexity is supported by minimizing cognitive load and maximizing ease-of-use. The relative importance of these two objectives can be expressed with two weights such that $0 \leq w_L, w_U \leq 1$ and $w_L + w_U = 1$ (L = cognitive load and U = ease-of-use).

Assigning weights in this fashion guarantees that the weights for all M objectives or attributes sum to one and are all within the range of 0–1. The Sloat Radar example has $M = 6$ attributes, and we need 6 importance weights, w_1, w_2, w_3, w_4, w_5 , and w_6 to correspond to each one of the attributes (availability, interoperability, ECCM, range, cognitive load, and ease-of-use). These 6 “global” weights can be derived from the “local” weights described in the previous paragraph. Availability stands alone in the hierarchy, and $w_A = w_1$ where w_1 is the global weight for availability. Performance is composed of three attributes, and $w_P = w_2 + w_3 + w_4$ where w_2 is the global weight for interoperability, w_3 is the global weight for ECCM, and w_4 is the global weight for range. Complexity is composed of two attributes, and $w_C = w_5 + w_6$ where w_5 and w_6 are the global weights for cognitive load and ease-of-use, respectively. Using the fact that $w_I + w_E + w_R = 1$ for the performance sub-hierarchy and $w_L + w_U = 1$ for the complexity sub-hierarchy, we can express the global weights for the 6 attributes:

$$w_1 = w_A \quad (4a)$$

$$w_2 = w_P \cdot w_I \quad (4b)$$

$$w_3 = w_P \cdot w_E \quad (4c)$$

$$w_4 = w_P \cdot w_R \quad (4d)$$

$$w_5 = w_C \cdot w_L \quad (4e)$$

$$w_6 = w_C \cdot w_U. \quad (4f)$$

Since each group of local weights sum to one, we know the global weights as calculated above will also sum to one:

$$w_1 + w_2 + w_3 + w_4 + w_5 + w_6 = w_A + w_P + w_C = 1.$$

Obtaining the weights requires solicitation of decision maker preferences using a different set of questions from that used above. There are many ways to do this, and we consider five methods.

4.2.1 Direct assessment

The most obvious way to gain the information required is to directly ask the decision maker for M numbers, w_i , such that $0 \leq w_i \leq 1$ for $\{1 \leq i \leq M\}$ and $\sum_{i=1}^M w_i = 1$. If there are no more than three or four effectiveness measures this is a feasible approach. Unfortunately many real-life problems have many more effectiveness measures and direct assessment becomes too difficult. The Sloat Radar example has six effectiveness measures, and this presents a formidable problem to the decision maker. All is not lost, however, if we utilize the structure given us in the objectives hierarchy.

Consider the Sloat Radar hierarchy. We can ask the decision maker to consider first just the relative importance of the three components that make up overall effectiveness: availability, performance, and user complexity. Directly assessing three importance weights at this level presents no insurmountable problem for the decision maker. Next we can ask the decision maker to consider the components of performance and tell us the relative importance of interoperability, ECCM capability, and range. Finally, we can ask the same question for the components of user complexity that requires only two importance weights be assigned. The results of this series of questions allows us to compute the complete set of six weights using Eqs. (4a) - (4f).

Using the hierarchy and this “divide and conquer” approach makes direct assessment feasible in many situations where otherwise it would appear overwhelming.

4.2.2 Equal importance

Direct assessment may result in the decision maker stating that “all individual measures are equally important.” In this case we must ask the decision maker if this statement applies to all M attributes or to each part of the hierarchy. Different importance weights result depending on the interpretation of the statement.

If the decision maker is referring to all M attributes, then clearly $w_i = 1/M$. In the Sloat Radar example this means that each measure would be given a weight of $1/6$. It is important to realize that this weighting has implications for the corresponding weights in the hierarchy. As discussed earlier, the importance weight for performance equals the sum of global weights for interoperability, ECCM, and range, and $w_P = 1/6 + 1/6 + 1/6 = 1/2$. The weight for complexity equals the sum of the weights for cognitive load and ease-of-use, and $w_C = 1/6 + 1/6 = 1/3$.

Clearly, equal importance across the bottom level measures does NOT imply equal weights throughout the hierarchy. Thus, it is very important to go back to the decision maker and ask: “Do you really believe performance is three times more important than availability?” “Do you really believe that minimizing complexity is twice as important as maximizing availability?” When phrased this way the decision maker may be led to reassess the importance weights.

If the decision maker is referring to the weights at each level in the objectives hierarchy, then we get a different set of weights for w_i . Figure 2 tells us that on the first level down availability, performance and complexity are all equally important:

$$w_A = w_P = w_C = 1/3.$$

This also means that all the component objectives for performance are equally important:

$$w_I = w_E = w_R = 1/3.$$

Finally, suppose the decision maker says that maximizing ease-of-use and minimizing cognitive load are equally important, then we know:

$$w_U = w_L = 1/2.$$

These local weights translate into the following global weights:

$$w_1 = 1/3, \quad w_2 = 1/3 \cdot 1/3 = 1/9, \quad w_3 = 1/3 \cdot 1/3 = 1/9, \\ w_4 = 1/3 \cdot 1/3 = 1/9, \quad w_5 = 1/3 \cdot 1/2 = 1/6, \quad w_6 = 1/3 \cdot 1/2 = 1/6.$$

A very different picture emerges. Availability is three times more important than interoperability, ECCM capability, and range. Availability is twice as important as cognitive load and ease-of-use.

The structure of the hierarchy is a powerful source of information that we will continually find helpful and informative.

4.2.3 Rank sum and rank reciprocal

Sometimes the decision maker will provide only *rank order* information. For example, in assessing the weights making up the performance measure, the decision maker may tell us: “Range is most important, interoperability is second most important, and ECCM is least important.” In such cases we can use either the *rank sum* or *rank reciprocal* method. Range has rank one, interoperability has rank two, and ECCM has rank three.

In the rank sum method, range has an un-normalized weight of 3, interoperability 2, and ECCM 1, based on the decision maker’s ranking. Dividing each of these weights by the sum $3 + 2 + 1 = 6$ returns the normalized weights: $w_R = 3/6$, $w_I = 2/6$, and $w_E = 1/6$.

