CHAPTER 17

Cost–Benefit Analysis for Human Systems Integration

WILLIAM B. ROUSE and KENNETH R. BOFF

17.1 INTRODUCTION

The past decade has been a period of very serious scrutiny of the activities of most enterprises. Business processes have been reengineered and enterprises have been downsized or, more popularly, rightsized. Every aspect of an enterprise must now provide value to customers, earn revenues based on this value, and pay its share of costs. Aspects of an enterprise that do not satisfy these criteria are targeted for elimination.

This philosophy seems quite reasonable and straightforward. However, implementation of this philosophy becomes rather difficult when the "value" provided is indirect and abstract. When anticipated benefits are not readily measurable in monetary units and only indirectly affect things amenable to monetary measurement, it can be very difficult to assess the worth of investing in such benefits.

There are a wealth of examples of such situations. With any reasonable annual discount rate, the tangible discounted cash flow of benefits from investments in libraries and education, for example, would be so small as to make it very difficult to justify societal investments in these institutions and activities. Of course, we feel quite justified arguing for such investments. Thus there obviously must be more involved in such an analysis than just discounted cash flow.

This chapter addresses types of human systems integration (HSI) investments that have these intangible characteristics in addition to more tangible attributes. One type is research and development (R&D). This type of investment is often made for the purpose of creating long-term value. It will certainly require years and may take decades before returns are fully realized. It is easy to see how R&D can be difficult to justify in terms of impacts on, for instance, this year's sales and profits or current operational readiness (Rouse and Boff, 1998).

Handbook of Human Systems Integration, Edited by Harold R. Booher.
ISBN 0-471-02053-2

Another type of investment with these intangible characteristics involves products and services that enhance human effectiveness and thereby HSI. This includes selection, training, system design, job design, organizational development, health and safety, and, in general, the wide range of things done to ensure and enhance the effectiveness of people in organizations ranging from businesses to military units. In particular, investments focused on increasing human potential, rather than direct job performance outputs, are much more difficult to justify than those with near-term financial returns (Rouse et al., 1997).

This chapter also addresses the complex interaction of R&D investments in human effectiveness as a central element of HSI. This is done by building on previous efforts by the authors and many others that address the two elements of this interaction. Investing in R&D to enhance human effectiveness presents a confluence of difficulties related to representing and quantifying benefits as well as attributing costs. Nevertheless, there is a widely shared sense that such investments are socially and economically important. It is difficult, however, to justify particular projects on the basis of such perceptions.

A primary difficulty involves the trade-off between the relatively short-term payoffs of direct improvements in job performance and the inherently long-term benefits of R&D efforts aimed at enhancing human effectiveness and HSI. Short-term investments usually involve less uncertainty and fewer risks. In contrast, revolutionary high-payoff innovations usually emerge from much earlier R&D investments. Thus small, certain, near-term returns compete with large, uncertain, long-term, and potentially very substantial returns. The methodology presented in this chapter enables addressing both types of investments.

In general, several issues underlie the difficulties in justifying the aforementioned types of long-term investments. As just noted, a fundamental issue concerns the associated uncertainties. Not only are the magnitudes and timing of returns uncertain; the very nature and characteristics of returns are uncertain. With R&D investments, for instance, the eventual payoffs from investments are almost always greater for unanticipated applications than for the originally envisioned applications (Burke, 1996). Further, organizations that make the original investments are often unable to take advantage of the eventual returns from R&D (Christensen, 1997).

Another central issue relates to the preponderance of intangible outcomes for these types of investments. For example, investments in training may enhance leadership skills of managers or commanders. Investments in organizational development can improve the cohesiveness of "mental models" of management teams or command teams and enhance the shared nature of these models. However, it is difficult to fully capture such impacts in terms of tangible, "bottom-line" metrics.

It is important to differentiate between intangible outcomes and those that are tangible but difficult to translate into monetary benefits or costs. For example, an investment might decrease pollution, which is very tangible, but it may be difficult to translate this projected reduction to estimated economic gain. This is a mainstream issue in economics and not unique to cost–benefit analyses.

A further issue concerns cost–benefit analyses across multiple stakeholders. Most companies' stakeholders include customers, shareholders, employees, suppliers, communities, etc. Government agencies often have quite diverse sociopolitical constituencies who benefit or stand to lose benefits in a myriad of ways depending on investment decisions. For example, government-sponsored market research may be part of a regional economic development plan or may be part of a broader political agenda focused on creating jobs. In general, diverse constituencies are quite likely to attempt to influence decisions in a variety of ways. These situations raise many basic questions relative to the importance of benefits and costs for the different stakeholders.

Yet another issue concerns the difference between assessing and predicting costs and benefits. It is certainly valuable to know whether past investments were justified. However, it would be substantially more valuable to be able to predict whether anticipated investments will later provide benefits that justify the initial investments. Of course, limits of our abilities to predict outcomes are not unique to cost–benefit analysis.

The types of investment problems addressed in this chapter are rife with many uncertainties, intangibles, and stakeholders and associated unpredictability. These issues are explored in this chapter in the context of alternative frameworks for performing cost–benefit analyses. This leads to clear conclusions about how best to methodologically handle these types of investments. Application of the resulting methodology is then illustrated in the context of three investment problems involving technologies for aiding, training, and ensuring the health and safety of personnel in military systems.

17.2 COST–BENEFIT FRAMEWORKS

There are a variety of frameworks for scrutinizing and justifying investments:

- *Cost–Benefit Analysis* Methods for estimating and evaluating time sequences of costs and benefits associated with alternative courses of action.
- *Cost–Effectiveness Analysis* Methods for estimating and evaluating time sequences of costs and multiattribute benefits to ensure that the greatest benefits accrue for given costs.
- *Life-Cycle Costing* Methods for estimating and evaluating costs of acquisition, operation, and retirement of alternative solutions over their total cycles of life.
- *Affordability Analysis* Methods for estimating and evaluating life-cycle costs compared to expected acquisition, operations, and maintenance budgets over the total life cycle of an alternative investment.
- *Return-on-Investment Analysis* Methods for projecting the ratio, expressed as a percentage, of anticipated free cash flow to planned resource investments.

This chapter focuses on cost–benefit analysis in a broad sense that includes many aspects of the other approaches. For more traditional treatments of cost–benefit analysis, as well as worked examples, see Layard and Glaister (1994) and Gramlich (1997).

Cost–benefit analyses are very straightforward when one considers fixed monetary investments made now to earn a known future stream of monetary returns over some time period. Things get much more complicated, however, when investments occur over time, some of which may be discretionary and when returns are uncertain.

Further complications arise when one must consider multiple stakeholders' preferences regarding risks and rewards. Additional complexity is added when returns are indirect and intangible rather than purely monetary. These complications and complexity are more common than are situations where the straightforward cost–benefit analyses are applicable. This section discusses alternative frameworks for addressing cost–benefit analyses and compares these alternatives relative to their abilities to address the issues considered in Section 17.1.

17.2.1 Traditional Economic Analysis

The time value of money is the central concept in this traditional approach. Resources invested now are worth more than the same amounts gained later. This is due to the costs of investment capital, which must be paid, or foregone, while waiting for subsequent returns on the investment. The time value of money is represented by discounting the cash flows produced by the investment to reflect the interest that would, in effect at least, have to be paid on the capital borrowed to finance the investment.

Equations (1–3) summarize the basic calculations of the discounted cash flow model. Given projections of costs c_i, $i = 0, 1, \ldots, N$, and returns r_i, $i = 0, 1, \ldots, N$, the calculations of *net present value (NPV), internal rate of return (IRR)*, or *cost–benefit ratio (CBR)* are quite straightforward elements of financial management (Brigham and Gapenski, 1988). The only subtlety is choosing a *discount rate* (DR) to reflect the current value of future returns decreasing as the time until those returns will be realized increases.

$$\mathrm{NPV} = \sum_{i=0}^{N} \frac{r_i - c_i}{(1+\mathrm{DR})^i} \tag{1}$$

$$\mathrm{IRR} = \mathrm{DR} \quad \text{such that} \quad \sum_{i=0}^{N} \frac{r_i - c_i}{(1+\mathrm{DR})^i} = 0 \tag{2}$$

$$\mathrm{CBR} = \frac{\sum_{i=0}^{N} c_i/(1+\mathrm{DR})^i}{\sum_{i=0}^{N} r_i/(1+\mathrm{DR})^i} \tag{3}$$

It is quite possible for *DR* to change with time, possibly reflecting expected increases in interest rates in the future. Equations (1)–(3) must be modified appropriately for time-varying discount rates.

