C H A P T E R 6

ANALYSIS CONCEPTS: COMBINING CRITERIA

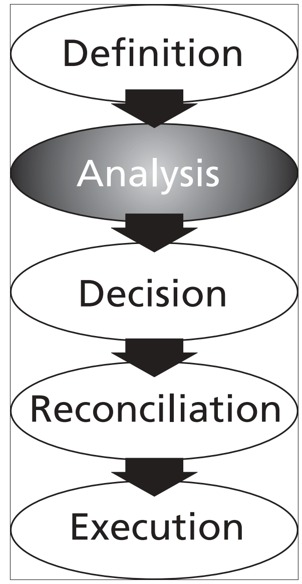

Science is built up with facts, as a house is with stones.
But a collection of facts is no more a science
than a heap of stones is a house.
-Jules Henri Poincaré: *La Science et l'Hypothèse,* 1908

NOW THAT WE HAVE DESCRIBED the important considerations about alternatives we wish to measure, we will take the next step and combine them to facilitate our comparisons. Methods for doing this range from simple to complex, depending on the nature of the problem. In this chapter, we will begin by making evaluations based on a single criterion and finish the chapter with more complex weighted models that incorporate as many criteria as our resources allow.

Combining Cost and Effectiveness

The heart of the Analysis Phase is the evaluation and comparison of alternatives. We confront several challenges combining cost and effectiveness that vary with the individual decision. For example, DoD recognizes the tension between cost and effectiveness and deliberately reflects that tension in its organization. The Combatant Commanders (CinCs) focus on warfighting capabilities (effectiveness) while the services must balance those requirements against programs and overall force structure (cost).

We know that the least costly and most effective alternatives are rarely the same, and each is attractive to different organizations. The least costly options are naturally appealing to our political leaders and headquarters staffs who are charged to conserve the public treasure. They rightly seek to fund as many programs as possible to support all the services' needs. On the other hand, military operators are biased toward the most effective system, regardless of cost, as they seek to win battles quickly, with the fewest casualties. However, an overly operational preference in programming can lead to excess capability (gold-plating) that drains resources from other urgent programs. Executive decision makers, then, are forced to make tradeoffs between cost and effectiveness. We strive to present those tradeoffs in structured terms to make these choices clear.

We cannot always compare the alternatives fairly using the same measures of cost, especially when procurement options are in different stages of completion. For instance, we may be comparing three missile alternatives to counter a new threat: (1) upgrading a missile already in the inventory; (2) procuring a prototype missile undergoing flight tests, for which research and development is complete; and (3) funding a missile proposal on the drawing board. Comparing the costs of these missiles equitably is difficult. The first missile will have by far the lowest unit cost in terms of procurement; however, an upgraded missile will not last as long as a new missile and it has a very limited potential for growth. The missile on the drawing board will have the longest life-cycle cost measured from the present. How shall we identify the time horizon for calculating life cycle costs for this procurement decision? Use the shortest? Or the longest? Instead, should we use the procurement price of each missile in constant dollars? Do we really care if costs are compared equitably? How we frame the measures of cost may determine the outcome of the analysis before we make any further comparisons.

We encounter a similar problem with comparisons based on effectiveness, especially when multi-mission capabilities are at issue. (Rarely do our alternatives achieve the same levels of effectiveness for each criteria.) Take the comparison of the effectiveness of the new C-17 transport aircraft with the C-5, C-141, and C-130. The C-5 and C-141 are used for inter-theater lift; the C-130's mission supports intra-theater lift. The C-17 was designed to do both. The measures of effectiveness the Air Force uses to compare the aircraft must include both airlift missions, and the alternative fleets must each accomplish both missions. However, effectiveness alone is not the issue; we must take cost of alternative fleets in consideration as well.

Fixing Cost or Effectiveness

One technique for combining cost and effectiveness is to fix cost or to fix effectiveness.[1] When we fix one of them, we compare our alternatives on the basis of the other. Thus, when we select a specific type of cost as the single type (or measure) of cost that we will compare our alternatives against, we are fixing cost. Of course, our challenge when fixing cost or effectiveness is to ensure the alternatives are truly equal in performance or value in the area we fix. We have a myriad of ways to measure cost. If we select just one, then we ignore the others which may vary significantly, e.g., if we fix life cycle cost, we only measure that type of cost; if we fix near-term cost, we eliminate downstream operations and maintenance costs from consideration.

