Draft July 1998

BENEFIT-COST ANALYSIS GUIDE

Treasury Board of Canada Secretariat

Ottawa

Table of Contents

Note: The Guide follows the steps of a benefit-cost analysis, from defining the problem and fair comparisons, to measuring costs and benefits, to dealing with uncertainty and risk. The final chapter discusses what should be done when the analyst must step outside the benefit-cost framework to consider other criteria.

Each chapter ends with a short summary of best practices and the Guide itself ends with a general summary. There is a glossary of terms, and terms flagged in bold type in the text can be found in that glossary. Terms in bold italic are defined where they are introduced. At the end of the Guide there is a bibliography of selected readings, arranged by topic.

Preface

This Guide provides a framework for benefit-cost analysis. It should be used in submissions to the Treasury Board when the matter has significant social, economic or environmental implications. A sound benefit-cost analysis should be at the heart of every business case presented to senior managers and to ministers. The new Guide is part of the Government of Canada's focus on evidence-based and analytical decision making.

However, the Guide is not the last word on benefit-cost analysis. It is written to be a useful tool for economists and non-economists alike. It provides the essential framework and, where possible, indicates best practice; however, it does not replace the need for training in assessing costs and benefits.

For managers

Managers can use the Guide to help them design and commission benefit-cost analyses and improve their understanding of the findings. Consistent application of the techniques set out in the Guide is key. Both managers and the Treasury Board need to be able to compare alternative courses of action in a standard and rigorous way if scarce resources are to be used to best advantage.

For analysts

This Guide is meant to be an authoritative statement of how a benefit-cost analysis should be undertaken for the Government of Canada. It replaces the previous Treasury Board *Benefit-Cost Analysis Guide* which was published in 1976. It provides a consistent framework for comparative analysis but, of course, it does not cover every aspect of the measurement of benefits and costs. Nor does it replace the professional expertise, which the analyst and the manager bring to each case. There is no cookbook for good decision making - a good framework is necessary, but insight is equally necessary.

Acknowledgements

This Guide was developed under the direction of two inter-departmental committees led by the Treasury Board Secretariat. One committee comprised experts in the technical topics; and the other comprised general managers and potential users of benefit-cost analyses. The principal author of the Guide was Kenneth Watson, Ph.D., contracted to Consulting of Audit Canada. Staff of the Treasury Board Secretariat and of Consulting and Audit Canada made significant contributions; as did the members of the advisory committees. This Guide would not have been possible without the generous intellectual contribution of all involved.

1. Introduction

1.1 Resource-allocation decisions Involve choices

Managers and analysts throughout the government are often asked to provide analysis in support of resource-allocation decisions that affect the government and, perhaps to a greater degree, also affect resources outside government.

Difficult choices are involved when resources are scarce. There is also increasing recognition that governments must be careful how they draw upon and regulate private-sector resources. Even in the apparently simple ‘do it’ or ‘don’t do it’ choice, there may be compelling arguments in favour of the ‘don’t do it’ **option**. Any action that consumes resources that could be put to another, and perhaps better, use must have a powerful justification. Frequently, there are several alternatives. In the past, the option of proceeding with a program was usually set against some theoretical alternative use of funds. Today, governments are often forced to finance new programs at the expense of existing ones

Sometimes the decisions go all the way to Cabinet; this happens, for example, when legislation or regulations change, or a substantial program initiative is involved. More often, senior management in the department will settle issues such as allocation of resources within a program. Sometimes program heads make the decisions at lower levels. Regardless of who makes the decision, the principles to be used are the same; what varies is level of investment in analysis justified by the resources at stake. As the title of this guide suggests, the principles are those of **benefit-cost analysis**.

1.2 Who is this guide for?

This guide is intended for two groups:

- analysts who conduct studies in support of decisions; and
- managers who use the results of the studies.

Analysts in government are, for the most part, not economists, although most have some **economic** training. To the greatest extent possible, this guide has been designed for the larger population, rather than specialist. Economists, however, should find that the standard framework offered here makes it easier to compare projects from different sources and easier to communicate results to managers who are familiar with the framework.

The terminology in this guide tends to be drawn from economics. There may be same differences in the usage of the same terms in other fields. For example, in this guide, **risk analysis** normally deals with *any* uncertainty whether the uncertain factor is negative or positive, and whether the uncertainty is in the probability of occurrence, in the magnitude of effect, or in the monetary value of the effect. In contrast, some fields tend to think of risk more narrowly as solely an *adverse* factor with emphasis on the probability of occurrence. Normally, the context will make it clear what usage is being followed. Although attempting to span wide areas within the policy-analysis community has potential pitfalls, it is evident that the guide will meet its intended purpose only if it reaches out to a wide readership. To facilitate this, the authors have defined terms as they arise and have provided a glossary (see Appendix A).

In summary, the guide has the following objectives:

- to provide an understanding of how benefit-cost analysis can help in decision-making ;
- to establish a general framework that will lend consistency to analyses, facilitating their comparison and ensuring good practices whether the analyses are performed by departmental specialists or by consultants;
- to serve as a self-instruction manual with concrete and detailed guidance on the basic elements of analysis; and

- to help analysts and managers determine when more sophisticated analysis might be required than can be generated internally and to standardize expectations about what work departmental specialists or consultants will provide.