The metrics in Equations (1)–(3) are interpreted as follows:

- The NPV reflects the amount one should be willing to pay now for benefits received in the future. These future benefits are discounted by the interest paid now to receive these later benefits.
- In contrast, IRR is the value of DR if NPV is zero. This metric enables comparing alternative investments by forcing the NPV of each investment to zero. Note that this assumes a fixed interest rate and reinvestment of intermediate returns at the internal rate of return.
- The CBR simply reflects the discounted cash outflows divided by the discounted cash inflows, or benefits.

17.2.2 Multiattribute Utility Models

Cost–benefit calculations become more complicated when benefits are not readily transformable to economic terms. Benefits such as safety, quality of life, and aesthetic value are very difficult to translate into strictly monetary values. Multiattribute utility models provide a means for dealing with situations involving mixtures of economic and noneconomic attributes.

Let cost attribute i at time j be denoted by c_{ij}, $i = 1, 2, \ldots, L$ and $j = 0, 1, \ldots, N$, and benefit attribute i and time j be denoted by b_{ij}, $i = 1, 2, \ldots, M$ and $j = 0, 1, \ldots, N$. The values of these costs and benefits are transformed to common utility scales using $u(c_{ij})$ and $u(b_{ij})$. These utility functions serve as inputs to the overall utility calculation at time j, as shown in Equation (4) (Keeney and Raiffa, 1976):

$$U(\underline{c}_j, \underline{b}_j) = U[u(c_{1j}), u(c_{2j}), \ldots, u(c_{Lj}), u(b_{1j}), u(b_{2j}), \ldots, u(b_{Mj})] \tag{4}$$

which provides the basis for an overall calculation across time using

$$U(\underline{C}, \underline{B}) = U[U(\underline{c}_1, \underline{b}_1), U(\underline{c}_2, \underline{b}_2), \ldots, U(\underline{c}_N, \underline{b}_N)] \tag{5}$$

Note that the time value of benefits depicted in Equations (1)–(3) is included in Equations (4) and (5) by dealing with the time value of costs and returns explicitly and separately from uncertainty.

An alternative approach involves assessing utility functions for discounted costs and benefits, possibly discounted as represented in Equations (1)–(3). With this approach, streams of costs and benefits are collapsed across time before the values are transformed to utility scales. The validity of this simpler approach depends on the extent to which people's preferences for discounted costs and benefits reflect their true preferences.

The mappings from c_{ij} and b_{ij} to $u(c_{ij})$ and $u(b_{ij})$, respectively, enable dealing with the subjectivity of preferences for noneconomic benefits. In other words, utility theory enables one to quantify and compare things that are often perceived as difficult to objectify. Unfortunately, models based on utility theory do not always reflect the ways in which human decision making actually works.

Subjective expected utility (SEU) theory reflects these human tendencies. Thus to the extent that one accepts that perceptions are reality, one needs to consider the SEU point of view when one makes expected-utility calculations. In fact, one should consider making these calculations using both objective and subjective probabilities to gain an understanding of the sensitivity of the results to perceptual differences.

Once one admits the subjective, one needs to address the issue of whose perceptions are considered. Most decisions involve multiple stakeholders, in other words, people who hold a stake in the outcome of a decision. It is therefore common for multiple stakeholders to influence a decision. Consequently, the cost–benefit calculation needs to take into account multiple sets of preferences. The result is a group utility model as shown in Equation (6) (Keeney and Raiffa, 1976; Kirkwood, 1979):

$$U = U[U_1(\underline{C}, \underline{B}), U_2(\underline{C}, \underline{B}), \ldots, U_K(\underline{C}, \underline{B})] \tag{6}$$

where K is the number of stakeholders.

Formulation of such a model requires that two important issues be resolved. First, mappings from attributes to utilities must enable comparisons across stakeholders. In other words, one has to assume that $u = 0.8$, for example, implies the same value gained or lost for all stakeholders, although the mapping from attribute to utility may vary for each stakeholder. Thus all stakeholders may, for instance, have different needs or desires for safety and, hence, different utility functions. They also may have different time horizons within which they expect benefits. For example, stakeholders of different generations, some perhaps not yet born, have different time horizons within which they expect to receive benefits. However, once the mapping from attributes to utility is performed and utility metrics are determined, one has to assume that these metrics can be compared quantitatively.

The second important issue concerns the relative importance of stakeholders. Equation (6) implies that the overall utility attached to each stakeholder's utility can differ. For example, it is often the case that primary stakeholders' preferences receive more weight than the preferences of secondary stakeholders. The difficulty of this issue is obvious. Who decides? Is there a super-stakeholder, for instance? Do the groups of stakeholders, or their representatives, simply vote on who gets how much weight? Such a procedure has its own theoretical problems, which cannot be addressed here.

Beyond these two more theoretical issues, there are substantial practical issues associated with determining the functional forms of $u(c_{ij})$ and $u(b_{ij})$ and the parameters within these functional relationships. This also is true for the higher level forms represented by Equations (4)–(6). As the number of stakeholders (K), cost attributes (L), benefit attributes (M), and time periods (N) increase, these practical assessment problems can be quite daunting.

17.2.3 Option-Pricing Theory

Many investment decisions are not made all at once. Instead, initial investments are made to create the potential for possible future and usually larger investments involving much greater benefits than likely for the initial investments. For example, investments in R&D are often made to create the intellectual property and capabilities that will support or provide the opportunity to subsequently decide whether or not to invest in launching new products or services. These launch decisions are contingent on R&D reducing uncertainties and risks as well as further market information being gained in the interim between the R&D investment decision and possible launch decision. In this way, R&D investments amount to purchasing options to make future investments and earn subsequent returns. These options, of course, may or may not be exercised.

Amram and Kulatilaka (1999), Boer (1998, 1999), Lint and Pennings (1998), and Luehrman (1998) advocate using option-pricing theory to analyze investments involving such contingent downstream decisions. Option-pricing theory focuses on establishing the value of an option to make an investment decision, in an uncertain environment, at a later date. Equations (7)–(9) summarize the basic calculations as outlined by Luehrman.

One of two central elements in option pricing is the value of the returns from the contingent investment decision should one choose to make this decision, i.e., exercise the option. This value is represented by the NPV (net present value) quotient. As indicated in Equation (7), this quotient is formed as the ratio of the present asset value, that it, the traditional NPV of the free cash flow projected to result from exercising the option, and the present value (PV) of the investment required to acquire these assets, i.e., the option exercise price, X. As shown by Equation (8), the latter present value decreases as the risk-free rate of return, r_f, increases and/or the time until the option must be exercised or expire increases:

$$\text{NPV quotient} = \frac{\text{present asset value}}{\text{PV}(X)} \tag{7}$$

$$\text{PV}(X) = \frac{\text{present required investment value}}{(1 + r_f)^t} \tag{8}$$

The use of a risk-free rate is premised on the assumption that PV(X) will be invested now and accrue interest at rate r_f for t time periods so that the exercise price, X, will be available when the option can be exercised. The risk-free rate is used because these funds are not at risk until investors decide to exercise the option. If they choose to let the option expire, they retain X for other purposes.

The second central element in option pricing is the cumulative volatility expressed, as shown in Equation (9), as the product of the standard deviation of returns per period times the square root of the number of periods:

$$\text{Cumulative volatility} = \sigma\sqrt{t} \tag{9}$$

where σ^2 is the variance of returns per time period (expressed as a fraction) and t equals the number of time periods. The inclusion of volatility in option-pricing models is central to realistically representing investments where it is seldom the case that future returns are certain.

It might be imagined that estimating σ^2 is quite difficult. However, it is common to use numbers in the 0.30 to 0.60 range (Luehrman, 1998) and sensitivity analysis to determine the extent to which particular option values depend on these estimates. It is important to keep in mind that the goal is making a well-informed investment decision, *not* making highly precise estimates of variables that only coarsely affect the decision at hand.

The values of the NPV quotient [Equation (7)] and cumulative volatility [Equation (9)] are used to ascertain Black–Scholes values. These values are computed from the partial differential equation that is the Black–Scholes option-pricing model (Black and Scholes, 1973). For certain classes of option-pricing problems (e.g., constant σ^2 and fixed t), Black–Scholes numbers have been precomputed and can be found tabulated in various publications (e.g., Amram and Kulatilaka,1999; Luehrman, 1998).