In the rank reciprocal method, the reciprocals of the original ranks give:

$$\begin{aligned} \text{Range} &= 1/1 \\ \text{Interoperability} &= 1/2 \\ \text{ECCM} &= 1/3 \end{aligned}$$

which sum to $11/6$. Dividing the reciprocal ranks by their sum gives three numbers that satisfy the summation condition and represent a valid set of weights: $w_R = 6/11$, $w_I = 3/11$, and $w_E = 2/11$.

Both the rank sum and rank reciprocal methods only require rank order information from the decision maker. The rank sum method returns weights that are less dispersed than the rank reciprocal method, and less importance is placed on the first objective in the rank sum method.

4.2.4 Pairwise comparison

Decision makers find it easier to express preferences between objectives where there are only two objectives. This has led to a procedure for soliciting preference information based on pairwise comparisons. This allows the importance weights over M measures to be established by comparing measures only two at a time. This is a special case of the swing weighting method, which is described subsequently.

4.2.5 Swing weighting

One method helpful with direct assessment that incorporates more than rank order information is *swing weighting*. Although more complex than the other methods, the swing weighting method can provide a more accurate depiction of the true importance the decision maker places on each objective. Swing weighting is also sensitive to the range of values that an attribute takes, and different values for attributes can result in different weights. The method has four steps. First, the decision maker is asked to consider the increments in *overall* effectiveness that would be represented by shifting, or “swinging,” each individual effectiveness measure from its least preferred value to its most preferred value. Second, the decision maker is asked to order these overall increments from least important to most important. Third, the decision maker quantitatively scales these increments on a scale from 0 to 100. Finally the sum-to-one condition is used to find the value of the least preferred increment.

Let us illustrate this process using the performance part of the Sloat Radar hierarchy. We need to find three importance weights: w_I , w_E , and w_R . First, suppose that the least preferred under performance are 0 km for range, 0 communication links for interoperability, and no ECCM. This “least-preferred alternative” with 0 km in range, no communication links, and no ECCM has a value of 0 for performance. The most preferred under performance are 1000 km for range, 5 communication links, and with ECCM. The decision maker believes that a swing in range from 0 to 1000 km will contribute most to performance, a swing from no communication links to five links will contribute second most to performance, and a swing from no ECCM capability to having ECCM capability will contribute the least to performance. Because range is the most preferred swing, we arbitrarily assign a value of 100 to a hypothetical alternative with 1000 km range, 0 communication links, and no ECCM. We ask the decision maker for the value of swinging from 0 to 5 communication links relative to the alternative with a value of 100 and the least-preferred alternative with a value of 0. The decision maker may say that a hypothetical alternative with 0 km range, 5 communication links, and no ECCM has a value of 67. Finally, the decision maker may conclude that a hypothetical alternative with 0 km range, 0 communication links, and having ECCM capability has a value of 33 relative to the other alternatives. Thus, we know the un-normalized weights are 100 for range, 67 for interoperability, and 33 for ECCM. Dividing each weight by the sum of the un-normalized weights gives the importance weight for each attribute: $w_R = 100/(100 + 67 + 33) = 1/2$, $w_I = 67/(100 + 67 + 33) = 1/3$, and $w_E = 33/(100 + 67 + 33) = 1/6$.

These weights express something very important in effectiveness analysis. When x_i is at its least preferred value, $v_i(x_i) = 0$. When it is at its most preferred value $v_i(x_i) = 1$. The weights, $\{w_i, 1 \leq i \leq M\}$, capture the relative importance of these changes. The ratio $w_i/w_{i'}$ expresses the relative importance between the changes from worst to best in the measures x_i and $x_{i'}$ for the i^{th} and i'^{th} attributes.

Swing weighting incorporates three types of preference information: ordinal (rank), relative importance, and range of variation of the individual effectiveness measures. Pairwise comparison uses rank and relative importance. The rank sum and rank reciprocal methods only use rank information.

5 EFFECTIVENESS ANALYSIS

The quantification of decision maker preferences completes our model of effectiveness. We can begin assessing or analyzing the overall effectiveness of each alternative. All the ingredients are present: The objectives hierarchy tells us what is important and defines the individual measures of effectiveness. The individual value functions tell us how the decision maker values marginal increments in these measures and scales them to a 0–1 interval. Finally, the importance weights tell us the relative importance of the individual attributes. The weights are combined with the values for a single alternative to calculate the overall effectiveness of the j^{th} alternative:

$$v(j) = w_1 \cdot v_1(x_1(j)) + w_2 \cdot v_2(x_2(j)) + w_3 \cdot v_3(x_3(j)) + \cdots = \sum_{i=1}^M w_i \cdot v_i(x_i(j)). \quad (5)$$

Table 2: Evaluation data for Sloat Radar

Alternative	Availability	Interoperability	ECCM	Range (km)	Cognitive load	Ease-of-use
Sloat 2	0.79	2 data links	No	250	Low	Medium
Sloat 3	0.87	2 data links	No	250	Low	Medium
SkyRay	0.98	4 data links	No	700	Medium	Low
Sweeper	0.97	2 data links	Yes	500	High	Low

The MOE is calculated as the weighted sum of all the individual value functions. Nothing of relevance is left out, and everything that enters the computation does so according to the preferences of the decision maker.

Evaluating the overall effectiveness of an alternative requires computing its $v(j)$. The result allows us to order all the alternatives from best to worst according to their MOE. The “best” alternative is now the alternative with the largest MOE. While this provides us a way to find “the” answer, we have a far more powerful tool at hand.

We can use Eq. (5) to investigate why we get the answers we do. For example, what makes the best alternative so desirable? What individual effectiveness measures contribute most to its desirability? An alternative may be the best because of very high effectiveness in only one single measure. This may tell us that we could be “putting all our eggs in one basket.” If this alternative is still under development, uncertainties about its development may make this alternative risky, and identifying the second best alternative may be important.

We can also assess the sensitivity of the answer to the importance weights. We can find out how much the weights must change to give us a different answer. For example, would a change of only 1% in one of the higher level weights change the ordering of the alternatives? How about a change of 10%? This is of practical significance because the decision maker assigns weights subjectively, and this always involves a lack of precision.