1.3 Why benefit-cost analysis?

Some think of benefit-cost analysis as a narrow financial tool. However this underestimates its versatility in addressing intangible values. Recent methodologies can help to estimate the value to Canadians of intangible benefits. At least we can often set a clear **cost** estimate against alternative ways of achieving an intangible benefit such as fairness in our immigration program.

Choices that confront policymakers have to be made. Quantitative analysis of probable outcomes of alternative courses of action can diminish the uncertainty and improve the decision-making process.

The current situation, in which programs seem to be perpetually under review, will probably persist for some years. Interest payments on the national debt now swallow such a large proportion of government revenues that funds available for program expenditures have shrunk considerably. Many programs that continue to offer useful services and **outputs** are being cut. Ministers are faced with difficult decisions; it is up to analysts to provide the most solid basis for those decisions.

But what criteria should be used?

Brief reflection on the question suggests that tax money should go to support programs where it will do the most good given the choices available. Defining *does the most good* and *given the choices available* captures the essential focus of benefit-cost analysis.

The basic elements are benefits, costs and choices. It is not a long step from ‘doing the most good’ to ‘creating the greatest (net) benefit.’ The same resources can not be committed to different ends. With a limited budget, we must to be certain that each project chosen has the largest possible value per dollar expended.

1.4 Where benefit-cost analysis fits into the decision-making process

New initiatives, especially those requiring legislation, go to Cabinet or, more commonly, a committee of Cabinet. There is a guide to drafting memoranda to Cabinet that specifies the general framework for the analysis necessary. It focuses on identifying the implications for particular segments of society and the evaluation of each option. Net benefit is not the only concern: **distributional effects** are often important.

Major redirection of programs often requires a submission to the Treasury Board, even when legislative change is not required and sufficient money already exist in the budget. Expenditure authorisations are usually tied to the performance of certain activities (not just overall goals), so major modifications have to come back to the Board for approval. Major Crown Projects (MCPs) and expenditures in any department that exceed defined limits also go before the Board. The Treasury Board *Submissions Guide* and the MCP policy both refer to the need for justification in terms of benefit-cost analysis. The present guide is a reference that not only sets down the basics of such analysis, but also establishes reporting conventions to ensure greater comparability of programs.

The annual process of approving department’s expenditure plans has been undergoing substantial change. Over time, the Government of Canada has put more emphasis on performance assessment. In the mid-nineties, the government introduced new instruments that have performance assessment as a key component: the departmental business plan (submitted to Treasury Board and Cabinet); the departmental performance report (submitted to Parliament); and the Treasury Board President’s report on performance review in the government (submitted to Parliament).

The business plans set out the departments’ strategies, objectives and performance commitments. For some departments, these commitments involve serious adjustments, because of the changes to size, scope and strategy they have had to make. The business plans also set out the departments’ commitments to review their major projects, programs and structural or resource changes. The Treasury Board of Canada Secretariat recommends that departments submit these reviews in the benefit-cost analysis format described in this guide.

For decisions that do not involve important policy issues, do not exceed ministers' delegated authority, or can be made at lower levels (for example by an assistant deputy minister), the use of this guide is still good practice. Many departments are adopting the language of the business case to analyse whether the expected return is worth the effort (benefits are greater than costs).

2. The benefit-cost analysis model

2.1 Introduction

Benefit-cost analysis is simply rational decision-making. People use it every day, and it is older than written history. Our natural grasp of costs and benefits is sometimes inadequate, however, when the alternatives are complex or the data uncertain. Then we need formal techniques to keep our thinking clear, systematic and rational. These techniques constitute a **model** for doing benefit-cost analysis. They include a variety of methods:

- identifying alternatives;
- defining alternatives in a way that allows fair comparison;
- adjusting for occurrence of costs and benefits at different times;
- calculating dollar values for things that are not usually expressed in dollars;
- coping with uncertainty in the data; and
- summing up a complex pattern of costs and benefits to guide decision-making.

It is important to keep in mind that techniques are only tools. They are not the essence. The essence is the clarity of the analyst's understanding of the options.

2.2 The benefit-cost analysis framework

Even when the measurements of costs and benefits are complete, they might not speak for themselves until they are put in a framework. Benefit-cost analysis provides that framework. It can be used wherever a decision is needed and is not limited to any particular academic discipline, such as economics or sociology, or to any particular field of public or private endeavour. It is a hybrid of several techniques from the management, financial and social sciences fields.

As far as possible, benefit-cost analysis puts both costs and benefits into standard units (usually dollars) so that they can be compared directly. In some cases, it is difficult to put the benefits into dollars, so we use **cost-effectiveness analysis**, which is a cost-minimization technique. For example, there might be two highway-crossing upgrade options that will result in the same saving of lives. In this case, we choose between the options on the basis of minimum cost.