The same problem occurs when we fix effectiveness; we have to select one (and only one) measure of effectiveness and then select the lowest cost option. For example, what single measure is the best way to measure the effectiveness of a tactical combat aircraft? In 1996, in their evaluation of the affordability of the Defense Department's tactical aviation procurement plan, the Congressional Budget Office selected aircraft age, represented by the proxy measure Technology Generation, and thereby fixed the effectiveness of the F/A-18E/F Super Hornet, the F-22 Raptor, and the Joint Strike Fighter at the same benchmark. When DoD identifies a set of specifications for a system and sends out a Request For Proposal to contractors, it is in effect fixing effectiveness and preparing to select the alternative with the lowest cost.

Cost and effectiveness are inextricably related to one another; however, the relationship is rarely the same for each alternative. For example, contractors may include different support packages in their proposals. Their unit price may vary with the quantity procured, creating step functions in the unit cost profile as they open another production facility. In such cases, we may

1. If cost and effectiveness are both equal, then the decision maker has no economic or performance basis for making a choice and can select any of the alternatives that meet minimum requirements.

not be able to fix cost or effectiveness easily. When we cannot fix cost or effectiveness, we might combine them to help us choose between the alternatives.

Cost and Effectiveness Ratios

One straightforward method for combining cost and effectiveness involves constructing a ratio. To do so, we must isolate one measure of cost and one measure of effectiveness that each represents our range of alternatives. By selecting one and only one MOC, we are effectively saying that the other costs are irrelevant to our decision; likewise, as we select an MOE we establish that there is an MOE that dominates all the others in importance.

The principal advantage of the cost-effectiveness ratio is its simplicity; this is also its principal disadvantage. Again, to use ratios, we must be confident the single MOC and single MOE we select truly dominate cost and effectiveness. We can compare alternatives easily and quickly using ratios displayed as a simple number (like cost per square foot), on graphs, or as tables. Again, this technique breaks down quickly as the complexity of the problem increases and a single MOC or MOE no longer dominates.

Weighted Models

One of the most powerful ways we can combine multiple criteria is by using weighted models. Basically, we arrange our criteria into a hierarchical model and evaluate the importance of each measure to the decision maker. We then assign weights to each criterion that reflect that relevance. After we build the overall hierarchical model, we evaluate a range of possible scores for each criterion using a utility scale. The utility scale is based on how much value we place on an alternative's improved performance (beyond the minimum requirement) for that criterion. Finally, we evaluate each alternative using the complete model, multiplying each utility score times the weight for that criterion then summing to create a total score for each alternative. We will describe each step in further detail below.

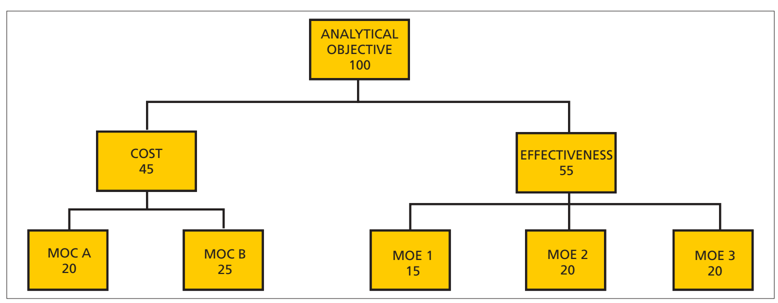

Figure 6-1. Weighted Model for Cost and Effectiveness.

Figure 6-1 is a typical weighted model with two primary categories of criteria, cost and effectiveness. The weights of a model may add up to 100, they may be normalized to total 1.0, or, if the model is built from the bottom up, they may combine to whatever number is the sum of the weights. Notice in figure 6-1 how the full weight descends from the analytic objective into successive tiers: first, Cost and Effectiveness, and then down to each of their subordinate measures. The weights of a subordinate tier must equal the weight of the parent tier above it.

Assigning the weights is a subjective process. It may be a group process, but it is the culmination of the experience of the people making the model. The manner in which we award weights may determine the outcome of the analysis, so we must have them approved by the decision maker and we may expect to have them challenged during the Reconciliation Phase. Thus, at this juncture, we are blending the mechanics of analysis with some very important professional judgments to make rational decisions.

UTILITY

After we design the structure of the weighted model and assign the weights to each criterion, we must decide how to assess our alternatives based on the lower-most criteria. This requires us to apply utility to the range of measures which defines our criteria. For example, if speed is one criterion, we must establish the range of speed for which we are interested and decide how important each incremental change is. We do this by constructing utility curves.

Utility is simply a way to show the decision maker's value-gained or lost by changes in the measurement of our criterion. Then, depending on where our alternative falls on the utility scale, we compare it with others. Conversion of all measures of cost and effectiveness to utility allows us to multiply values by the weights of criteria and convert numerous measures to one aggregate number for each alternative.