Most important of all is the ability to assess the effect of uncertainty in the future condition. The decision maker’s preferences are a function of the future condition. If the decision maker believes the future condition will change, then the preferences, value function, and weights may change. The effects of uncertainties in the problem formulation can be readily evaluated once we have our model.

We illustrate each of these situations using the Sloat Radar example. Suppose there are four alternatives: (1) “do nothing” (keep the two existing radars at Sloat which is labeled as Sloat 2); (2) purchase a third radar for Sloat of the same type (which is labeled as Sloat 3); (3) purchase the new SkyRay radar; and (4) purchase the new Sweeper radar. The latter two alternatives are new, with better range, availability, interoperability and ECCM. These come at the cost of higher cognitive load and less ease-of-use. The costs of procurement for the two new radars are also higher than purchasing an existing radar. The data on which each of these four alternatives are depicted in Table 2.

5.1 The Components of Success

The decision maker preferences for the importance weights are $w_A = 0.60$, $w_P = 0.35$, $w_C = 0.05$, $w_I = 0.50$, $w_E = 0.0$, $w_R = 0.50$, $w_L = 0.50$, and $w_U = 0.50$. Individual value functions for availability and range are specified by value functions similar in shape and form to that portrayed in Figure 4. The individual value function for interoperability is similar but specified for integer values, 0 – 5. The individual value function for ECCM capability is binary, taking the value 0 for no ECCM capability and the value 1 for having ECCM capability. Cognitive load and ease-of-use are evaluated using a constructed scale of three categories: low, medium, and high. For cognitive load, $v_L(\text{low}) = 1$, $v_L(\text{medium}) = 0.5$, and $v_L(\text{high}) = 0$. For ease-of-use $v_U(\text{low}) = 0$, $v_U(\text{medium}) = 0.5$, and $v_U(\text{high}) = 1$. Combining these evaluations and preferences gives the values depicted in Table 3.

Multiplying each of the global weights by the each of the values and adding them together calculates an MOE for each radar, as pictured in Figure 7. SkyRay is the most effective alternative and we know why. It has the highest value for interoperability and range. Sweeper is second best because it has the second highest value for range. All alternatives possess approximately the same availability rating. Interoperability and range are the attributes that are most important in discriminating among these alternatives.

Table 3: Values and weights for Sloat Radar

Alternative	Availability	Interoperability	ECCM	Range (km)	Cognitive load	Ease-of-use
Sloat 2	0.962	0.375	0	0.413	1.0	0.5
Sloat 3	0.982	0.375	0	0.413	1.0	0.5
SkyRay	0.998	0.938	0	0.909	0.5	0
Sweeper	0.997	0.375	1	0.774	0	0
Global weights	0.60	0.175	0	0.175	0.025	0.025

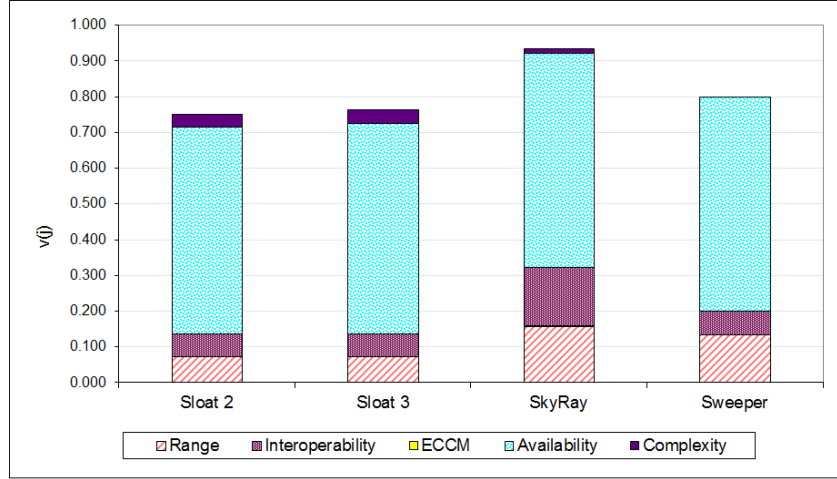


Figure 7: Components of overall effectiveness for Sloat Radar

5.2 Sensitivity to Preferences

The original importance weights produce an ordering of alternatives in which SkyRay is most effective, Sweeper next most effective, followed by Sloat 3 and then Sloat 2. Is this ordering robust to reasonable changes in the weights? If not, which weights are most influential? Questions like these are answered using the model to recompute the MOE of each alternative with different importance weights.

First let us fix $w_P = 0.35$ and vary both w_A and w_C over a reasonable range of values subject to the condition that $w_A + w_P + w_C = 1$. The result is shown in Figure 8. SkyRay remains the most effective alternative for the range of weights examined, which means that the most effective alternative is fairly robust to changes in the importance weights for availability relative to complexity. The ordering of the other alternatives changes if $w_A \leq 0.55$ or, correspondingly, when $w_C \geq 0.10$. As availability becomes less important relative to complexity, Sweeper becomes the least effective radar, and Sloat 3 becomes the second most effective alternative. When determining the second most effective alternative, the decision maker needs to ask himself: “How confident am I that $w_A > 0.55$ when $w_P = 0.35$?”

Attention now focuses on the influence of w_P . We repeat the above analysis for w_A and w_C for different values of w_P . If $w_P = 0.25$, SkyRay remains for the most effective alternative as long as $w_A \geq 0.5$ and $w_C \leq 0.25$. As depicted in Figure 9, if $w_P = 0.2$, Sloat 3 becomes the most effective if $w_A \leq 0.55$ or $w_C \geq 0.25$. The question to ask is: “How likely is it that $w_A \leq 0.55$, $w_P \leq 0.2$, and $w_C \geq 0.25$?” If the decision maker responds that he will not put that much importance on complexity relative to availability and performance, then SkyRay remains the most effective alternative.

5.3 Uncertainty in the Future Conditions

A new future condition or planning scenario may change the decision maker’s preferences and consequently, the ordering of the alternatives. Investigating this type of sensitivity is very important when there is uncertainty in the future condition. Let us assume the situation described in Section 5.1 and depicted in Figure 7 corresponds to a particular planning scenario we will call the baseline.