> The feature that distinguishes benefit-cost analysis from cost-effectiveness analysis is the attempt benefit-cost analysis makes to go as far as possible in quantifying benefits and costs in money terms. However, benefit-cost analysis rarely achieves the ideal of measuring all benefits and costs in money terms ... so the distinction is merely a difference in degree and not in kind.
>
> - Treasury Board, *Benefit-Cost Analysis Guide*, 1976

2.3 The steps in benefit-cost analysis

There is no 'cookbook' for benefit-cost analysis. Each analysis is different and demands careful and innovative thought. It is helpful, however, to have a standard sequence of steps to follow. This provides consistency from one analysis to another, which is useful to both the analysts doing the study and the managers reading the report.

> Obviously, the ... steps cannot be performed by the analyst in isolation and will require consultations with the decision-maker and others, the gathering of a wide variety of

> information, and the use of a number of analytical techniques. It is important that, as the analyst proceeds, the decision-maker is kept in touch with the form of the analysis and the assumptions being made.
>
> - Treasury Board, *Benefit-Cost Analysis Guide*, 1976

A set of standard steps is listed below. Each step is explained in the chapter indicated.

1. Examine needs, consider constraints, and formulate objectives and targets. State the point of view from which costs and benefits will be assessed. (See this chapter.)

2. Define options in a way that enables the analyst to compare them fairly. If one option is being assessed against a **base case**, ensure that the base case is optimised. (See Chapter 3.)

3. Analyze **incremental** effects and gather data about costs and benefits. Set out the costs and benefits over time in a spreadsheet. (See Chapter 4.)

4. Express the cost and benefit data in a valid standard unit of measurement (for example, convert **nominal dollars** to **constant dollars**, and use accurate, undistorted prices). (See Chapter 5.)

5. Run the **deterministic model** (using single-value costs and benefits as though the values were certain). See what the deterministic estimate of **net present value (NPV)** is. (See Chapter 6.)

6. Conduct a **sensitivity analysis** to determine which variables appear to have the most influence on the NPV. Consider whether better information about the values of these variables could be obtained to limit the uncertainty, or whether action can limit the uncertainty (negotiating a labour rate, for example). Would the cost of this improvement be low enough to make its acquisition worthwhile? If so, act. (See Chapter 7.)

7. Analyse risk by using what is known about the ranges and **probabilities** of the costs and benefits values and by simulating expected outcomes of the investment. What is the **expected net present value (ENPV)**? Apply the standard **decision rules.** (See chapters 8 and 9.)

8. Identify the option, which gives the desirable distribution of income (by income class, gender or region - whatever categorisation is appropriate). (See Chapter 10.)

9. Considering all of the quantitative analysis, as well as the qualitative analysis of factors that cannot be expressed in dollars, make a reasoned recommendation.

This sequence is the preferred way to structure the benefit-cost analysis report.

2.4 Why is a point of view important?

A good way to start a discussion of benefit-cost analysis is by noting that the benefit-cost analyst must work consistently from a clear point of view. Whose costs and benefits are being assessed? The analyst is not restricted to a single point of view. The government might take the narrow fiscal point of view, for example, or a broad social point of view, or both. Whatever the point of view chosen, each analysis must take a single point of view and it must be stated clearly at the outset.

It is obvious that a cost from one person's point of view can be a benefit from another's. What is obvious when stated, however, is sometimes not obvious in the midst of an analysis. It is not at all uncommon to see lists of benefits or costs that are apples and oranges as far as a consistent point of view is concerned. Should taxes levied be counted as a benefit or a cost? Should jobs created be considered a benefit or a cost to the project? The answers depend on the point of view.

If there is a single decision-maker, then an analysis from one point of view is often adequate. If the interests of more than one person or group are affected, then several analyses might be necessary. Consider the decision to construct a recreational facility in a national park. The analyst who wants to provide advice to his or her minister might need to know how the project would look from the general social point of view (all costs and benefits to Canadians), from the fiscal point of view of the park authority, from the provincial point of view, and from the point of view of local environmental groups.

The point of view defines the ‘in group’ and the ‘out group.’ The in group consists of those people whose costs and benefits are to be taken into account in the analysis. For example, suppose that the in group comprises all the citizens of a town called Bin. In that case, if some of the resources of the Bin citizens are used up, there is a cost to be counted. If some of their resources are given to people outside Bin, there is a cost to be counted as well. If one citizen of Bin, however, gives resources to another citizen of Bin without anything being used up, then the total resources of Bin citizens are not affected and no cost or benefit is to be counted. A **transfer payment** has been made (see Section 4.2.1).

As well as identifying costs and benefits correctly, one must choose *parameters that are consistent with the point of view* of the analysis. For example, the appropriate **discount rate** depends on what point of view is being taken in the analysis (see Section 5.5).

2.5 The components of benefit-cost analysis

All public-investment decisions can be modelled in the same standard way, using as the general framework for analysis the same four components:

- a parameters table;
- an incremental-effects model;
- a table of costs and benefits over time; and
- a table of possible investment results and a statistical and graphical analysis of NPV and investment risk.

These components are depicted in Figure 2.5.1.