Utility must be dimensionless because we are going to add utility scores together. We derive utility scores from measurements and assessments that often do have dimensions, e.g., dollars, speed, range, man-days. To make the conversion from the direct assessment to utility scores to use in the model, we create utility curves that show value as a function of performance, cost, and sometimes risk.

As you might imagine, the key to this process is the care with which each alternative's performance on each measure is converted into the common currency of utility. Using a fighter aircraft example, suppose we wish to compare two alternatives. Aircraft **A** has a top speed of Mach 1; aircraft **B** has a top speed of Mach 2. If we determine that Mach 2 has twice the utility of Mach 1 for a fighter, the utility score of aircraft **B**'s speed should be twice that of aircraft **A**. If Mach 2 is only 10 percent more useful, then aircraft **B**'s speed-utility score should be 10 percent greater than that of aircraft **A**. You can see that by mapping speed in regard to its utility, you can generate a utility curve—a picture of the utility of each speed.

We can do the same for every criterion; every measurement or assessment falls along a utility curve that expresses its usefulness or value to our organization and the decision maker. The very act of identifying utility benefits our organization because the participants must have (or gain) a thorough and universal understanding of our core values and missions as we establish the utility curve of value versus performance. We can evaluate utility in a variety of valuable and creative manners, numerical and otherwise.

UTILITY DISPLAYS

Graphs are the most common method for displaying utility. A graph is a simple model that shows pictorially the relationship between two sets of numbers. Since we are translating the measures of our criteria into utility, graphs are especially powerful for converting performance and cost data into utility by showing how changes in performance or cost relates to value or usefulness.

To build a utility curve, first we establish a range of values for the criterion and identify the range of likely values for the set of alternatives. If we are selecting among armored fighting vehicles, one of our criteria may be Maximum Off-Road Speed. We set a minimum requirement (threshold) of 30 MPH to ensure the new vehicle can keep up with our current armored vehicles and those that we project will be in the inventory during the new vehicle's service life. We would like a top speed (objective) of 60 MPH to enable it to dash from point to point. These are the endpoints for our range of values for speed: 30-60 MPH that we will put on the horizontal axis as shown in figure 6-2.

Next, we decide how much resolution we want to distinguish performance levels among the alternatives: our options include every MPH difference (or even less), 5 MPH blocks, 10 MPH blocks, etc. After discussion within our organization and with the analyst, and with the approval of the decision maker, we opt for 5 MPH increments because they are significant enough to differentiate between alternatives without being too detailed. These blocks are often called bins. Each bin, in this example, has five values that we score equally, i.e., to us, there is no more value in a vehicle that travels at 45 MPH than there is for one that travels at 41 MPH.[2] This is another crucial intersection of experience and analysis. Our experience, and that of the operators, subjectively determines the utility of different levels of performance, just as it did when we established weights.

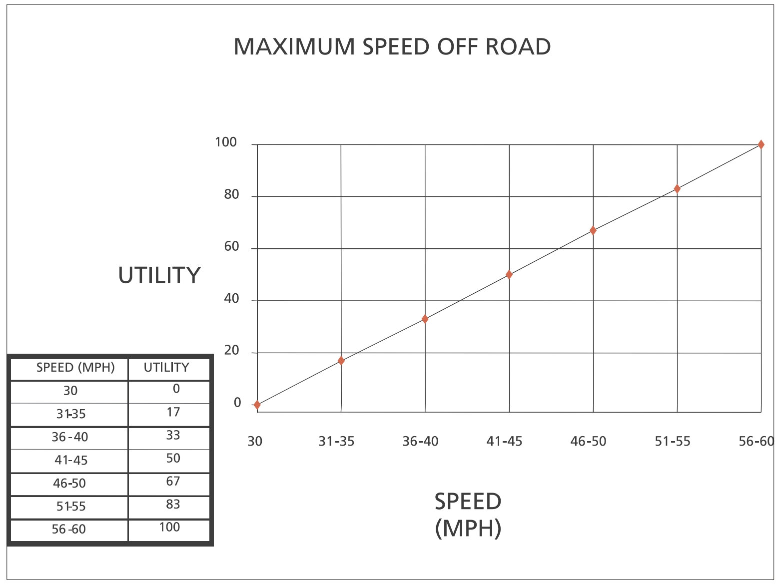

Figure 6-2. Straight-Line Utility Curve.