The Black–Scholes values, expressed as percentages, increase with increasing NPV quotient and increasing cumulative volatility. This percentage is multiplied times the present asset value in Equation (7) to determine the value of the option. Thus, in general, the value of an option to later decide on an investment increases with r_f, σ^2, and t. In particular, in the presence of high volatility and high risk-free returns, the longer one can wait to decide, the more valuable the option. However, the present asset value in Equation (7) decreases with time. Thus, depending on the specific cash flow and investment projections as well as the parameters chosen, the option value may increase, decrease, or possibly increase to a maximum and then decrease. Sensitivity analysis is a good way to gain an understanding of this range of possibilities.

The resulting option value is totally premised on the assumption that waiting does not preempt deciding later. In other words, the assumption is that the decision to exercise an option cannot be preempted by somebody else deciding earlier. In typical situations where other actors (e.g., competitors) can affect possible returns, it is common to represent their impact in terms of changes of projected cash flows (Amram and Kulatilaka, 1999). In many cases, competitors acting first will decrease potential cash flows that will decrease the option value. It is often possible to construct alternative competitive scenarios and determine an optimal exercise date.

A central attraction of this model is the explicit recognition that the purpose of an investment now (i.e., purchasing an option) is to ensure the option to make a subsequent investment later (i.e., exercise the option). For example, one invests in creating new technologies for the option of later incorporating these technologies in product and service lines. The significance of the contingent nature of this decision makes an option-pricing model a much better fit than a traditional discounted cash flow model.

However, not all long-term investment decisions have substantial contingent elements. For example, one may invest in training and development to later have the option of selecting among talented managers for elevation to executive positions. There are minimal investments associated with exercising such options; almost all of the investment occurs up front. Thus option-pricing models are not useful for such decisions.

17.2.4 Knowledge Capital Approach

Tangible assets and financial assets usually yield returns that are important elements of a company's overall earnings. It is often the case, however, that earnings far exceed what might be expected from these "hard" assets. For example, companies in the software, biotechnology, and pharmaceutical industries typically have much higher earnings than companies with similar hard assets in the aerospace, appliance, and automobile industries, to name just a few. It can be argued that these higher earnings are due, for example, to greater knowledge capital among software companies. However, since knowledge capital does not appear on financial statements, it is very difficult to identify and, better yet, to project knowledge earnings.

A recent article by Mintz (1998) summarizes a method developed by Baruch Lev for estimating knowledge capital and earnings. This article in *CFO* drew sufficient attention to be discussed in The *Economist* (1999) and reviewed by Strassman (1999). In general, both reviews applauded the progress represented by Mintz's article but also noted the shortcomings of his proposed metrics.

The key, Mintz and Lev argue, is to partition earnings into knowledge earnings and hard asset earnings. Equation (10) accomplishes this by first projecting normalized annual earnings from an average of three past years and estimates for three future years using readily available information. Earnings from tangible and financial assets are calculated from reported asset values using industry averages of 7 and 4.5 percent for tangible and financial assets, respectively. Knowledge capital is then estimated by dividing knowledge earnings by a knowledge capital discount rate, as shown in Equation (11). Based on an analysis of several knowledge-intensive industries, Mintz and Lev use 10.5 percent for this discount rate.

$$\begin{aligned}\text{Knowledge earnings} &= \text{normalized annual earnings}\\ &\quad - \text{earnings from tangible assets}\\ &\quad - \text{earnings from financial assets} \qquad (10)\end{aligned}$$

$$\text{Knowledge capital} = \frac{\text{knowledge earnings}}{\text{knowledge capital discount rate}} \qquad (11)$$

Using this approach to calculate knowledge capital, Mintz compares 20 pharmaceutical companies to 27 chemical companies. He determines, for example, knowledge capital–book value ratios of 2.45 for pharmaceutical companies and 1.42 for chemical companies. Similarly, the market value–book value ratios are 8.85 for pharmaceutical companies and 3.53 for chemical companies. Considering this correlation between knowledge capital and market value, Strassman (1999) points out that Mintz's estimates do not fully explain the full excess of market values over book values.

The key issue within this overall approach is being able to partition earnings. While earnings from financial assets should be readily identifiable, the distinction between tangible and knowledge assets is problematic. Further, using industry average return rates to attribute earnings to tangible assets does not allow for the significant possibility of tangible assets having little or no earnings potential. Finally, of course, simply attributing all earnings "left over" to knowledge assets amounts to giving knowledge assets credit for everything that cannot be explained by traditional financial methods.

Nevertheless, the knowledge capital construct appears to have potential application to investments involving, for example, R&D or training and development. The purpose of these two types of investments seems to obviously be that of increasing knowledge capital. Further, companies that make investments for this purpose do seem to create more knowledge capital. The key for cost–benefit analyses is being able to project investment returns in terms of knowledge capital and, in turn, project earnings and separate these earnings into knowledge earnings and hard earnings. Further, one needs to be able to do this for specific investment opportunities, not just the company as a whole.

17.2.5 Comparison of Frameworks

Table 17.1 provides a comparison of the four frameworks just reviewed. It is important to note that this assessment is not really an apples-to-apples comparison. Multiattribute utility theory provides much more of a general framework than the other three approaches, which emphasize financial metrics. Nevertheless, these four approaches represent the dominant alternatives.

Traditional economic analyses are clearly the most narrow. However, in situations where they apply, these analyses are powerful and useful. Most of the investment situations addressed in this chapter do not fit these narrow characteristics. For example, if R&D investments in the human effectiveness aspects of HSI are viewed within a traditional framework, with typical discount rates, no one would ever invest anything in such R&D. But people do make such investments and, thus, there must be more to it than just NPV, IRR, and CBR. [In fact, Cooper et al. (1998b) have found that companies relying solely on financial metrics for R&D investment decisions tend to be the poorest performers of R&D in terms of subsequent market success.]

One view is that R&D reduces uncertainty and buys time before committing very substantial resources to productization, process development, etc. Option-pricing theory seems to be a natural extension of traditional methods to enable handling these complications. As noted earlier, several authors have advocated this approach for analyses of R&D investments.

The knowledge capital approach provides another, less mathematical way of capturing the impacts of R&D investments in human effectiveness aspects of HSI. The difficulty of this approach, which is probably inherent to its origins in accounting and finance, is that it does not address the potential impacts of alternative investments. Instead, it serves to report the overall enterprise score after the game.

Multiattribute utility models can, in principle, address the full range of complications and complexity discussed thus far. Admittedly, the ability to create a rigorous multiattribute utility model depends on the availability of substantial amounts of information regarding stakeholders' preference spaces, probability density functions, etc. However, in

TABLE 17.1 **Comparison of Cost–Benefit Frameworks**

Issue/Framework	Traditional Economic Analysis	Multiattribute Utility Models	Option-Pricing Theory	Knowledge Capital Approach
Representation of uncertainties	Focuses on expected revenues and costs without consideration of variances.	Probabilistic uncertainties and stakeholders' preferences regarding uncertainties are central to models.	Volatility of returns is a central construct within this model.	Focuses on actual and expected earnings without consideration of variances.
Intangible vs. tangible outcomes	All outcomes must be converted to monetary units.	Preferences regarding intangible outcomes can be incorporated.	All outcomes must be converted to monetary units.	All outcomes must be converted to monetary units.
Multiple stakeholders in costs/benefits	One-dimensional nature of costs and benefits implies one stakeholder.	Formulations for multiple stakeholders are available and limitations are understood.	One-dimensional nature of costs and benefits implies one stakeholder.	One-dimensional nature of costs and benefits implies one stakeholder.
Assessing vs. projecting costs/benefits	Depends on abilities to project monetary costs and benefits.	Depends on abilities to project attributes of utility functions.	Depends on abilities to project monetary costs and benefits.	Difficult to project impact of particular investments.

the absence of such information, a much more qualitative approach can be quite useful, as is discussed later in this chapter.

The value of the multiattribute utility approach also depends on being able to compare overall utilities of alternative investments, which in turn depends on being able to compare different stakeholders' utilities of the alternatives. This ability to transform a complex, multidimensional comparison into a scalar comparison is laden with assumptions. The saving grace of the approach, in this regard, is that it makes these assumptions quite explicit and, hence, open to testing. This does not, of course, guarantee that they will be tested.

Expected-utility calculations serve to show how one alternative is better than another, rather than providing absolute scores. Thus, differences of expected utilities among alternatives are usually more interesting than the absolute numbers. In fact, the dialog among stakeholders that is often associated with trying to understand the sources of expected-utility differences can provide crucial insights into the true nature of differences among alternatives.