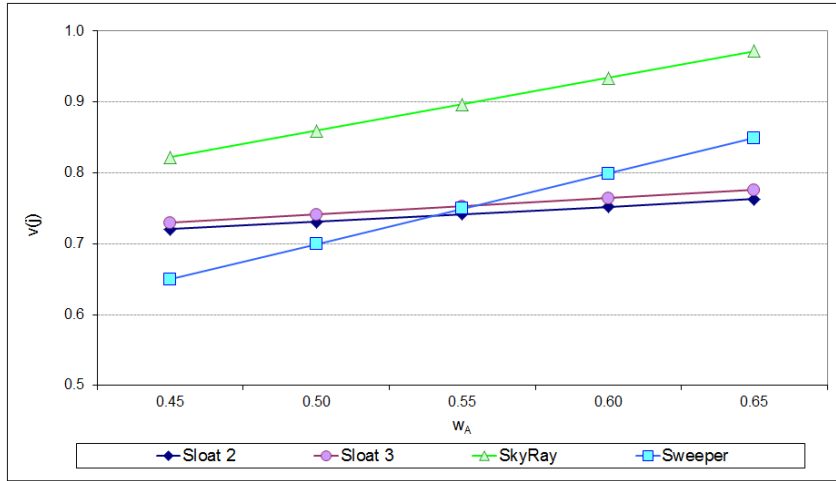


Figure 8: Sensitivity to w_A and w_C with $w_P = 0.35$




Figure 9: Sensitivity to w_A and w_C with $w_P = 0.2$

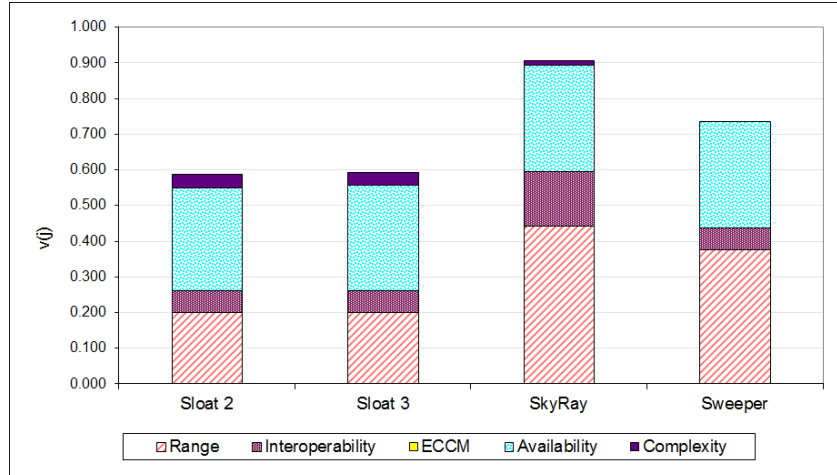


Figure 10: Supersonic attack scenario




Figure 11: ECM scenario

Now suppose new intelligence estimates suggest the potential enemy will re-equip its bomber fleet with supersonic attack aircraft. This new scenario may change decision maker preferences to place more importance on performance because of the enemy's increase in capability. For example, the importance weights defining the measure of overall effectiveness may change to $w_A = 0.30$, $w_P = 0.65$, and $w_C = 0.05$. Furthermore, the importance weights defining the components of performance could change to $w_I = 0.25$, $w_E = 0.0$, and $w_R = 0.75$. This preference structure gives the result depicted in Figure 10. The ordering of the alternatives is unchanged and the picture resembles that in Figure 7. There are differences, however, in the composition of the MOE for each alternative. Range plays a more important role, and availability plays a less important role. Nevertheless, the original ordering of alternatives is robust to this change in the future condition.

Now assume that intelligence reports indicate the potential enemy cannot afford to reequip its bomber fleet with new aircraft. Instead, they have decided on an avionics upgrade to give the existing bomber fleet radar jamming capability or electronic counter measures (ECM). Under this future condition the decision maker may still favor performance over availability, and $w_A = 0.30$, $w_P = 0.65$, and $w_C = 0.05$. ECCM becomes important, and the weights within performance may change to $w_I = 0.10$, $w_E = 0.75$, and $w_R = 0.15$. The resulting MOE are depicted in Figure 11. Sweeper is now the most effective alternative for this future condition because it is the only radar with ECCM capability.

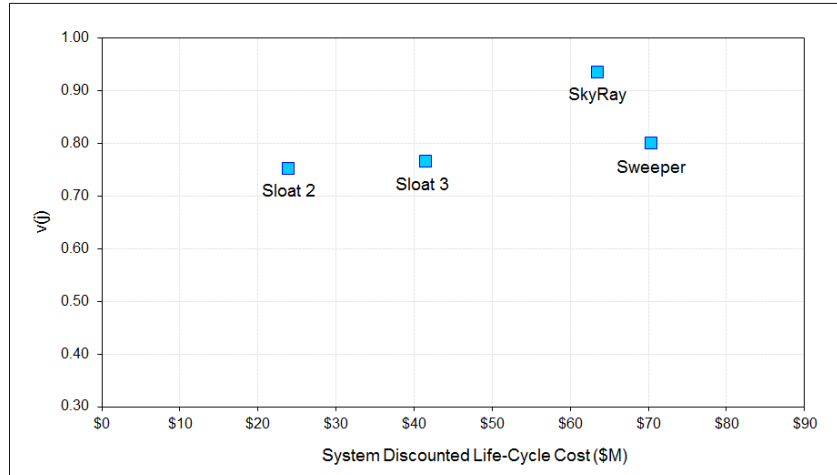


Figure 12: Alternatives in cost-effectiveness space

6 COST-EFFECTIVENESS

Computing the MOE for each alternative allows the decision maker to order the alternatives from most effective to least effective. This is only half the story, though, because cost also matters. Ultimately the decision maker will have to integrate cost information with effectiveness and engage in *cost-effectiveness analysis*. This section presents the conceptual elements used in the analysis and outlines the process followed in the analysis.