Figure 2.5.1: The general flow of benefit-cost analysis

Parameters table	
Parameter 1: Population growth rate	2% per annum
Parameter 2: Social discount rate	10% per annum
Parameter 3: Apple price increase	5% per annum

Incremental effects model		**Period**	
	t_0-t_1	t_1-t_2	t_2-t_3
Events	___	___	___
Consequences	___	___	___

Table of costs and benefits over time (simplified illustration)

	Costs		Benefits			
Period	**Materials**	**Labour**	**Sales**	**Net (nominal $)**	**Net (constant $)**	**Present values**
t_0-t_1	($100)	($67)	$40	($127)	($124)	($113)
t_1-t_2	($212)	($34)	$90	($156)	($148)	($123)
t_1-t_3	($455)	($84)	$600	$67	$57	$43
					NPV =	**($193)**

Investment results table

NPVs for various combinations of input values

($785)
($768)
($700)
.
.
.
$350
$369
$382

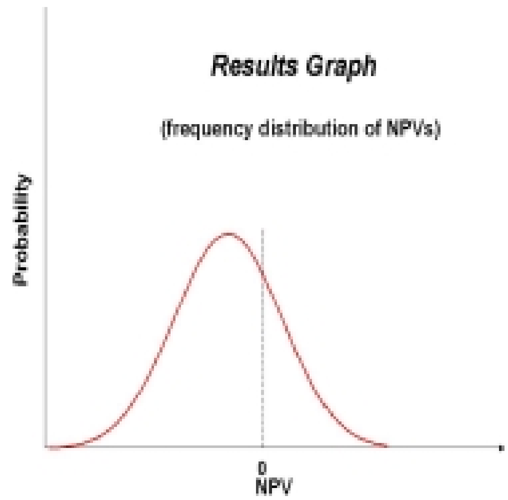

The first component of this investment model is the *parameter table*, which is a list of variables used to calculate the costs and benefits. For example, both costs and benefits of a project might be influenced over time by the population growth rate of the community. Rather than retype the population growth rate every time it appears in a formula within the table of costs and benefits, it is better to list it in the parameter table and refer to it in other parts of the spreadsheet as it is needed. Although not absolutely essential, the use of a parameter table also facilitates all kinds of 'what if' analyses, including sensitivity analysis and risk analysis (see Chapters 7-9). It simplifies the analyst's task when changing the value of the parameter, a key requirement of risk analysis. Instead of searching through the whole model for all the places where the population growth rate was used (and perhaps missing some!), the analyst can change the value in the parameter table, and all its uses in the benefit-cost model will change automatically and simultaneously.

The second component is the *incremental-effects model*. In business or industrial contexts, this is sometimes called the production model. It sets out the expected events and consequences over time. The nature of the events will depend on the project - from illnesses (an immunization project), to sales (an export-incentive project), to letters sorted (a post office **capital**-investment project). These events are often subject to uncertainty, so we tie them in with the parameter table in the same way that we tie in the table of costs and benefits later.

The third component of the model is the *table of costs and benefits over time.* This is a list of all costs and benefits, with the values for each noted for every period within the **investment horizon.** These values are best-expressed in nominal dollars so that adjustments normally calculated in nominal dollars can be made (adjustments for taxes, for example). Nominal dollars cannot be added or subtracted across periods, however, so they must at some stage be converted to constant dollars, and then to **present values**, before they can be summed up. (For details on constructing a table of costs and benefits, see Section 2.6.) There are two ways of doing this. The first way is to calculate the full table of costs and benefits in nominal dollars, then another table in constant dollars, and then in present values. The second way is a little easier and more concise: the analyst adds all benefits and subtracts all costs within each period to obtain a single nominal-dollar net for each period and then converts this nominal-dollar net cash flow to constant dollars and present values (by convention, the analyst is allowed to add and subtract nominal dollars within a single period, although this is an approximation of true value because the worth of a dollar might change if the period is lengthy). (Nominal and constant dollars are discussed further in Chapter 5.)

The final component of the model is the *investment results table*. Each time the benefit-cost model is run, it estimates an NPV of the investment. If it is a deterministic model, in which all the **inputs** have fixed values, then the result of each run will always be the same NPV. If it is a risk-analysis model, in which the parameters' values vary within a stated range according to probabilities, the estimated NPV will also vary. The result of many runs of the model will be a list of possible NPVs, and this list itself will have to be analyzed statistically to determine the probable true NPV. This statistical analysis will show the maximum and minimum values of the NPV and the probabilities that the NPV is within various ranges. With this information, the analyst can apply decision rules to ascertain whether the project is a good one and whether it is the best alternative.

2.6 Constructing tables of costs and benefits

By far the greatest amount of time in a benefit-cost analysis is spent in constructing the tables of costs and benefits over time. To construct these tables, the analyst identifies the full set of relevant costs and benefits, estimates their quantities for each period, and calculates their values by applying their prices to their quantities in each period. This needs to be done carefully. A benefit-cost analysis is no better than its data.

There are no shortcuts. It is seldom, if ever, accurate to construct one year of costs and benefits and to assume that this year is repeated 25 times in constant dollars out to investment horizon t_{25}. The world does not work like this. Not only do prices change, but also relative prices change — land becomes more expensive, computing power becomes cheaper, commodities follow a price cycle and so on. One benefit-cost analyst might not have the expertise to estimate all the quantities and prices needed in the analysis and may have to draw upon other specialists for data estimates.