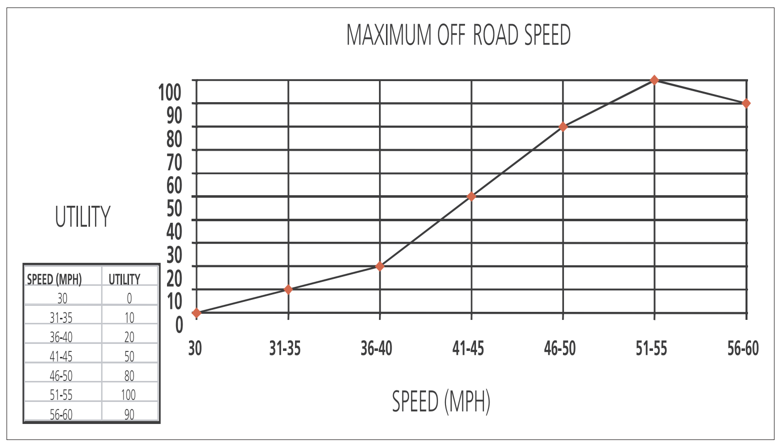

Figure 6-3. Curvilinear Utility Function.

For the vertical axis, we select an arbitrary scale of Utiles (unit-less measures of utility) from 0 to 100. We now build a data table by asking ourselves—or others, such as operators in the field—how much they value speed above 30 MPH, in 5 MPH increments as shown in the lower left corner of figure 6-3. If every increase in speed has the same value as the previous increment and the next, we will get a straight-line utility function as we show in figure 6-2.

Linear utility curves are good representations for many of the criteria we examine. They accurately describe cost profiles in which one additional dollar is worth the same to the decision maker as was the last dollar. There are, however, many circumstances that lead us to complicate our curves to better reflect reality.

Continuing the example above, we will probably find while discussing off-road speed with the operators that one 5 MPH increment is *not* as good as the next. For tactical vehicles, a maximum speed of 30 MPH is acceptable but not very desirable. Most troops desire speeds of at least 40 MPH, and many find speeds above 50 MPH significantly preferable. Towards the upper range of speeds, however, there is a leveling off and then a decrease in utility because operators never really need speeds over 55 MPH—and speeds beyond 55 MPH increase accident rates (according to our sampling). Revising our data table, we generate the curve in figure 6-3. Notice

2. If we are using a computer-based system, we could use the formula for the line instead of making bins. In the case in figure 6-3, Utility = 3.33 * speed (MPH) - 100; any speeds below 30 MPH result in negative numbers, another way of saying that alternative does not meet our minimum requirements.

that the bins remained in five MPH increments, but their utility does not change as evenly as before in the data table. The shape of the curve reflects our values; some incremental improvements are worth more than others.

Both the curves in figures 6-2 and 6-3 are smooth and continuous (without breaks). Sometimes we have utility curves that have abrupt changes in value (discontinuities) that create breaks in functions. This happens when a change in the criterion's value influences another criterion or triggers an event, such as when we receive a price break for buying in larger quantities.

Sometimes cost exhibits this behavior, e.g., if we budgeted a certain amount of money for a project, the utility of cost will decline smoothly as cost rises—lower cost has higher utility so our curve has a negative slope—until we reach our budget limit, marked by the star in figure 6-4. If we choose an alternative that exceeds the budget, the decision maker will have to obtain more funds, a distasteful but not impossible proposition. The utility of cost becomes linear again, once we cross that increased budget point, because each one-dollar is once more worth the same as the next. It is the act of having to get more money that creates the huge drop in utility; once we pass this hurdle, we return to the original slope where each one-dollar increase is worth the same utility as the next.

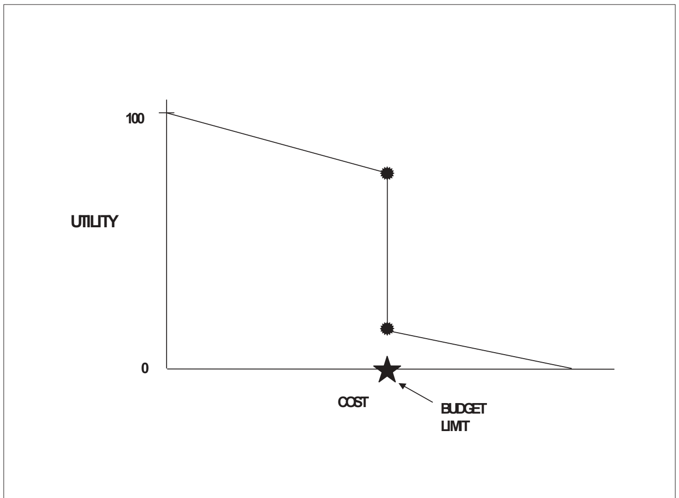

Figure 6-4. Discontinuous Utility Curve.