6.1 Conceptual Framework

The decision maker ultimately pursues two overall objectives when searching for a solution: (1) maximize effectiveness and (2) minimize cost. An alternative j will be evaluated in terms of its MOE, $v(j)$, and its discounted life-cycle cost, $c(j)$. Drawing a picture of effectiveness and cost on a scatter plot can help us think in these two dimensions at the same time. Cost is plotted on the horizontal axis, and effectiveness is plotted on the vertical axis. If the four alternatives in the Sloat Radar example have the following discounted life-cycle costs and MOEs (with the importance weights from the baseline) in the table below, Figure 12 depicts this information graphically.

Alternative	Cost (millions of dollars)	Effectiveness (MOE)
Sloat 2	23.9	0.752
Sloat 3	41.6	0.765
SkyRay	63.5	0.934
Sweeper	70.4	0.799

This framework makes it easy for the decision maker to see which alternative is cheapest or most expensive and which alternative is most effective or least effective. It shows the distance between alternatives—how much more effective and/or costly is one alternative versus another. Such a picture provides all the information the decision maker needs to choose the most cost-effective alternative, the exact meaning of which requires understanding the many ways the solution can be interpreted.

6.2 Solution Concepts

Multiple ways exist to define a solution that minimizes cost and maximizes effectiveness. The solutions are distinguished from each other based on the amount of additional information required of the decision maker. First, we present two concepts that do not require any more information from the decision maker. Second, we develop solution concepts that require a little additional information from the decision maker. Finally, we conclude with a definition for the most cost-effective solution. This last concept requires eliciting additional decision maker preference information and is the most demanding.




Figure 13: Superior or dominant solution




Figure 14: Efficient solution

6.2.1 Superior solution

A superior solution is a feasible alternative that has the lowest cost and the greatest effectiveness. It does not matter if you are a decision maker who places all emphasis on minimizing cost or a decision maker who places all emphasis on maximizing effectiveness. Both types of decision makers would select the superior solution if one exists. Figure 13 illustrates this concept. Alternative 1 is superior to all the others, and it is called the *dominant* solution because this alternative is the cheapest and the most effective. We do not need any preference information beyond that which we have already obtained in order to define effectiveness. No alternatives exist in the “north-west” quadrant relative to alternative 1 because no alternative is either cheaper or more effective than alternative 1.

6.2.2 Efficient solution

The efficient solution concept builds upon the superior solution concept. An *efficient* solution is one that is not dominated by, or inferior to, any other feasible alternative, and an alternative is not efficient because there exists another alternative that is superior to it. An example of this type of solution is found in Figure 14. Alternatives 4 and 6 are dominated by, or are inferior to, alternatives 3 and 5, respectively. Consequently, alternatives 4 and 6 should not be selected. Alternatives 2, 3, 5, and 7 are all efficient.

This solution concept does not necessarily yield unique answers. If only one efficient solution exists, it is the superior solution. If more than one efficient solution exist, a set of non-unique solutions or alternatives exists. This may not be a bad circumstance, however, because the decision maker has flexibility, and several alternatives can be defended as efficient.

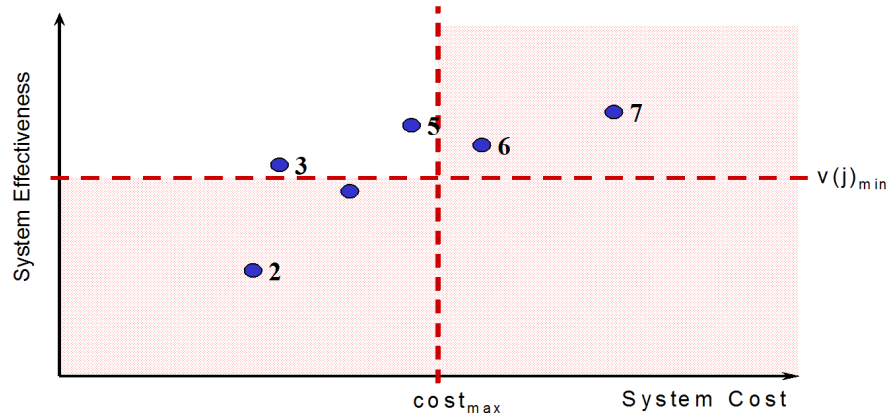


Figure 15: Satisficing solution

6.2.3 Satisficing solution

A satisficing solution is a feasible solution that is “good enough” in the sense that it exceeds a minimum level of effectiveness and does not exceed a maximum cost. This solution concept requires the decision maker to state a desired MOE and a maximum life-cycle cost. An alternative that satisfies both of these requirements simultaneously is satisfactory. Like the efficient solution concept, the satisficing solution concept often yields non-unique answers. Figure 15 shows a situation where both alternatives 3 and 5 are satisfactory solutions because both alternatives are cheaper than the maximum cost (cost_{\max}) and more effective than the minimum effectiveness level ($v(j)_{\min}$). Because both alternatives are satisfactory, a decision maker will need another solution concept to decide between the two alternatives.

6.2.4 Marginal reasoning solution

If a superior solution does not exist, marginal reasoning may be used to select an alternative. This solution concept begins with an efficient set of alternatives, such as those from Figure 14, which are reproduced in Figure 16. Suppose the decision maker is considering alternative 3. Alternatives 2 and 5 are its closest neighbors: alternative 2 is similar in cost and alternative 5 is similar in effectiveness. Two questions can be asked: (1) Is the marginal increase in cost worth the marginal gain in effectiveness when moving to alternative 5? and (2) Is the marginal savings in cost worth the marginal decrease in effectiveness when moving to alternative 2? If the decision maker prefers alternative 5 over alternative 3, the additional cost is worth paying to obtain the additional increase in effectiveness. If the decision maker prefers alternative 3 to alternative 5, the marginal increase in effectiveness is not worth the marginal increase in cost. If the decision maker cannot decide between alternative 3 and alternative 5, the marginal increase in cost is balanced by the marginal increase in effectiveness, or equivalently, the cost savings just compensate for the decrease in effectiveness. It is a “toss up” between the two.

This reasoning can be applied to each alternative in the efficient set. After all the alternatives have been considered, the decision maker will arrive at one of two situations: (1) one alternative is identified as the most preferred alternative; or (2) two or more alternatives are identified as equivalent and, as a group, are more preferred than than all the rest.