In some cases, the analyst will be working with pro forma financial operating statements for a proposed program or project, or with business income and expense statements. These sorts of data often require some adjusting to fit a benefit-cost framework. One financial framework is not better than another is. Each has its own internal consistency, but data from one framework might not fit in another.

The greatest difference between benefit-cost **cash flows** and business cash flows is that the latter may include accrued values, **depreciation** and similar allowances. Benefit-cost analysis does not use accruals, depreciation allowances or other 'non-cash' items. In benefit-cost analysis, each cost and benefit is fully recognised at the time it occurs (not accrued beforehand), timing is dealt with through **discounting**, and changes in the values of **assets** are dealt with by including **residual values** at the investment horizon.

In benefit-cost analysis, accounts receivable and payable are not recognized until the cash is actually received or paid. Working capital is not a cost, although the change in working capital during a particular period is either a cost (if working capital decreases) or a benefit (if it increases). Production costs are recognized fully at the time they occur. Changes in inventory may signal either costs or benefits, but the actual measurement of these is through production costs and sales. Benefit-cost cash flows are simple tables with everything recognized when it occurs. Although this is a simple concept, it can be uncomfortable at times for a financial officer who is used to accrual accounting.

2.7 Accounting conventions

Once you are familiar with the general form of the benefit-cost model, it is useful to think about the conventions that are used to set it up in a standard way. One convention is not necessarily better than another is, but standardization is needed if the model is to be a general tool for comparison. Conventions are important for many aspects of the model, such as the investment horizon, time-of-occurrence assumptions, and the **numeraire** — a common unit of measurement.

2.7.1 The investment horizon

The investment horizon is the end of the period over which costs and benefits will be compared to ascertain whether an investment is a good one. If costs and benefits can be identified for the whole economic life of the project and uncertainties are low, then the full economic life provides the best investment horizon. If not, there may be logical points in the economic life of the project at which to terminate the investment analysis. For instance, there may be a point at which relatively certain costs and benefits are suddenly followed by much less certain values. Recapitalization of a building, for example, tends to occur in cycles: 5-7 years for repainting and new carpets; 15-17 years for service systems such as heating; and 25-50 years for major structural components. These thresholds where major new uncertainties occur might suggest an appropriate investment horizon. It is important, however, that the investment horizon not be chosen to deliberately favour the project. Even if you are considering a single option, then it might be advisable to analyse the project within various investment horizons to see whether a change in investment horizon affects the outcome. If you are comparing alternatives, you should use the same investment horizon in the analysis of each (see Section 3.3.1).

2.7.2 Time-of-occurrence assumptions

Costs and benefits occur at different points within the standard period being used (within the year, for example). Therefore, you need a convention for establishing where all costs and benefits will be assumed to fall within the period. Normally, the analyst selects one of three possibilities: either at the beginning of each period, in the middle, or at the end. Underlying this practice is the need to have a reasonably simple pattern of costs and benefits over time so that changing nominal dollars to constant dollars, and to present values, is not too difficult.

Figure 2.7.2 : Time of occurrence assumptions

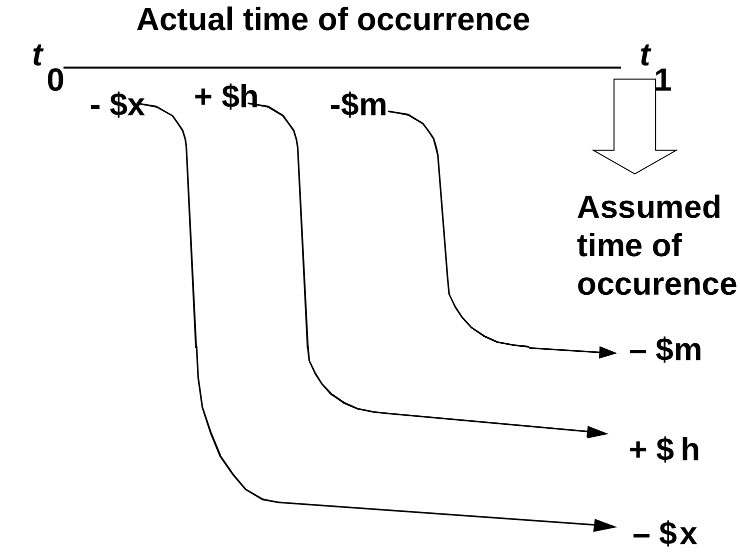

This practice assumes that if all benefits and costs are shifted to the same point within each period, then, on balance, the overall outcome will not be affected. This is generally a reasonable assumption, except in cases where there is an unusually large cost or benefit being shifted a long way - say from near the beginning of a period to the end. A common example of this is a large, initial lump-sum investment; an analyst using end-of-period accounting will assume that this investment occurs at the end of the first year (and therefore is adjusted for inflation once and discounted once), when in fact it occurs at the beginning. This artificial adjustment of a large cost can make a major difference to the outcome of the analysis if there is no comparable benefit being similarly adjusted in the same model.