Overall, one must conclude that multiattribute utility models provide the most generalizable approach. This is supported by the fact that multiattribute models can incorporate metrics such as NPV, option value, and knowledge capital as attributes within the overall model; indeed, the special case of one stakeholder, linear utility functions, and NPV as the sole attribute is equivalent to the traditional financial analysis. Different stakeholders' preferences for these metrics can then be assessed and appropriate weightings determined. Thus, use of multiattribute models does not preclude also taking advantage of the other approaches; the four approaches, therefore, can be viewed as complementary rather than competing. For these reasons, the multiattribute approach is carried forward in the remainder of this chapter.

17.3 COST–BENEFIT METHODOLOGY

Cost–benefit analysis should always be pursued in the context of particular decisions to be addressed. A valuable construct for facilitating an understanding of the context of an analysis is the value chain from investments to returns. More specifically, it is quite helpful to consider the value chain from investments (or costs), to products, to benefits, to stakeholders, to utility of benefits, to willingness to pay, and finally to returns on investments. This value chain can be depicted as follows:

Investments (costs)	to	Resulting products over time
Products over time	to	Benefits of products over time
Benefits over time	to	Range of stakeholders in benefits
Range of stakeholders	to	Utility of benefits to each stakeholder
Utility to stakeholders	to	Willingness to pay for utility gained
Willingness to pay	to	Returns to investors

The process starts with investments, which result or will result in particular products over time. Products need not be end products; they might be knowledge, skills, or technologies. These products yield benefits, also over time. A variety of people—or stakeholders—have a stake in these benefits. These benefits provide some level of utility to each stakeholder. The utility perceived or anticipated by stakeholders affects their willingness to pay for these benefits. Their willingness to pay affects their "purchase" behaviors, which result in returns for investors.

The central methodological question concerns how one can predict the inputs and outputs of each element of this value chain. This question is addressed elsewhere in some detail for R&D management (Rouse et al., 1997; Rouse and Boff, 1997) and for human effectiveness (Rouse and Boff, 1997). Briefly, a variety of models have been developed for addressing this need for prediction. These models are very interesting and offer much potential. However, they suffer from a central shortcoming. With few exceptions, there is an almost overwhelming lack of data for estimating model parameters as well as a frequent lack of adequate input data. Use of data from baselines can help, but the validity of these baselines depends on new systems and products being very much like their predecessors. Overall, the paucity of data dictates development of a more qualitative methodology whose usefulness is not totally determined by availability of hard data. The remainder of this section outlines such a methodology.

As indicated in the earlier comparison of four frameworks for addressing cost–benefit analysis, the most broadly applicable of these alternatives are multiattribute utility models. The remainder of this section describes the following seven-step methodology:

- Step 1: Identify stakeholders in alternative investments.
- Step 2: Define benefits and costs of alternatives in terms of attributes.
- Step 3: Determine utility functions for attributes (benefits and costs).
- Step 4: Decide how utility functions should be combined across stakeholders.
- Step 5: Assess parameters within utility models.
- Step 6: Forecast levels of attributes (benefits and costs).
- Step 7: Calculate expected utility of alternative investments.

It is important to note that this methodology is by no means novel and builds upon works by many others related to multiattribute analysis (e.g., Keeney and Raiffa, 1976; Sage, 1977; Hammond et al., 1998; Matheson and Matheson, 1998; Sage and Armstrong, 2000).

Step 1: Identify Stakeholders The first step involves identifying the stakeholders who are of concern relative to the investments being entertained. Usually this includes all the people in the value chain summarized earlier. This might include, for example, those who will provide the resources that will enable a solution, those who will create the solution, those who will implement the solution, and those who will benefit from the solution.

Step 2: Define Benefit and Cost Attributes The next step involves defining the benefits and costs involved from the perspective of each stakeholder. These benefits and costs define the attributes of interest to the stakeholders. Usually, a hierarchy of benefits and costs emerges, with more abstract concepts at the top, e.g., viability, acceptability, and validity (Rouse, 1991), and concrete measurable attributes at the bottom.

Step 3: Determine Stakeholders' Utility Functions The values that stakeholders attach to these attributes are defined by stakeholders' utility functions. The utility functions enable mapping disparate benefits and costs to a common scale. A variety of techniques are available for assessing utility functions (Keeney and Raiffa, 1976).

Step 4: Determine Utility Functions Across Stakeholders Next, one determines how utility functions should be combined across stakeholders. At the very least, this involves assigning relative weights to different stakeholders' utilities. Other considerations such as desires for parity can make the ways in which utilities are combined more complicated. For example, Equation (6) may require interaction terms to assure all stakeholders some utility.

Step 5: Assess Parameters of Utility Functions The next step focuses on assessing parameters within the utility models. For example, utility functions that include diminishing or accelerating increments of utility for each increment of benefit or cost involve rate parameters that must be estimated. As another instance, estimates of the weights for multistakeholder utility functions have to be estimated. Fortunately, there are a variety of standard methods for making such estimates.

Step 6: Forecast Levels of Attributes With the cost–benefit model fully defined, one must next forecast levels of attributes or, in other words, benefits and costs. Thus, for each alternative investment, one must forecast the stream of benefits and costs that will result if this investment is made. Quite often, these forecasts involve probability density functions rather than point forecasts. Utility theory models can easily incorporate the impact of such uncertainties on stakeholders' risk aversions. On the other hand, information on probability density functions may not be available or may be prohibitively expensive. In these situations, beliefs of stakeholders and subject matter experts can be employed, perhaps coupled with sensitivity analysis (see step 7) to determine where additional data collection may be warranted.

Step 7: Calculate Expected Utilities The final step involves calculating the expected utility of each alternative investment. These calculations are performed using specific forms of Equations (4)–(6). This step also involves using sensitivity analysis to assess, for example, the extent to which the rank ordering of alternatives, by overall utility, changes as parameters and attribute levels of the model are varied.

Use of the Methodology Some elements of the cost–benefit methodology just outlined are more difficult than others. The overall calculations are quite straightforward. The validity of the resulting numbers depends, of course, on stakeholders and attributes having been identified appropriately. It further depends on the quality of the inputs to the calculations.

These inputs include estimates of model parameters and forecasts of attribute levels. As indicated earlier, the quality of these estimates is often compromised by lack of available data. Perhaps the most difficult data collection problems relate to situations where the impacts of investments are both uncertain and very much delayed. In such situations, it is not clear which data should be collected and when they should be collected.

A recurring question concerns the importance that should be assigned to differences in expected-utility results. If alternative A yields $U(A) = 0.648$ and alternative B yields

$U(B) = 0.553$, is A really that much better than B? In fact, are either utilities sufficiently great to justify an investment?

These questions are best addressed by considering past investments. For successful past investments, what would their expected utilities have been at the time of the investment decisions? Similarly, for unsuccessful past investments, what were their expected utilities at the time? Such comparisons often yield substantial insights.

Of course, the issue is not always A versus B. Quite often the primary question concerns which alternatives belong in the portfolio of investments and which do not. Portfolio management is a fairly well-developed aspect of new product development (Cooper et al., 1998a; Gill et al., 1996). Well-known and recent books on R&D/technology strategy pay significant attention to portfolio selection and management (Roussel et al., 1991; Matheson and Matheson, 1998; Boer, 1999; Allen, 2000). In fact, the conceptual underpinnings of option-pricing theory are based on notions of market portfolios (Amram and Kulatilaka, 1999).

Most portfolio management methods rely on some scoring or ranking mechanism to decide which investments will be included in the portfolio. Expected utility is a quite reasonable approach to creating such scores or ranks. This is particularly useful if sensitivity analysis has been used to interactively explore the basis and validity of differences among alternatives.

A more sophisticated view of portfolio management considers interactions among alternatives in the sense that synergies between two alternatives may make both of them more attractive (Boer, 1999; Allen, 2000). Also correlated risks between two alternatives may make both of them less attractive. A good portfolio has an appropriate balance of synergies and risks.

In principle, at least, the notions of portfolio synergy and risk can be handled within multiattribute utility models. This can be addressed by adding attributes that are characteristics of multiple rather than individual alternatives. In fact, such additional attributes might be used to characterize the whole portfolio. An important limitation of this approach is the likely significant increase in the complexity of the overall problem formulation. Indeed, this is an issue in general when multiattribute utility models are elaborated to better represent problem complexities.