6.2.5 Importance weights for effectiveness and cost

The most formal solution concept asks the decision maker to determine the relative importance of cost to effectiveness. This means there exists *in the mind of the decision maker* a payoff function that combines the two issues of concern: (1) maximizing effectiveness and (2) minimizing cost. The first objective is represented by the MOE $v(j)$. The second objective is a function of the cost measure $c(j)$ and represents a less-is-better preference relation. The decision maker needs to determine a value function for cost similar to the example given in Section 4.1 where a value function for the weight of a radio is constructed. The value function for cost $v_c(c(j))$ can be an exponential




Figure 16: Marginal reasoning solution

function as in Eq. (2) or a linear function as in Eq. (3b) in which less is better. Figure 27 in the Appendix depicts a linear value function for cost in which $c_{\text{ideal}} = \$0$ and $c_{\text{too much}} = \$90$ million.

The payoff function for the cost-effectiveness of the j^{th} alternative is denoted by $V(j)$ and follows a simple additive form:

$$V(j) = W_E \cdot v(j) + W_C \cdot v_c(c(j)) \quad (6)$$

where W_E and W_C are the importance weights for effectiveness and cost, respectively. Upper case letters distinguish these weights from those used in defining $v(j)$. W_E represents the importance to the decision maker of maximizing effectiveness and W_C represents importance to the decision maker of minimizing cost. The weights must satisfy $0 \leq W_E \leq 1$, $0 \leq W_C \leq 1$, and $W_E + W_C = 1$. The relative importance of these two conflicting objectives is captured by the ratio W_C/W_E .

Eq. (6) represents the preferences of a rational decision maker. When confronted by a choice between two alternatives equal in effectiveness, the decision maker will choose the less expensive alternative because $v_c(c(j))$ decreases as the cost $c(j)$ increases. Similarly, when confronted by a choice between two alternatives equal in cost, the decision maker will choose the alternative with the greater MOE because $v(j)$ is greater for the alternative that is more effective.

The overall cost-effectiveness function $V(j)$ defines a preference structure that allows us to develop valuable insights. We can see this graphically by rearranging the cost-effectiveness function and drawing a straight line to represent $V(j)$ on the cost-effectiveness graph. Express $v(j)$ as a function of $V(j)$ and $c(j)$:

$$v(j) = \frac{1}{W_E} V(j) - \frac{W_C}{W_E} \cdot v_c(c(j)).$$

For simplicity, we assume a linear individual value function—see Eq. (3b)—for $v_c(c(j))$ in which less is better. If we assume $c_{\text{ideal}} = 0$, then $v_c(c(j)) = [c(j) - c_{\text{too much}}]/[-c_{\text{too much}}] = 1 - k \cdot c(j)$ where $k = 1/c_{\text{too much}}$. For a fixed value of $V(j)$, we see that $v(j)$ is a linear function of $c(j)$:

$$\begin{aligned} v(j) &= \frac{1}{W_E} V(j) - \frac{W_C}{W_E} \cdot [1 - k \cdot c(j)] \\ &= \frac{1}{W_E} V(j) - \frac{W_C}{W_E} + \frac{W_C}{W_E} \cdot k \cdot c(j). \end{aligned}$$

This is the equation for a straight line when we interpret $c(j)$ as the independent variable and $v(j)$ as the dependent variable. The intercept for the line is given by $V(j)/W_E - W_C/W_E$ and the slope is given by $k \cdot W_C/W_E$. This line represents all combinations of cost and effectiveness corresponding to a given level of overall cost-effectiveness, $V(j)$. As $V(j)$ increases (for example, $V^{**} > V^*$ in Figure 17), this “isoquant” line shifts up and to the left. As we move closer to the northwest corner of the cost-effectiveness plot, we move to greater levels of overall cost-effectiveness.

Suppose the decision maker is much more interested in minimizing cost than maximizing effectiveness. This decision maker would select weights where $W_C \gg W_E$. The slope of the lines representing constant overall




Figure 17: Cost-effectiveness linear preference




Figure 18: $W_C \gg W_E$

cost-effectiveness would be very steep. This situation is depicted in Figure 18. Alternative 3 is the best in this situation because it lies on the highest achievable isoquant of cost-effectiveness. A decision maker who places more emphasis on maximizing effectiveness would choose weights that result in less steep lines of constant overall cost effectiveness, which could yield a picture like Figure 19. Now alternative 5 is the most cost-effective. Finally, consider a decision maker who places considerable importance on the maximization of effectiveness, which implies that $W_E \gg W_C$. The slope of the lines is very small resulting in nearly flat isoquants of cost-effectiveness as in Figure 20. Alternative 7 is now the most cost-effective alternative.

The above three cases illustrate the importance of the efficient set, and the most cost-effective alternative is selected from the set of efficient alternatives. The most cost-effective depends on the decision maker preferences for cost reduction versus effectiveness maximization. The answer to the decision problem requires elicitation of decision maker preferences—we do not know what we mean by “cost-effectiveness” until we incorporate decision maker preferences over cost and effectiveness.

6.3 Cost-effectiveness for Sloat Radar

Figure 12 shows no superior solution exists among the four alternatives of the Sloat Radar problem and Sweeper is dominated by SkyRay. The efficient set is composed of Sloat 2 (the “do nothing” alternative), Sloat 3 (purchase a third radar of the type already installed), and SkyRay. The question is which of the three non-dominated alternatives in the efficient set is most cost-effective? The answer cannot be given until we solicit additional preferences from the decision maker.