In the preceding examples, we have used quantifiable data for making our value assessments. While the utility of a particular value may be debated between decision makers, we can usually agree on a general function that describes rational decision makers' preferences. With qualitative criteria, however, we find ourselves marrying subjective traits (the value to the decision maker) and value judgments (the subjective data) on both axes of the utility curve.

We have several other options for compiling and displaying qualitative data. We can use a carefully defined descriptive scale, like the common traffic light scheme of green, yellow, and red for good, fair, and bad (see the next box). We may use qualitative terms from Very Low to Very High with any number of gradations in between. We can display our assessments directly in these same terms or, if we are going to incorporate them into a weighted model, we can convert them into utiles or numerical values, e.g., we convert Very High to 10 points; Medium to 5 points, and Very Low to 0 points.

ASSESSING OBJECTIVES: U.S. EUROPEAN COMMAND'S THEATER OBJECTIVES AND MEASURES

Figure 6-5 is the mechanism used by the U.S. European Command to display its quarterly progress towards achieving the Commander-In-Chief's Theater Objectives.[3] In 1995, then-CINCEUR General George A. Joulwan, U.S. Army, established nine theater objectives. He pronounced they were meaningless unless his staff could measure and report his command's progress toward achieving them.

3. Typically, USEUCOM has nine theater objectives such as Quality of Life for the U.S. military community, Strengthening NATO, Supporting Middle East peace initiatives, etc. The fishbone mechanism is based on the cause and effect diagrams from Michael Brassard's *The Memory Jogger Plus* published by Goal/QPC, Methuen, Mass., rev. 1996.

Basically, in conjunction with the rest of its Theater Security Planning System, CINCEUR's headquarters staff consolidates inputs, called Indicators, from its components and subordinate commands to assess each criterion, in this case measures of effectiveness. The indicators are often numerical, e.g., the number of exercises scheduled compared to the number executed, and the number of exercises this year compared to last year. A decision maker (a Division Chief at the O-6 level) evaluates each MOE in terms of Satisfactory (green), Concern (yellow), or Unsatisfactory (red).[4] Note in this case the colors (rather than numbers) express utility. Translating the individual indicators into the common currency of utility enables the decision maker to easily understand the meaning of the indicators.

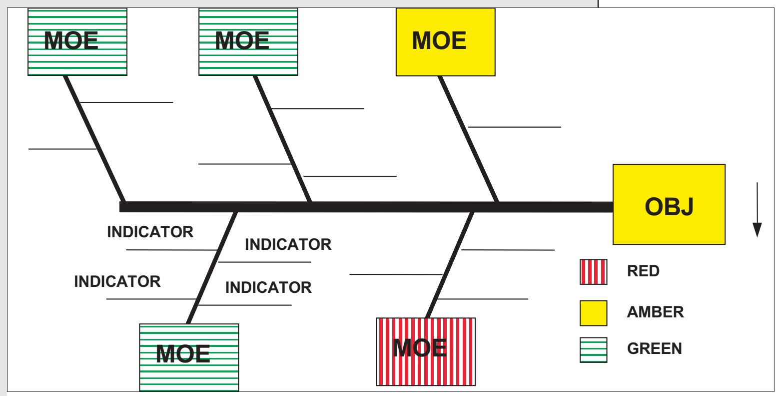

Figure 6-5: USEUCOM Red/Yellow/Green Measurement Chart

The Director (an O-8 decision maker) responsible for each objective makes an over-all assessment of the objective, assigns a color, and indicates an upward, downward, or neutral trend. The staff briefs the status charts to the CINC and his Component Commanders quarterly. Naturally, most of the discussion revolves around the areas that are not green; the briefing books include, behind each fishbone display, the MOE assessments and indicator data that support each rating.

After we have built utility curves for each criterion, we are ready to display them with our model. Recall that each criterion that ends a branch of the weighted model must be associated with a utility curve.

Figure 6-6 combines a weighted model with its utility curves. Note some of the characteristics of the curves. Where cost is concerned, one dollar is worth the same as the next to us for both MOCs so we have the highest score (maximum utility) for the lowest possible cost and follow a straight line down to zero for the highest affordable cost. The utility curves under the MOEs illustrate a variety of shapes that reflect the varied nature of the particular measures. In each case, we decide how much value we associate with exceeding the minimum requirement for that measure and whether one incremental change has the same value as the next. Thus the first MOE uses a curve, the second is a step function (yes or no), and the third uses bins to convert performance into utility.