This sort of problem has led to hybrid conventions. For instance, in a table of costs and benefits, common convention is to assume that the costs and benefits listed in the first column of numbers occurred at a single point in time, t_0, rather than in a period. The next column lists costs and benefits for the period t_0-t_1. Then, if t_0 is selected as the base point in time, as is commonly done, the initial lump-sum investment is unchanged by adjustments for inflation or discounting. We recommend this procedure when there is a large initial investment.

It is important to know which period conventions your software uses so that you can avoid inappropriate calculations. Each benefit-cost report you submit should state the period convention used.

2.7.3 The numeraire - a common unit of value

Before they can be summed up, all costs and benefits must be expressed in a common unit of value. This involves three main things: expressing them all in a common numeraire (say Canadian dollars of investment funds); adjusting for **inflation** where necessary (converting to constant dollars); and expressing all in present values (adjusting for differences in the time of occurrence of costs and benefits).

The costs and benefits must be in a common monetary unit before they can be compared. Most investment analysis uses a dollar of investment as the unit of measurement. However, some public-sector models use a dollar of consumption or a dollar of foreign exchange as the numeraire. All are acceptable, but clarity and consistency are essential. If price **distortions** are widespread in a particular economy, the benefit-cost analysis may use border prices or world prices as the numeraire, that is, as the best measure of true value of costs and benefits. For most purposes of the Government of Canada, a dollar of investment expressed in Canadian prices is an adequate numeraire. It also has the advantage of being easy to understand.

Both constant dollars and present values (they are *not* the same things!) are defined at a particular point in time. Any point will do, but the most frequent choices for t_0 are the time at which the analysis is being done, the start of the project, or the start of a new fiscal year. Costs that are incurred before t_0 would, of course, be inflated, rather than deflated, to an equivalent t_0 value.

2.8 The report

In general, reports should contain at least the following:

- a description of the need, problem or opportunity;
- a description of the options with an explanation of why they were chosen and why it is fair to compare them;
- a statement of the point of view of the analysis;
- a statement of assumptions and **scenarios;**
- a deterministic analysis;
- a cost-benefit analysis and a risk analysis;
- a discussion of equity effects and other non-economic effects; and
- a ranking of the options.

Best practice - the general benefit-cost model

- Benefit-cost analysis can be applied to a wide range of decisions made by the Government of Canada.
- Every benefit-cost analysis must state the point of view from which benefits and costs will be assessed.
- There is no cookbook for benefit-cost analysis, but a standard set of steps is a useful starting point.
- Each benefit-cost analysis should contain a parameter table; an incremental-effects model; a table of costs and benefits over time; a table of possible investment results; and a statistical and graphical analysis of expected NPV and investment risk.

3. Defining fair options

Unfair comparisons are odorous.

- Shakespeare, *Much Ado About Nothing*

3.1 Why are fair options difficult to define?

Two requirements must be met if a comparison of alternatives in a benefit-cost analysis is to be fair. The first is to ensure that *all* the relevant alternatives are considered. The proposed investment should be compared with the best alternative uses of the resources. It is not enough to assume that a proposed project has only to make a return equal to the discount rate to be acceptable. There may be alternative projects that would do even better!

The second requirement is that the alternatives to be compared must each be defined in a consistent and fair way. In particular, one cannot make a simple comparison between two investment alternatives if they are at different **scales**, occur at different times, or involve different ownership. This is sometimes not obvious from some texts that imply that any alternatives, however structured, can be compared simply by calculating their NPVs. This is not the case. There is an underlying assumption that may not be valid.

The dubious assumption is that unused resources earn a normal **rate of return** (that is, they earn a rate of return equivalent to the discount rate). On this assumption their present value will be zero (for example, if the unused asset earns 10 per cent per annum and is discounted by 10 per cent per annum, then it is a 'wash' as far as NPV is concerned). Consequently, adding or subtracting the unused resources cannot affect the total present value of any investment under consideration. For example, consider investment *A*, which costs $70 and has an NPV of $30, and investment *B*, which costs $200 and has an NPV of $31. According to the NPV decision rule (see Section 6.1), investment *B* is better. However, this assumes that the $130 we have left over if we invest in *A* has an NPV of zero and therefore doesn't affect the investment decision. This assumption is implicit, in a sense hidden, and might not be valid.

In real situations, the resources left over if one chooses the smaller investment or an investment further in the future may not earn a rate of return equivalent to the discount rate. Indeed, it is not uncommon in government for unused resources to have a negative return in the short run because of the holding costs. Consider the case where a department is deciding whether to renovate an older building it already owns or to lease a new building. If it decides to lease, what will happen to the older building? The traditional NPV decision rule assumes that the resources tied up in the old building (essentially its existing market value) will earn a return equal to the discount rate. Of course, this is unlikely. It is more likely that the building will continue to sit empty, waiting for some other use, and an NPV calculation that ignores this reality is likely to give poor guidance to decision makers.

3.2 An optimised base Case

It is important to identify the set of most promising options. To do this, the benefit-cost generalist works with experts in the substantive field. For example, if the Government of Canada wanted to improve winter transit time for freighters in the St. Lawrence Seaway, then the options would be identified mainly by experts in marine transportation, ice-conditions identification, communications and related technical fields. The focus is on what will or will not work, putting aside for the moment the question of financial attractiveness. No option should not be eliminated at this stage of the benefit-cost analysis on the grounds of politics or equity before its net economic values are known.