Beyond these technical issues, it is useful to consider how this cost–benefit methodology should affect decision making. To a very great extent, the purpose of this methodology is to get the right people to have the right types of discussions and debates on the right issues at the right time. If this happens, the value of people's insights from exploring the multiattribute model usually far outweighs the importance of any particular numbers.

The practical implications of this conclusion are quite simple. Very often, decision making happens within working groups who view computer-generated, large-screen displays of the investment problem formulation and results as they emerge. Such groups perform sensitivity analyses to determine the critical assumptions or attribute values that are causing some alternatives to be more highly rated or ranked than others. They use "What if . . .?" analyses to explore new alternatives, especially hybrid alternatives.

This approach to investment decision making helps to substantially decrease the impact of limited data being available. Groups quickly determine which elements of the myriad of unknowns really matter—where more data are needed and where more data, regardless of results, would not affect decisions. A robust problem formulation that can be manipulated, redesigned, and tested for sanity provides a good way for decision-making groups to reach defensible conclusions with some level of confidence and comfort.

17.4 THREE EXAMPLES

Human effectiveness concerns enhancing people's direct performance (aiding), improving their potential to perform (training), and ensuring their availability to perform (health and safety). These are central issues in HSI. Investments in human effectiveness also have the potential of increasing returns on other investments by, for example, enabling people to take full advantage of new technologies.

Three examples of aiding, training, and health and safety investments are discussed in this section—VCATS (aiding), DMT (training), and PTOX (health and safety). These examples focus on enhancing human effectiveness and HSI in military systems, particularly Air Force systems. The applicability of these technologies, and the relevance of the following analysis of the impacts of these technologies, to other military services and to non-military problems should also be readily apparent.

17.4.1 Visually Coupled Targeting and Acquisition System

The Visually Coupled Targeting and Acquisition System (VCATS) provides aiding to military aircraft pilots. The VCATS includes a helmet-mounted tracker and display (HMT/D), associated signal processing sensor/transducer hardware, interchangeable panoramic night vision goggle with head-up display (PNVG-HUD), and extensive upgrades to the aircraft's operational flight program software (Rastikis, 1998). The VCATS enables the pilot to cue and be cued by on-board and off-board systems, sensors, and weapons as well as be spatially and temporally coupled with the control processes implemented with the HMT/D and PNVG-HUD. The system is particularly effective in helping pilots to cue weapons and sensors to targets, maintain "ownship" formation situation awareness, and avoid threats via provision of a real-time, three-dimensional portrayal of the pilots' tactical and global battlefield status. In general, the VCATS enables pilots to acquire targets and threats faster. This results in improvements in terms of (1) how far, (2) how quickly, and (3) how long—for both initial contacts and countermeasures.

To a great extent, the case for advanced development has already been made for the VCATS and current support is substantial. However, the transition from advanced development to production involves ensuring that the options created by the VCATS and validated by combat pilots are exercised. The case has also been argued for ongoing investments in basic research and exploratory development to ensure that the VCATS has future technology options, particularly for migration to multirole fighter aircraft. The maturity of the program should help in making this case in terms of benefits already demonstrated. However, in the current budget climate, there is also substantial risk that VCATS research may be viewed as essentially "done." This raises the potential for negative decisions regarding further investments.

17.4.2 Distributed Mission Training

Distributed mission training (DMT) involves aircraft, virtual simulators, and constructive models that, collectively, provide opportunities for military pilots to gain experiences deemed important to their performance proficiency relative to anticipated mission requirements (Andrews, 2000). The desired training experiences are determined from competencies identified as needed to fulfill mission requirements. These competency requirements are translated to training requirements stated in terms of types and durations of experiences deemed sufficient to gain competency.

The case to be made for DMT involves investments to address research issues and technology upgrades of near-term capabilities. The primary options-oriented argument is that investments in R&D in DMT will create contingent possibilities for cost savings in training due to reduced use of actual aircraft. More specifically, DMT options, if exercised, will provide cash flows of savings that justify the investments needed to field this family of technologies.

A much more subtle options-oriented argument concerns the training experiences provided by DMT that could not otherwise be obtained. Clearly, the opportunity to have relevant training experiences must be better than not having these experiences. The option, therefore, relates to proficiency versus possible lack of proficiency.

As straightforward as this may seem, it quickly encounters the difficulty of projecting mission impacts—and the value of these impacts—of not having proficient personnel. One possible approach to quantifying these benefits is to project the costs of using real aircraft to gain the desired proficiencies. While these costs are likely to be prohibitive, and thus never would be seriously considered, they nevertheless characterize the benefits of DMT.

17.4.3 Predictive Toxicology

Predictive toxicology (PTOX) is concerned with projecting the impacts on humans from exposure to operational chemicals (individual and mixtures). The impact can be characterized in terms of the possibility of performance decrement and consequent loss of force effectiveness, possible military and civilian casualties, and potential long-term health impacts. Also of concern are the impacts of countermeasures relative to sustaining immediate performance and minimizing long-term health impacts [Office of Science and Technology Policy (OSTP), 1998].

The case to be made involves investment in basic research and exploratory development programs, with longer term investment in an advanced development program to create deployable predictive toxicology capabilities. The requisite R&D involve developing and evaluating models for predicting performance and health impacts of operational chemicals. Advanced development will focus on field sensing and prediction, termed deployment toxicology. The nature of the necessary models is strongly affected by the real-time requirements imposed by deployment.

17.4.4 Applying the Methodology

The remainder of this section primarily addresses steps 1 to 4 of the cost–benefit methodology in the context of these three examples related to human effectiveness aspects of HSI. These steps constitute the "framing" steps of the methodology, rather than the "calculation" steps. Appropriate framing of cost–benefit analyses is critical to subsequent calculations being meaningful and useful.

Step 1: Identify Stakeholders This step involves identifying people, usually types of people, and organizations that have a stake in costs and benefits. All three of the examples involve three classes of stakeholders: warfighters, developers, and the public. A key issue concerns the relative importance of these three types of stakeholders. Some would argue that warfighter preferences dominate decisions. Others recognize the strong role that developers and their constituencies play in procurement decisions. Yet another argument is that the dominating factor is value to the public, with the other stakeholders being secondary in importance.

Warfighters as stakeholders include military personnel in general, especially for PTOX. Warfighters of particular importance include aircraft pilots, personnel who support flight operations, and military commanders. Developers as stakeholders include companies and their constituencies, e.g., stockholders, employees, and communities. The public's interests are represented by several agents, including Congress, the executive functions within the military services, and the military procurement establishment. Pilots and other military personnel are users of the technologies of interest, developers are the providers, and the public's agents are the customers for these technologies. There are obvious trade-offs across the interests of users, providers, and customers.

Step 2: Define Benefit and Cost Attributes Benefits and costs tend to fall in general classes. Examples of benefits for military organizations and contractors include the following:

Enhanced impact	Increased lethality, survivability, and availability
Enhanced operability	Decreased response time and increased throughput
Enhanced design	New techniques and larger pool of experienced people
Increased opportunities	New tactics and countermeasures

Example cost attributes applicable to military procurement include the following:

Investment costs	Capital investments and R&D costs
Recurring costs	Operating and General and Administrative (G&A) costs
Time costs	Time from development to fielding to competent use
Opportunity costs	Other costs/benefits foregone

These general classes of benefits and costs can be translated into specific benefit and cost attributes for the three classes of stakeholders in VCATS, DMT, and PTOX. Benefits for warfighters (users) include enhanced performance (e.g., response time), confidence in performance, and health and safety in varying combinations for the three examples. Costs for these stakeholders include learning time and changing their ways of doing things to ensure compatibility between new and legacy technologies.

Benefits for companies and their constituencies (providers) include R&D funds received, subsequent intellectual property created, and competitive advantages that result. Also important are jobs and economic impacts in the community. Direct costs include bid and proposal costs as well as opportunity costs. Less direct costs include, for instance, economic development resources and incentives provided to the companies by their communities.

The primary benefit sought by the public's agents (customer) is mission performance per dollar. It can easily be argued for all three examples that mission performance is increased. Unfortunately, it is difficult to attach a value to this increase. For example, what is the value of being able to generate 5 percent more sorties per time period? The answer depends on whether more sorties are needed.

Few would argue with the importance of successfully meeting mission requirements. However, if the types of innovations represented by these examples enable exceeding mission requirements, what are such increases worth? This is a politically sensitive question. If better performance is of substantive value, why was not this level of performance specified in the original requirements?