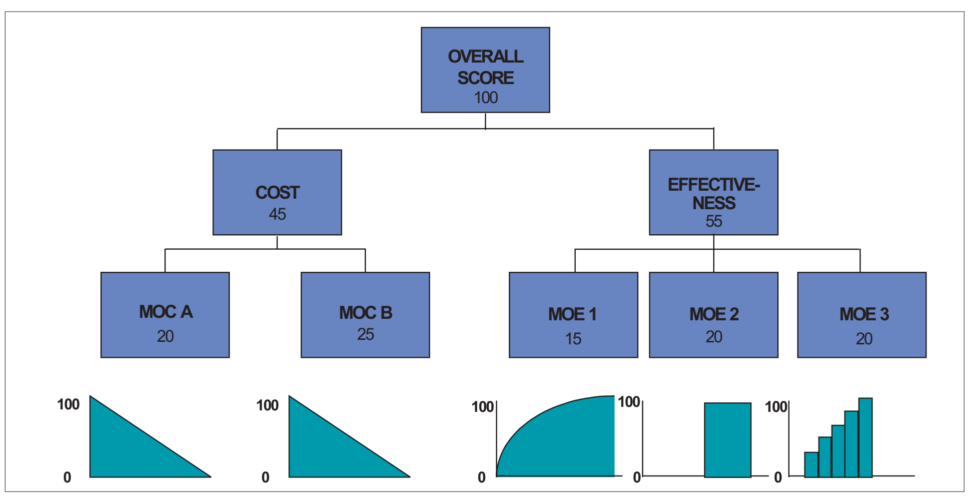

Figure 6-6: Model with Weights and Utility Curves.

4. It is a real challenge to convince organizations that reporting bad news about actions beyond their control (a red rating) will not be held against them, e.g., if the Arabs and Israelis refuse to negotiate with each other, we do not blame the action officers in J-5 Plans and Policy!

EVALUATING ALTERNATIVES USING A WEIGHTED MODEL

After we assign weights and decide upon the utility curve for each criterion, we evaluate the alternatives using the model, one criterion at a time. We arrive at a score for each alternative for each criterion; the score is its utility, which we multiply by the weight of the criterion as we show below.

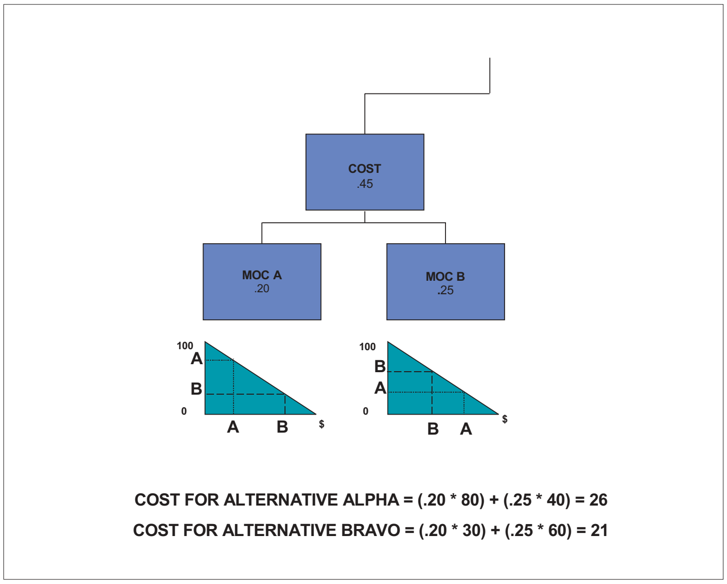

Figure 6-7: Calculating an Alternative's Score for Two Criteria.

To calculate the score, first we find its performance measurement or our assessment along the horizontal axis, and then we read upward until we intercept the utility curve. At the intersection, we read across to the left to find the utility value, as we show in figure 6-7, e.g., alternative Alpha has a utility value of 80 for MOC A. Then we multiply the utility value (80) times the weight (.20) to arrive at the score for this criterion: 16 for alternative Alpha for MOC A. In the figure, alternative Alpha scores 26 of 45 possible points for Cost while Alternative Bravo scores 21. We repeat this process for every measure at the end of every branch and then sum the scores to arrive at a total score for each alternative. In reality, the analyst will do most of this stubby pencil drill, but whether he or she uses a pencil or a Cray computer, the underlying principles are the same.