When you are making a 'go' or 'no go' decision on a single-project proposal, it is important to optimise the baseline (without project) case before you calculate the incremental costs and benefits of the proposed project. Failure to do this has been a common source of error. The status quo is not necessarily the appropriate baseline. It is often possible to improve results without major capital investment. If so, it is the improved status quo that provides the baseline against which you should measure the proposed project. In other words, it is not 'before' or 'after' that we are interested in but rather the best 'with' and the best 'without' the project.

When you are identifying the alternatives to be analysed, keep in mind that public-investment decisions share three important characteristics. First, to some extent, these decisions may be *irreversible*; once committed, the resources cannot be recovered. Second, the outcome of the investment may be uncertain because the input data are *uncertain*. Third, there may be some leeway with respect to the *timing* of the investment.

We will consider the issue of uncertainty in Chapters 7 to 9. The issues of irreversibility and timing, which were known to analysts in the past but not given much attention, have become important considerations in the 1990s. The ability to wait before making an irreversible investment is important. A simple case is an initial investment that gives rise to benefits that increase over time, such as an investment in a new road where the traffic is expected to increase. Investing immediately might have a positive NPV, but this might mask negative NPVs in the early years because the benefits of later years predominate. In this case, waiting for the optimum moment to invest is important. Generally, keeping your alternatives open while waiting for new information that might affect your decision gives you an option that may be valuable indeed.

3.3 How to construct fair options

The only way you can be sure that the options whose present values are being compared are really fair alternatives is to standardize them for time, for scale and for already-owned components.

3.3.1 Standardize the options for timing

If there are two investment alternatives with different time frames, they must be standardized normally by choosing the longer time frame for both. If one project starts earlier and the other finishes later, then the earlier start and the later finish normally define the standardized time frame. All resources need to be accounted for, in all alternative time frame projects, for the full timeframe. Sometimes, indivisible components determine how the timeframe can be standardised. For example, if you are choosing between making a gravel road (surface life of 6 years) or a making a blacktop road (surface life of 15 years), what investment horizon would allow a fair comparison? The following are two possibilities:

Comparison 1: 2 applications of blacktop vs. 5 applications of gravel (30-year horizon)
Comparison 2: 1 application of blacktop vs. 3 applications of gravel (18-year horizon)

The first seems to be a fair comparison: (2 ? 15) = (5 ? 6). The second does not: (1 ? 15) ≠ (3 ? 6). What happens to the blacktop road after its surface life expires (years 16 to 18)?

3.3.2 Standardize the options for scale

Standardizing the options for scale is similar to standardizing the options for time, as illustrated above. If you have two investment alternatives involving different levels of investment, then you must specifically account for the resources left over after making the smaller investment rather than just assuming they generate a zero NPV.

3.3.3 Standardize the options for already-owned components

If one investment option uses a resource that is already owned by the government, then the analyst must also show what would happen to this resource for each of the alternative investment options. For example, if the government owns a building and one alternative is to renovate it, then the analysis of each of the other alternatives must also incorporate what happens to that building. You cannot blindly assume that the building will earn any particular rate of return if it is not used for the proposed purpose. In fact, several mistakes are common. One mistake is to assume without examination that already-owned resources earn the standard rate of return (the discount rate). Another is to treat them as though they were costless. A third mistake is to consider them only in the one investment scenario where they are most relevant and assume they do not exist for the alternative scenarios. The rule is that any already-owned asset that appears in one alternative should appear in all alternatives, and the return it generates in each case should be examined specifically.

The correct way to treat already-owned assets is to recognize (in all the investment alternatives) their full **opportunity cost** at the beginning and at the end of the standardized investment period. Their opportunity cost is generally best measured by the net market value of the asset including, where relevant, sales costs, site clean-up costs, and the time required to sell the asset.

3.3.4 Fair options diagrams

In practice, it can be difficult to conceptualize the investment alternatives, standardizing for scale, for timing and for already-owned components. A diagram can clarify a complex set of investment alternatives. We call such a diagram a ‘fair options’ diagram.

An example is Figure 3.3.4, which depicts three options for accommodating 350 staff. Note that the options must all be at the same scale. The best measure of scale is that, although the four options do not encompass exactly the same floor space, each could accommodate 350 staff. All options cover the same period (four years), and all consider what will happen to already-owned assets.