A marginal reasoning solution may begin with Sloat 2 and ask the decision maker whether he is willing to




Figure 19: $W_E > W_C$




Figure 20: $W_E \gg W_C$

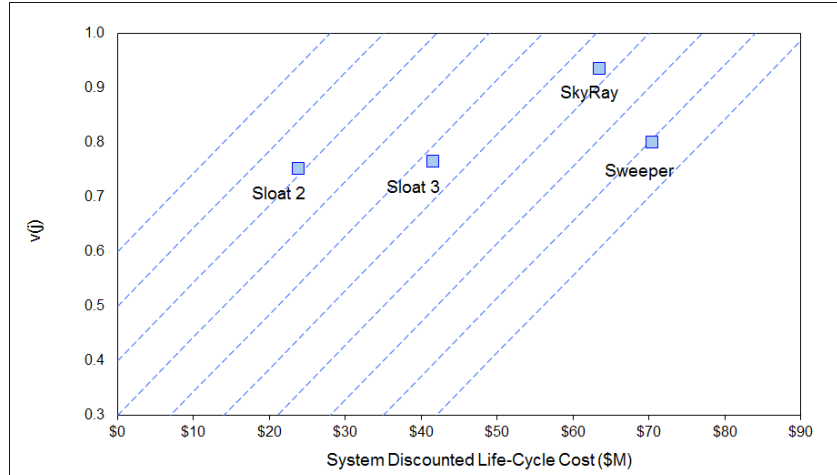


Figure 21: Cost-effectiveness: $W_C/W_E = 1.0$

spend an additional \$17.7 million to achieve an increase in 0.013 in effectiveness (the difference in MOEs between Sloat 3 and Sloat 2). The increase in effectiveness occurs because Sloat 3 is available more frequently than Sloat 2. If the decision maker responds that he is unwilling to spend the extra money for a small increase in effectiveness, we can ask him if he is willing to spend an additional \$39.6 million to achieve an increase in 0.182 in effectiveness (the difference in MOEs between SkyRay and Sloat 2). SkyRay is more effective than Sloat 2 because it has more range, is available more often, and has greater interoperability. Framing the questions in this manner can help the decision maker understand precisely what additional capability he is getting for an increased cost.

The more formal manner of answering which alternative is the most cost-effective requires elicitation of the final set of importance weights, W_C and W_E . Suppose the decision maker believes cost and effectiveness are equally important. These preferences translates to $W_C/W_E = 1.0$ resulting in the picture given in Figure 21. The decision maker would select the “do nothing” alternative because at least one isoquant line lies to the right of Sloat 2 and to the left of Sloat 3 and SkyRay.

If the decision maker believes maximizing effectiveness is twice as important as minimizing cost, $W_C/W_E = 0.5$, which results in the picture given in Figure 22. Sloat 2 is still most cost-effective because at least one isoquant line separates Sloat 2 from Sloat 3 and SkyRay. If the decision maker believes maximizing effectiveness is four times more important than cost reduction, $W_C/W_E = 0.25$. These preferences give the picture shown in Figure 23. SkyRay is the most cost-effective because at least one isoquant line lies below SkyRay and above Sloat 2 and Sloat 3.

This method also can serve to find the ratio between W_C and W_E which would make the decision maker indifferent between two alternatives. For example, when $W_C/W_E = 1/3$ we obtain the situation in Figure 24. Here Sloat 2 and SkyRay are almost of equal cost-effectiveness. The analysis of cost-effectiveness afforded by the model reduces the management question to one of asking the decision maker: “Do you feel that effectiveness is more than three times as important as cost?” If the answer is yes, SkyRay should be selected. If the answer is no then do nothing (Sloat 2) is the best alternative. The model helps to focus attention on critical information.

Now that we have found the critical value of $W_C/W_E \simeq 1/3$, it is of interest to consider the effects of uncertainty in the future condition. Two other planning scenarios besides the baseline were considered during the MOE discussion. If the supersonic attack scenario is highly likely, then we use different weights in the definition of effectiveness. The $v(j)$ values change for all alternatives and Figure 25 shows SkyRay is the most cost-effective if $W_C/W_E \leq 0.57$. The decision maker should be asked: “Is effectiveness more than 1.75 times as important than cost?” If he answers yes, SkyRay is the best alternative. The ECM planning scenario changes the importance weights in effectiveness, and the resulting cost-effectiveness is shown in Figure 26. If $W_C/W_E = 1/3$, Sweeper is the best alternative. Under this scenario, Sweeper remains the best alternative as long as the decision maker values effectiveness at least 1.4 times as much as cost, or $W_C/W_E \leq 0.74$.