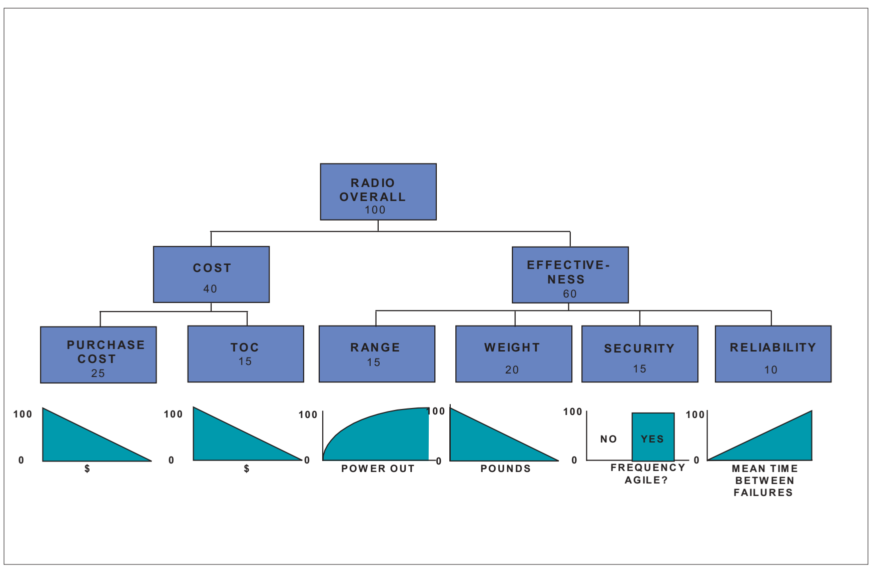

Figure 6-8: Weighted Model for a Portable Radio.

In figure 6-8, we have a weighted model and utility curves for a hypothetical portable radio requirement. The ideal alternative in a model like this would score 100; the worst (but still acceptable) would score 0; the vast majority of alternatives fall in between. The output from the model should allow us to identify preference and provide us with insights on how the alternatives compare and complement each other. These insights should include the strengths and weaknesses of each alternative and where the tradeoffs lie in selecting between options. After running the model, we should be able to see how we can modify a lesser-ranked alternative to improve its value to the decision maker. The results should also indicate where and how we could form hybrid alternatives to optimize the strengths of different options.

RADIO	PRICE	TOC	RANGE	WEIGHT	SECURITY	RELIABILITY	TOTAL
POPIEL 1995	23	15	0	20	0	0	58
WHAMMO 3000	11	7	12	16	15	6	67
ZONKER 101	4	14	15	6	15	10	64

Table 6-1. Radio Alternatives.

In table 6-1, we display the results of applying the model in figure 6-8 to three alternatives. Based on this model, the Whammo 3000 is the best choice—assuming that our weights correctly reflect the decision maker's perspective about what is important and that the utility curves reflect how we value differences in performance. Note, too, that the Whammo 3000 did not score highest for any single criteria.

Combining Risk and Uncertainty with Cost and Effectiveness

While combining cost and effectiveness is usually a straightforward process that becomes obvious from the Definition Phase and its analytic objectives, incorporating risk offers several options. As we discussed in Chapter 5, we can combine risk and cost by buying out risk and we can combine risk with performance by using expected values, multiplying probability times outcome. We can build our risk assessments into our models, or we may choose to evaluate risk separately after we evaluate our alternatives in terms of cost and effectiveness. The importance of including a consideration of risk and uncertainty is that it informs the Decision Maker as to what we don't know, as well as what we do know. Further, it helps him/her understand the negative consequences of each alternative. This more complete information package is what he/she needs to make the best informed decision possible.

CASE STUDY: THE ANALYSIS PHASE - COMBINING CRITERIA
USMC MEDIUM-LIFT REQUIREMENTS: THE V-22 OSPREY AND HELICOPTERS

The Institute for Defense Analyses used a fixed cost approach to conduct their V-22 and helicopter comparison. As we mentioned at the end of Chapters 3 and 4, they fixed cost at two levels in FY88 Dollars. First, at $33B, the cost they estimated for the Marine Corps desired fleet that would lift half of their assault force in the first wave, and, second at $24B that was the funding DoD was prepared to allocate for replacement helicopters. We reproduce the study's Table 4 again for easy reference:

Marine Corps Medium-Lift Assault Aircraft	Number at Cost Level I ($33B FY88)	Number at Cost Level II ($24B FY88)
V-22	502	356
New Helicopter	634	450
CH-47M	673	527
CH-60 (S)/CH-53E+	287/347	240/283
CH-46E+/CH-53E+	317/336	251/258
Puma/CH-53E+	330/322	260/246
EH-101/CH-53E+	252/335	200/256

Was this an appropriate methodology for this study? Cost was certainly a dominant aspect of this problem: indeed Secretary of Defense Cheney's overriding concern was the V-22's near-term cost in the face of many competing DoD programs. This study does a good job of demonstrating the value DoD received in return for its dollars, i.e., IDA can show, based on its assumptions (including those surrounding scenarios) and measures of effectiveness, that the V-22 generates more medium-lift per dollar. That is important and useful information. The only significant weakness of fixing cost in this manner is it removes the possibility that a less expensive helicopter fleet could accomplish the mission at a lower cost than $24B, however, that is unlikely because we can assume DoD explored that possibility when they identified their medium-lift proposal.