A good way to avoid this difficulty is to take mission requirements as a given and determine how much money could be saved in meeting these requirements by adopting the technologies in question. For example, could requirements be met with fewer aircraft, pilots, and support personnel? As shown in Table 17.2, the cost savings due to these decreases can be viewed as benefits of the technologies. It also might be possible for the VCATS, DMT, or PTOX to enable meeting mission requirements with less capable systems, rather than just fewer systems. This possibility provides substantial opportunities for increased benefits due to these technologies.

Note that this philosophy amounts to trying to provide a given level of defense for the least investment. Another approach might be to attempt to provide the most defense per investment dollar. However, this immediately begs the question of how much defense is enough. Unlike the business world where value is defined by the marketplace and, hence, can provide a basis for optimization—see Nevins and Winner (1999) for a good example—there is no widely agreed-upon approach to measuring military value and optimizing accordingly.

The rationale for the benefits indicated in Table 17.2 for each of the three examples includes the following:

- The VCATS enables pilots to compete with threats, increase the number of wins versus losses, and counter threats (e.g., missiles) in ways that they could not do otherwise. Consequently, it must be possible to meet *fixed* mission requirements with fewer aircraft and associated infrastructure. These benefits can be translated into financial returns in terms of cost avoidance.
- Distributed mission training provides opportunities to practice behaviors that would not otherwise be practiced, for the most part due to the costs of practice. This decreases the probability of not performing acceptably given inadequate training. Also, DMT provides training experiences that would not otherwise be possible. For example, in the DMT environment, pilot "kills" actually disappear. In contrast, field exercises often "reuse" kills because of the costs of getting adversaries into the exercise in the first place.
- Predictive toxicology enables larger proportions of deployed forces to be fully functional and less dependent on medical surveillance or medication and earlier intervention before the onset of problems. In principle, this should enable reducing the size of deployed force, which is critical for increasingly likely expeditionary military missions (Fuchs et al., 1997).
- In addition, PTOX provides cost avoidance due to downstream health impacts. The ability to predict the "body burden" of toxicity during deployment should enable removing personnel from risk once the burden is approaching predetermined limits. These capabilities are likely to also be very important for nonmilitary operations such as disaster clean-up.

It is *not* essential that the savings indicated in Table 17.2 actually occur. For example, it may be that the number of aircraft is not decreased, perhaps due to factors far beyond the

TABLE 17.2 Public Benefits and Costs for Three Examples

	VCATS	DMT	PTOX
Benefits	Fewer aircraft and associated personnel to meet mission requirements due to better performance and fewer aircraft losses	Fewer aircraft and associated personnel to meet mission requirements due to better performance, fewer aircraft losses, and fewer aircraft for training	Fewer personnel to meet mission requirements and decreased medical costs due to fewer people affected by toxic materials, fewer people lost to toxic effects, fewer people to care for people affected, and decreased downstream medical costs
Costs	Initial investment (option price) for proposed R&D costs and, later, contingent investment (exercise price) for subsequent fielding of technology	Initial investment (option price) for proposed R&D costs and, later, contingent investment (exercise price) for subsequent fielding of technology	Initial investment (option price) for proposed R&D costs and, later, contingent investment (exercise price) for subsequent fielding of technology

scope of these analyses. However, one can nevertheless attribute to these technologies the benefits of having provided opportunities to meet mission requirements in less costly manners. Technologies that provide such opportunities are valuable; the extent of this value is the extent of the opportunities for savings.

This argument puts all three examples on common ground. The benefits of all alternative technologies can be expressed as reduced costs to meet requirements. From an options-pricing perspective, these savings can be viewed as free cash flow returned on investments in these technologies. The "option price" is the R&D costs. The "exercise price" is the subsequent costs of fielding the technologies. Thus, assuming costs savings can be projected (albeit with substantial volatility), the options value of investing in these technologies can be calculated.

Step 3: Determine Stakeholders' Utility Functions Different stakeholders' preferences over the benefit and cost attributes will vary substantially with specific situations. However, there is a small family of functional relationships that captures most, if not all, expressed preferences (Keeney and Raiffa, 1976). Thus while context-specific tailoring is needed, it can be performed within a prescribed (and preprogrammed) set of functions, both within and across stakeholders. Similarly, alternative parameter choices can be prescribed in terms of choices of weightings.

An important aspect of cost–benefit analyses, as advocated in this chapter, is the likely nonlinear nature of utility functions. In particular, diminishing returns and aspiration levels tend to be central to stakeholders' "preference spaces." In other words, while linear functions imply that incremental increases (or decreases) of attributes always yield the same incremental changes in utility, nonlinear functions lead to shifting preferences as attributes increase (or decrease). Figure 17.1 portrays a range of example utility functions; *a* and *d* illustrate linear relationships; *b* and *e* show accelerating relationships; and *c* and *f* show diminishing relationships.

To illustrate how these types of relationships can be employed to represent the preferences of users, providers, and customers, the general forms of each type of

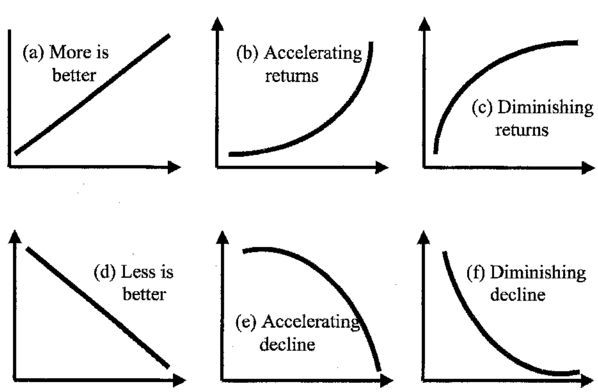

Figure 17.1 Example utility functions: (*a*) More is better, (*b*) accelerating returns, (*c*) diminishing returns, (*d*) less is better, (*e*) accelerating decline, and (*f*) diminishing decline.

stakeholder's utility function are shown in Equations (12)–(14):

$$U_{\text{user}} = U[u(\text{performance}), u(\text{confidence}), u(\text{cost of change})] \tag{12}$$

$$U_{\text{provider}} = U[u(\text{resources}), u(\text{advantage}), u(\text{cost of pursuit})] \tag{13}$$

$$U_{\text{customer}} = u(\text{option value}) \tag{14}$$

where, as noted earlier, users are primarily concerned about impacts of investments on their performance, their confidence in their performance, and the costs of changing their ways of performing; providers are concerned with the investment resources supplied to develop the technologies in question, the competitive advantages created by the intellectual property created, and the costs of pursuing the investment opportunities; and, finally, customers are focused on the financial attractiveness of the investments as reflected in the option values of the alternatives, which are based on projected cash flows (i.e., costs savings), volatility of cash flows, magnitudes of investments required, and time periods until returns are realized.

Considering the elements of Equations (12)–(14), the appropriate functional forms from Figure 17.1 are likely to be as follows:

- *u*(performance) is an accelerating returns function (Fig. 17.1*b*):
  - The VCATS is least concave since relatively modest performance improvements are of substantial utility.
  - Distributed mission training is moderately concave since training on otherwise untrained tasks must produce substantial improvements to yield high utility.
  - Predictive toxicology is most concave since major decreases in performance risk are needed to ensure high utility increases of personnel availability.
- *u*(confidence) is a linear function (Fig. 17.1*a*) since greater confidence is always better but there are unlikely to be significant thresholds.
- *u*(cost of change) is an accelerating decline function (Fig. 17.1*f*) since low to moderate costs are easily sustained while larger costs present difficulties.
- *u*(resources) is an accelerating returns function (Fig. 17.1*b*) since moderate to large resources are needed to make opportunities attractive.
- *u*(advantage) is a linear function (Fig. 17.1*a*) since greater advantage is always better but there are unlikely to be significant thresholds.
- *u*(cost of pursuit) is an accelerating decline (Fig. 17.1*f*) since low to moderate costs are easily sustained while larger costs present difficulties.
- *u*(option value) is a linear function (Fig. 17.1*a*) since customers will inherently gain the expected value across a large number of investments.

It is important to note the importance of this last assumption. If customers'—that is, the public's—utility function were not linear, it would be necessary to entertain assessing the specific form of their function. Unlike users and providers, the public is not so easily identified and interviewed.