Figure 22: Cost-effectiveness: $W_C/W_E = 0.5$

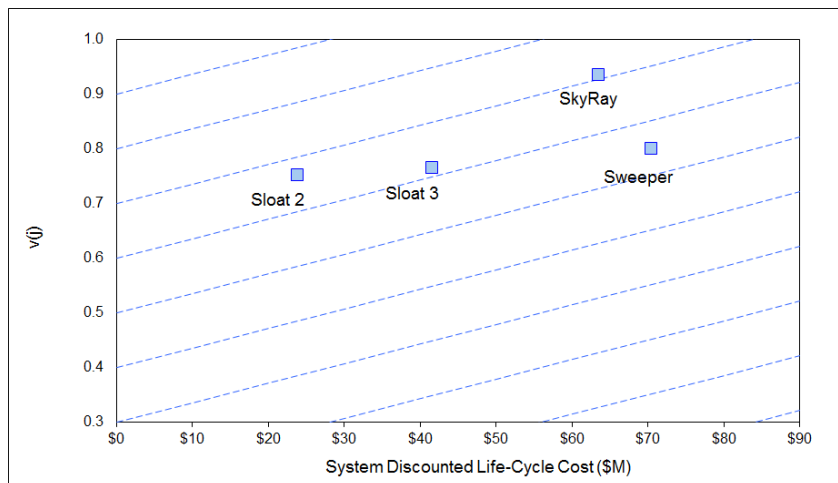


Figure 23: Cost-effectiveness: $W_C/W_E = 0.25$




Figure 24: Cost-effectiveness: $W_C/W_E = 0.33$




Figure 25: Cost-effectiveness for supersonic attack scenario: $W_C/W_E = 0.57$

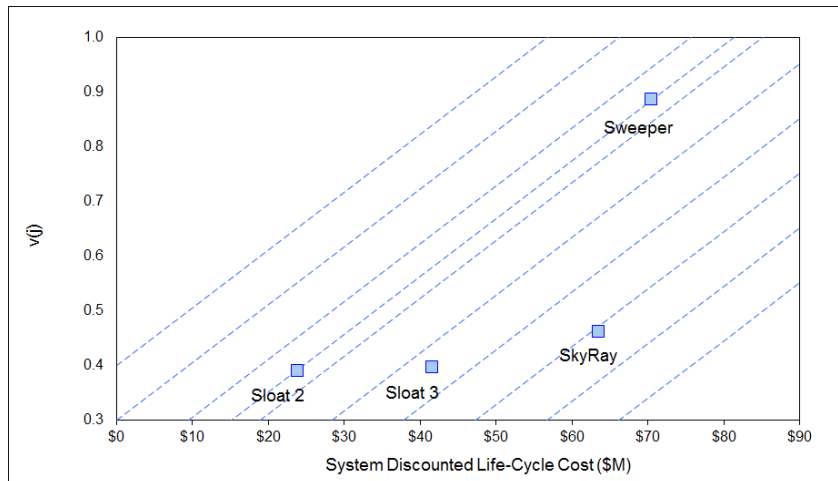


Figure 26: Cost-effectiveness for ECM scenario: $W_C/W_E = 0.74$

7 SUMMARY

Multiple objective decision problems arise very frequently. Their solution requires the decision maker to first determine what really matters and list the issues or consequences of concern. This process of discovery should be pursued through the construction of an objectives hierarchy. It is the single most important step towards a solution. Without doing this we run the risk of not knowing “what is the real problem” and of not asking “the right question.”

The lowest levels of the hierarchy define the individual measures of effectiveness. These are the natural, constructed, or proxy measurement scales by which we begin to quantify the effectiveness of an alternative. The proper integration of these measures into a single MOE requires quantification of decision maker preferences. First we need to know decision maker preferences over marginal changes in the individual effectiveness measures—the answer to “How much is enough?” These changes can represent increases (when more is better) or decreases (when less is better). This information allows us to convert the individual effectiveness measures into individual value measures on a 0 – 1 scale. Second we need to know decision maker preferences across the individual measures—the answer to “How important is it?” This information allows us to combine the individual scaling functions using a weighted sum defined by the importance weights.

Next, we combine the resulting MOE with the discounted life-cycle cost. This provides all the information needed to conduct the analysis of cost-effectiveness. Viewing things in a two-dimensional framework expedites thinking about the solution by allowing visual inspection and making use of humans’ ability at pattern recognition. In this framework, five solution concepts apply: (1) the superior solution, (2) the efficient solution, (3) the satisficing solution, (4) marginal reasoning, and (5) weighting cost versus effectiveness. The first concept should always be sought, and a superior solution (if one exists) is the best alternative regardless of the preferences for cost reduction versus effectiveness maximization. The efficient solution and the satisficing solution often yield non-unique answers and gives the decision maker flexibility over which alternative to select. The set of efficient solutions is intrinsically important because because the most cost-effective alternative will come from this set. Selecting the most cost-effective alternative requires eliciting additional decision maker preferences. We must know the decision maker’s preferences over cost minimization and effectiveness maximization before determining the most cost-effective alternative.

APPENDIX: Individual Value Functions

There are many functions for describing decision maker preferences over marginal changes in a single effectiveness measure. The exponential function provides a good approximation to many preferences:

$$v_i(z_i) = \frac{1 - e^{-az_i}}{K}$$

where either

$$z_i = x_i - x_{\min}$$

(the more-is-better case) or

$$z_i = x_{\max} - x_i$$

(the less-is-better case) and $K = 1 - e^{-a[x_{\max} - x_{\min}]}$ is a constant that standardizes the range of variation so $0 \leq v_i(z_i) \leq 1$.

This function is flexible enough to include the linear value function. It can be shown using calculus that, in the limit as $a \rightarrow 0$, this function becomes

$$v_i(z_i) = \frac{z_i}{x_{\max} - x_{\min}}.$$

This describes a decision maker who places equal value on marginal changes of equal amounts. For example, let $a = 0.001$, $x_{\min} = 0$ and $x_{\max} = \$90$ million. We obtain the function pictured in Figure 27. This represents a decision maker who values a \$10 million cost savings the same, no matter if it is a reduction in cost from \$90 to \$80 million or from \$70 to \$60 million.

This exponential form can be expanded to include a quadratic term,

$$v_i(z_i) = \frac{1 - e^{-az_i - b_1 z_i^2}}{K_1},$$

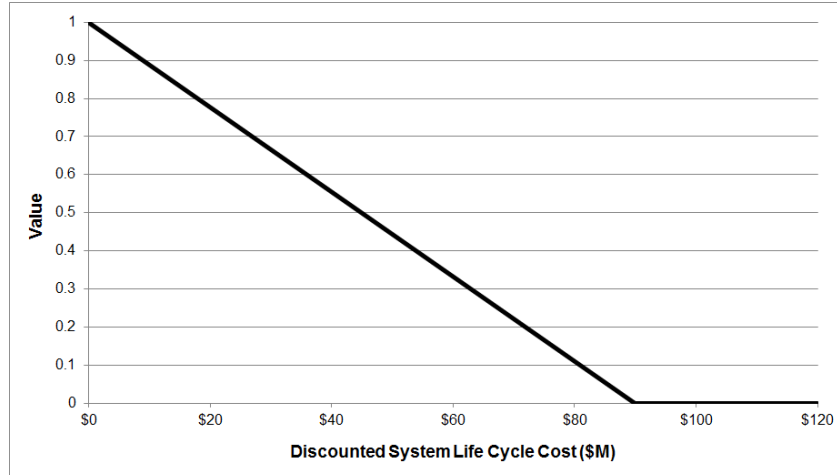


Figure 27: Linear value function for cost

or a cubic term,

$$v_i(z_i) = \frac{1 - e^{-az_i - b_1 z_i^2 - b_2 z_i^3}}{K_2},$$

where $K_1 = 1 - e^{-a[x_{\max} - x_{\min}] - b_1[x_{\max} - x_{\min}]^2}$ and $K_2 = 1 - e^{-a[x_{\max} - x_{\min}] - b_1[x_{\max} - x_{\min}]^2 - b_2[x_{\max} - x_{\min}]^3}$. The quadratic form with $a = 0$ allows value functions with an “S-shape” to be represented. The cubic form with $a = b = 0$ permits us to incorporate value functions that appear almost like “step” functions.

Higher order terms can be specified but in practice are almost never needed. The additional terms permit a wider range of preference behavior to be modeled but are hardly ever worth the increased complexity.