But what if IDA had fixed effectiveness instead? Naturally, they would have to focus on the assault scenario and the Marines would have to define a required build-up rate of combat power for the air assault some distance from the amphibious ships. The first challenge would be to identify a scenario; at this time, there was none on the shelf. The Marines were embracing Over-The-Horizon assault and just beginning to explore the ramifications of avoiding the build-up on the beach prior to moving to the objective, the sequence central to their earlier doctrine. Developing a new scenario, in the middle of this contentious medium-lift issue, was fraught with problems. The Marines (and IDA) would almost certainly have been accused of bias by V-22 skeptics whether justified or not—fixing effectiveness would expand the controversy rather than reduce it, as this study was intended to do, because stakeholders would argue about the scenario before they even got to the analysis.

There was no worthwhile cost-effectiveness ratio to create in this situation because there were two cost issues—near-term and life cycle cost—and Congress specified multiple effectiveness issues. IDA could have created a weighted model, but that may have added more complexity than was desirable given their target audience of DoD leadership and Congress. A weighted model would have been an amalgamation of what we already see: the two types of cost on one side and the eight scenario evaluations on the other. Assigning the weights would have been very contentious: imagine trying to reach agreement of the relative importance of cost and effectiveness. Given the circumstances, we think IDA chose the most appropriate methodology available.

Risk was incorporated in a very limited manner by IDA in its effectiveness evaluations and largely ignored as they considered cost. As events have shown, this was an important, albeit deliberate, omission. In the logistics sustainment scenario in particular, the effectiveness of the aircraft was based on mechanical reliability. To reduce weight and conserve space, the V-22 uses a flight control hydraulic system that produces 5,000 psi system pressure, more than double the pressure of most helicopters. Although IDA halved the time between failures the V-22 manufacturer projected, the doubled failure rate was still well below that of helicopters. With a new technology aircraft like the V-22, there was no way to create objective probabilities, however IDA's subjective probability was over-optimistic.

There was also a great deal of risk in the schedule and cost projections for the V-22 that fell outside IDA's study. One can argue whether it was their place to challenge Bell-Boeing's cost and schedule figures, but there was little doubt even in 1990 that they were optimistic. Essentially, they left those risks for the decision makers to tackle independently, probably for reasons of practicality. IDA was under severe time pressure to complete this study and by making assumptions and limiting it to what was measurable, they simplified its completion.

Summary

The methods we choose for combining effectiveness, cost, and risk criteria depend upon on the nature of the problem. We desire to make our comparisons as simple as possible by reducing the number of variables, which is most practical when we can fix either cost or effectiveness. If neither can be fixed but both have a dominating measure we can establish a cost-effectiveness ratio, either as a single number or displayed on a graph.

For more complex situations, we use weighted models that allow us to use as many criteria as time and money permit. As we construct these models, we in DoD establish the weights often without detailed foreknowledge of the alternatives. This leads to a more general model, but we may find, when we apply it to the alternatives, that some of our criteria are not very helpful discriminating between options because they all score equally in that area. In other circumstances, we may know the alternatives before we build the weighted model. This permits us to choose criteria we know will highlight differences, but the model may require revision if new alternatives are added that vary from the existing options in a new way.

For each criterion, we build a utility curve that we will use to translate measures and assessments into common, dimensionless units that reflect value or usefulness. We evaluate each alternative for each criterion using the utility curves. After multiplying each score by its weight, we sum them to reach a total score for each alternative. The data that we use to evaluate the alternatives, and to establish the utility scores, is largely beyond our control in the sense that it should represent the truth about each alternative, either by objective measure or careful subjective assessment. The results of the model can be affected by changing either the weights of the criteria, the shapes of the utility curves, or the values of the alternatives themselves. Because there are many subjective evaluations built into any weighted model, we will insist on knowing them when the analysts provide us the model.

We may incorporate risk with cost and effectiveness or we can study them in isolation after we have evaluated cost and effectiveness. However we decide to address them, we should make our intentions clear to the decision maker.