With the identification of the stakeholders (step 1) and framing of the cost–benefit attributes (step 2), the process of determining the form of stakeholders' utility functions (step 3) can draw upon considerable standard "machinery" of decision analysis. The specific versions of the functional forms discussed above are likely to vary with the VCATS, DMT, and PTOX. However, the overall formulation chosen is quite general.

Step 4: Determine Utility Functions Across Stakeholders Another important aspect of the utility functions is their typical lack of alignment across stakeholders. Specifically, either different stakeholders care about different things or possibly they care about the same things in different ways. For example, customers may be very price sensitive while users, who seldom pay prices themselves, are usually much more concerned with impacts on their job performance

For the types of investment problems considered in this chapter, preferences typically differ across time horizons and across people with vested interests in different investment opportunities. Thus far in the formulation of the three examples, the stakeholders do not have attributes in common. However, they are nevertheless likely to have competing preferences since, for example, the alternative providing the greatest performance impact may not have the largest option value.

Differing preferences across stakeholders are often driving forces in pursuing cost–benefit analyses. These differing preferences can be aggregated, and traded off, by formulating a composite utility function such as

$$U = U[U_{\text{user}},\ U_{\text{provider}},\ U_{\text{customer}}] \tag{15}$$

Often Equation (15) will be linear in form with weights assigned to component utility functions to reflect the relative importance of stakeholders. Slightly more complicated are multilinear forms that include products of component functions, e.g., $U_{\text{user}} \times U_{\text{customer}}$. Multilinear formulations tend to ensure that all stakeholders gain nonzero utility because, otherwise, zero in either term in a product yields zero overall.

Considering trade-offs across stakeholders, it is important to note that the formulation of the analysis can often be usefully expanded to include a broader set of stakeholders. These additional stakeholders may include other entities that will benefit by advances of the technologies in question, although they may have little or no stake in the immediate application for the technology. It is also quite possible that stakeholders such as "the public" have multiple interests, e.g., military effectiveness and public safety from toxic risks.

Broadening the analysis in this way is likely to have differing impacts on the assessment for the three examples due to the natures of the technologies and issues being pursued. The three examples differ in this regard in the following ways:

- The VCATS addresses a rather esoteric set of issues from the public's perspective.
- Distributed mission training addresses an issue with broad general support from the public but narrower specific constituencies.
- Predictive toxicology addresses strong cross-cutting health and safety issues of substantial concern to the public.

These differences suggest that PTOX would gain a larger ΔU than DMT, and DMT would in turn gain a larger ΔU than the VCATS, by broadening the number of stakeholders and issues. Quite simply, the "spin-off" benefits of PTOX are likely to be perceived as much greater by a larger number of stakeholders.

However, if the formulation is further broadened to consider the likelihood that the desired technologies will emerge elsewhere if investments are not made in these efforts, the ΔU impacts will likely be the opposite. The PTOX research and development are being pursued by several agencies. Distributed mission training has broad applicability for both military and nonmilitary applications and consequently is being pursued by other parties. The VCATS, in contrast, is highly specialized and is unlikely to emerge from other sources.

These two possibilities for broadening the formulation, in terms of stakeholders and issues, clearly illustrate the substantial impact of the way in which cost–benefit assessments are framed. If the framing is too focused, important spin-off benefits may not be included. On the other hand, framing the analysis too broadly may raise issues that are difficult to quantify, even roughly, and includes stakeholders whose preferences are difficult to assess. Of course, many modeling efforts face such difficulties (Sage and Rouse, 1999).

Steps 5–7: Calculation of Overall Cost–Benefit The remaining steps of the cost–benefit methodology involve assessing parameters of utility functions, forecasting levels of attributes, and calculating expected utilities. Performing these steps obviously depends on having data on stakeholders' preferences and projected/targeted attribute levels. Discussion of such data is well beyond the scope of this chapter and, in light of the nature of the examples, it would be difficult to publish the requisite data.

The needed data can, in many instances, be quite difficult to compile. It can be particularly difficult to relate returns on human effectiveness investments to organizational impacts. Relationships between human and organizational performance are needed. These relationships should answer the following types of questions:

- How do improvements in human performance (e.g., via aiding) translate to increased organizational impacts? Specifically, how does a 2-second improvement in pilot response time due to the VCATS affect mission performance?
- How do improvements of human potential to perform (e.g., via training) translate to actual performance and consequent increased organizational impacts? Specifically, how does increased practice via DMT impact subsequent performance and, in turn, translate to improved mission performance?
- How do improvements in human availability to perform (e.g., via health and safety) translate to actual performance and consequent increased organizational impacts? Specifically, how does prevention of toxic exposure, due to PTOX, affect immediate unit performance and thereby affect mission performance?

These can be difficult questions. However, they are not inherently cost–benefit questions. Instead, they are fundamental system design questions (Sage and Rouse, 1999). If answers are possible, then cost–benefit analyses are more straightforward.

For the VCATS, DMT, and PTOX examples, it may be possible to translate human performance improvements to organizational impacts via mission models. Such models are typically used to determine, for example, the "logistics footprint" needed to support a targeted sortie generation rate or, as another illustration, the combat wins and losses likely with competing defensive measures and countermeasures. Such models can be used, perhaps with extensions, to project the impacts of faster responses due to the VCATS, improved task performance due to DMT, and increased personnel availability due to PTOX.

It is important to note, however, that even if such projections are not available, the multiattribute methodology presented here can still be employed. However, the validity of cost–benefit assessments and predictions will then depend upon subjective perceptions of attribute levels and the relative importance of attributes. Any limitations of this more subjective approach reflect underlying limitations of knowledge rather than inherent limitations of the methodology.

Once $U[U_{user}, U_{provider}, U_{customer}]$ is fully specified, both functionally and in terms of parameters of these functions, one is in the position to project attribute levels (e.g., option values), calculate the expected utility of the alternative investments (e.g., VCATS, DMT, and PTOX), and perform sensitivity analyses. This provides the basis for making investment decisions. There are several ways that these cost–benefit assessments can be used to inform decision making.

The most common way of using expected-utility cost–benefit assessments is to rank order alternative investments in terms of decreasing $U[U_{user}, U_{provider}, U_{customer}]$ and then allocate investment resources from highest ranked to lowest ranked until resources are exhausted. This approach allows the possibility of alternatives with mediocre $U_{customer}$ making the cut by having substantial U_{user} and $U_{provider}$. To avoid this possibility, one can rank order by $U[U_{user}, U_{provider}, U_{customer}]$ all alternatives with $U_{customer} > U_{co}$, which implies a minimum acceptable option value.

If resources are relatively unconstrained, one can invest in all alternatives for which $U_{user} > U_{uo}$, $U_{provider} > U_{po}$ and $U_{customer} > U_{co}$. This reflects situations where all stakeholders prefer investment to no investment. Of course, one can also rank order these alternatives by $U[U_{user}, U_{provider}, U_{customer}]$ to determine priorities for investment. However, if resources are truly unconstrained, this rank ordering will not change the resulting investment decisions.

17.4.5 Summary

The three examples discussed in this section have portrayed a cross section of human effectiveness investments to enhance HSI, ranging from aiding to training to health-and-safety investments. The discussion has shown how this range of investment alternatives can be fully addressed with an overarching multiattribute utility, multistakeholder cost–benefit formulation. The stakeholder classes of user, provider, and customer are broadly applicable. The classes of attributes discussed also have broad applicability.

These examples have also served to illustrate the merits of a hybrid approach. In particular, option-pricing theory has been used to define the issue of primary interest to customers, ensuring that investments make financial sense, and this issue has then been incorporated into the overall multiattribute formulation. This enabled including in the formulation a substantial degree of objective rigor as well as important subjective attributes and perceptions. As a result, rigor is not sacrificed but instead is balanced with broader, less quantifiable considerations.

It is useful to note that the knowledge capital construct was not employed in the formulation for these three examples, despite the intuitive appeal of the notion that investments in human effectiveness increase knowledge capital (Davenport, 1999). While the formulation reported here could have included increases of knowledge capital as possible benefits, there is no basis for predicting such impacts. Subjective estimates could, of course, be employed. However, this construct is not defined with sufficient crispness to expect reliable estimates from subject matter experts.

The discussion of these examples of human effectiveness investments has served to illustrate the value of an overall cost–benefit formulation. The generality of this formulation allows it to be applied to analyses of a wide variety of HSI investment decisions. The types of information needed to support such analyses are defined by this formulation. While the availability of information remains a potential difficulty, this formulation nevertheless substantially ameliorates the typical problems of comparing ad hoc analyses of competing investments. Also of great importance, this formulation enables cross-stakeholder comparisons and trade-offs that, for the lack of a suitable methodology, are usually ignored or resolved in ad hoc manners.

17.5 CONCLUSIONS

It is difficult to make the case for long-term investments that will provide highly uncertain and intangible returns. This chapter has reviewed alternative ways to characterize such investments and presented an overall methodology that incorporates many of the advantages of these alternatives. This methodology has been illustrated in the context of R&D investments in human effectiveness aspects of HSI.

Central to the cost–benefit analysis methodology presented is a multiattribute, multi-stakeholder formulation. This formulation includes nonlinear preference spaces that are not necessarily aligned across stakeholders. The nonlinearities and lack of alignment provide ample opportunities for interesting trade-offs.

It is important to stress the applicability of this methodology to nearer term HSI investments, which may or may not involve R&D. While the time frame will certainly affect choices of attributes (for instance, option values may not be meaningful for near-term investments), the overall cost–benefit methodology remains unchanged. This chapter focused on long-term R&D investments because such analyses are the most difficult to frame and perform.

It is also useful to indicate that cost–benefit analysis, as broadly conceptualized in this chapter, can be a central element in assessment activities related to life-cycle costing (e.g., affordability) and program/contract management (e.g., earned-value management (Earned Value Management Center, 2000)). For the former, attributes reflecting life-cycle costs can easily be incorporated. For the latter, costs and benefits can be tracked and compared to original projections. This does, of course, require that benefits be attributable to ongoing processes and not just outcomes.

This cost–benefit methodology, when coupled with appropriate methods and tools for predicting attribute levels (Sage and Rouse, 1999), can enable cost–benefit predictions and, thereby, support investment decision making. Using attributes such as option values and potentially knowledge capital can make it possible to translate the intuitive appeal of R&D and human effectiveness investments into more tangible measures of value.

Note also that the methodology includes many of the elements necessary to developing a business case for HSI investments. Markets (stakeholders), revenues (benefits), and costs are central issues in business case development and in this methodology. However, this methodology also supports valuation of investments with broader constituencies (e.g., the public) and ranges of issues (e.g., jobs created) than typically considered in business cases.

Finally, we have also found that use of the methodology presented here provides indirect advantages in terms of causing decision-making groups to clarify and challenge underlying assumptions. This helps decision makers avoid being trapped by common delusions that would mislead them relative to likely costs/benefits (Rouse, 1998).

REFERENCES

Allen, M. S. (2000). *Business Portfolio Management: Valuation, Risk Assessment, and EVA Strategies*. New York: Wiley.

Amram, M., and Kulatilaka, N. (1999). *Real Options: Managing Strategic Investment in an Uncertain World*. Boston: Harvard Business School Press.

Andrews, D. H. (2000). Distributed Mission Training. In W. Karwowski (Ed.), *International Encyclopedia of Ergonomics and Human Factors*. Philadelphia, PA: Taylor & Francis.

Black, F., and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 637–659.

Boer, F. P. (1998, September/October). Traps, Pitfalls, and Snares in the Valuation of Technology. *Research Technology Management*, pp. 45–54.

Boer, F. P. (1999). *The Valuation of Technology: Business and Financial Issues in R&D*. New York: Wiley.

Brigham, E. F., and Gapenski, L. C. (1988). *Financial Management: Theory and Practice*. Chicago, IL: Dryden.

Burke, J. (1996). *The Pinball Effect: How Renaissance Water Gardens Made the Carburetor Possible and Other Journeys through Knowledge*. Boston: Little, Brown.

Christensen, C. M. (1997). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston: Harvard Business School Press.

Cooper, R. G., Edgett, S. J., and Kleinschmidt, E. J. (1998a). *Portfolio Management for New Products*. Reading, MA: Addison-Wesley.

Cooper, R. G., Edgett, S. J., and Kleinschmidt, E. J. (1998b). Best Practices for Managing R&D Portfolios. *Research Technology Management*, *41*(4), 20–33.

Davenport, T. O. (1999). *Human Capital: What it is and Why People Invest It*. San Francisco: Jossey-Bass.

Economist. (1999, June 12). A Price on the Priceless: Measuring Intangible Assets. *Economist*, pp. 61–62.

Earned Value Management Center. (2000). Available: http://evms.dcmdw.dla.mil.

Fuchs, R., McCarthy, J., Corder, J., Rankine, R., Miller, W., and Gawron, V. (1997). *United States Air Force Expeditionary Forces*. Washington, DC: Air Force Scientific Advisory Board.

Gill, B., Nelson, B., and Spring, S. (1996). Seven Steps to New Product Development. In M. D. Rosenau, Jr. (Ed.), *The PDMA Handbook of New Product Development* (Chapter 2). New York: Wiley.

Gramlich, E. M. (1997). *A Guide to Cost-Benefit Analysis*. Prospect Heights, IL: Waveland.

Hammond, J. S., Keeney, R. L., and Raiffa, H. (1998). *Smart Choices: A Practical Guide to Making Better Decisions*. Boston, MA: Harvard Business School Press.

Keeney, R. L., and Raiffa, H. (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: Wiley.

Kirkwood, C. W. (1979). Pareto Optimality and Equity in Social Decision Analysis. *IEEE Transactions on Systems, Man, and Cybernetics*, *9*(2), 89–91.

Layard, R., and Glaister, S. (Eds.). (1994). *Cost-Benefit Analysis*. Cambridge, UK: Cambridge University Press.

Lint, O., and Pennings, E. (1998). R&D as an Option on Market Introduction. *R&D Management*, *28*(4), 279–287.

Luehrman, T. A. (1998, July/August). Investment Opportunities as Real Options. *Harvard Business Review*, pp. 51–67.

Matheson, D., and Matheson, J. (1998). *The Smart Organization: Creating Value through Strategic R&D*. Boston, MA: Harvard Business School Press.

Mintz, S. L. (1998, February). A Better Approach to Estimating Knowledge Capital. *CFO*, pp. 29–37.

Nevins, J. L., and Winner, R. I. (1999, April). *Ford Motor Company's Investment Efficiency Initiative: A Case Study*, Paper P-3311. Alexandria, VA: Institute for Defense Analyses.

Office of Science and Technology Policy (OSTP). (1998, August). *A National Obligation: Planning for Health Preparedness for and Readjustment of the Military, Veterans, and Their Families after Future Deployments*, Presidential Review Directive No. 5. Washington, DC: Office of Science and Technology Policy.

Rastikis, L. (1998). Human-Centered Design Project Revolutionizes Air Combat. *CSERIAC Gateway*, *XIV*(1), 1–6.

Rouse, W. B. (1991). *Design for Success: A Human-Centered Approach to Designing Successful Products and Systems*. New York: Wiley.

Rouse, W. B. (1998). *Don't Jump to Solutions: Thirteen Delusions That Undermine Strategic Thinking*. San Francisco: Jossey-Bass.

Rouse, W. B., and Boff, K. R. (1997). Assessing Cost/Benefits of Human Factors. In G. Salvendy (Ed.), *Handbook of Human Factors and Ergonomics* (Chapter 49). New York: Wiley.

Rouse, W. B., and Boff, K. R. (1998). R&D/Technology Management: A Framework for Putting Technology to Work. *IEEE Transactions on Systems, Man, and Cybernetics*—Part C, *28*(4), 501–515.

Rouse, W. B., Boff, K. R., and Thomas, B. G. S. (1997). Assessing Cost/Benefits of R&D Investments. *IEEE Transactions on Systems, Man, and Cybernetics—Part A*, *27*(4), 389–401.

Rouse, W. B., Kober, N., and Mavor, A. (Eds.). (1997). *The Case for Human Factors in Industry and Government*. Washington, DC: National Academy Press.

Roussel, P. A., Saad, K. N., and Erickson, T. J. (1991). *Third Generation R&D: Managing the Link to Corporate Strategy*. Cambridge, MA: Harvard Business School Press.

Sage, A. P. (1977). *Systems Methodology for Large-Scale Systems*. New York: McGraw-Hill.

Sage, A. P., and Armstrong, J. (2000). *An Introduction to Systems Engineering*. New York: Wiley.

Sage, A. P., and Rouse, W. B. (Eds.). (1999). *Handbook of Systems Engineering and Management*. New York: Wiley.

Strassman, P. A. (1999, September). *Does Knowledge Capital Explain Market/Book Valuations*? Available: www.strassman.com/pubs/km.