Figure 3.3.4: An example of a fair options diagram

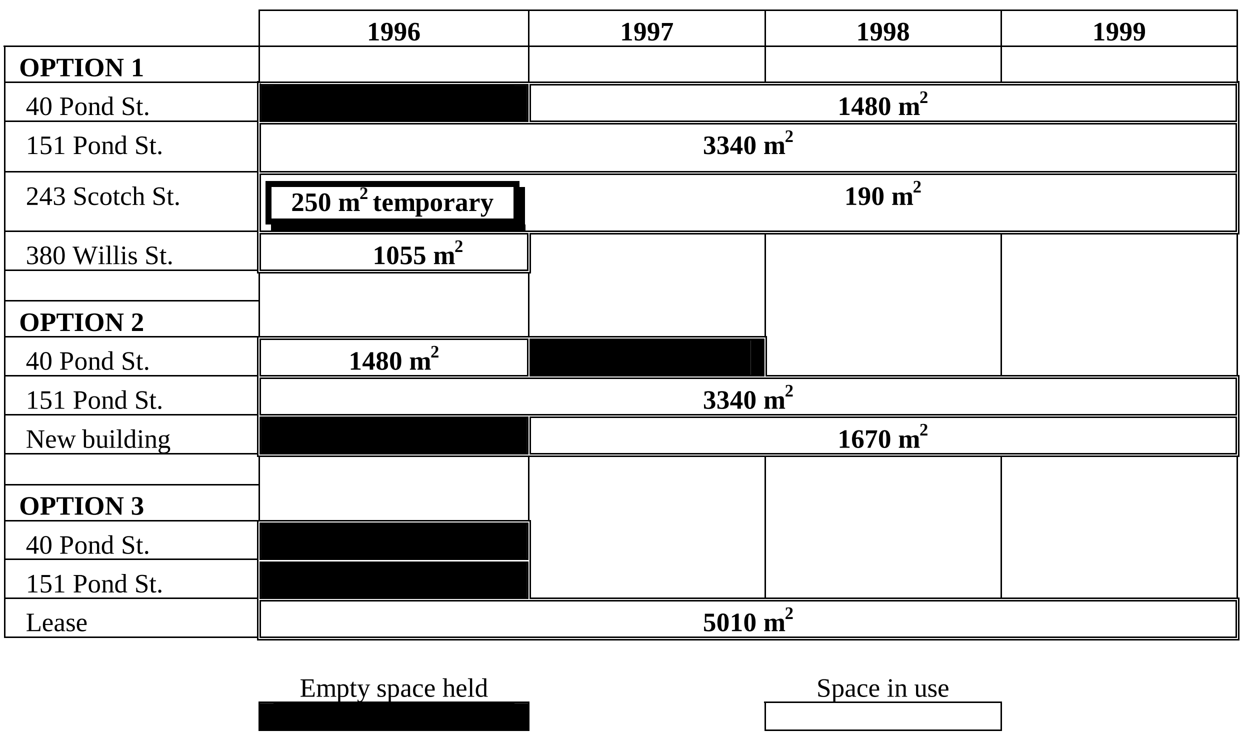

The three options shown in Figure 3.3.4 have been standardized in the following ways:

- the time frame is the same for each investment (1996-1999).
- the project size is the same (350 people accommodated).
- the already-owned asset (40 Pond St.) appears in all options.
- there is a plausible plan for handling already-owned assets that are not needed in a particular option.

3.3.5 Non-essential components of options

Options must be self-standing, as well as fair. That is, they must be complete and spare. ***Spare*** means there should not be anything in the option that is not essential to it. For example, suppose the federal government owned some land close to an airport and that land was expected to increase in market value over the next several years. At the same time, the government is considering whether to build a training centre. In this situation, the training facility may be a poor investment, but the overall NPV of the project may still look good because of the inclusion of the increasing value of the land. This is not legitimate. You must be sure that all components of an option are indeed essential; otherwise, the benefit-cost analysis may be misleading.

3.4 Incremental effects analysis

Before you can undertake a financial or economic analysis of a proposed project or program, you need a clear understanding of the incremental events and consequences to be expected. In general, you will need input from subject-matter experts. If the project is to decrease the risk of oil spills, the analysis team should include engineering and scientific experts to estimate the expected frequency of spills, assess the potential consequences of a spill for marine ecosystems, and to evaluate the extent to which the project will be able to prevent spills or contain damage once a spill occurs. Similarly, if the project is to build a road by-pass, the analysis team should include traffic engineers to estimate the incremental improvements in safety and travel time that will result. In benefit-cost analysis, then, two ‘subject matter’ skills will always be needed:

- expertise in estimating the expected frequency of events; and
- expertise in assessing the potential consequences of events.

The benefit-cost analyst brings two additional skills to bear on the information provided by the subject-matter experts:

- expertise in valuing outcomes in dollars; and
- expertise in making fair comparisons between benefits and costs.

One should not exaggerate the difference between the two sets of skills, however. Both the specialist and the generalist rely on similar analytical skills. The important point is that a team effort is often required for a full analysis.

Best practice - defining fair options

- For all public investments, a full set of the most promising options should be compared.
- When a single proposal is being considered, it must be compared with a ‘baseline case’ and the baseline case must be optimized.
- The option to delay a project to wait for better information, or for better starting conditions, can have considerable value.
- Each option must be standardized for scale, timing and already-owned components to permit a fair comparison. A fair options diagram can clarify a complex set of investment options.

4. Measuring and valuing costs and benefits

4.1 Introduction

This chapter explains how to measure costs and benefits to obtain information to enter into the benefit-cost framework. The general benefit-cost framework can be learned in a week's study, but the measurement of costs and benefits is a limitless topic and often requires a wide range of expertise. Because the coverage in this guide is necessarily brief, analysts will often have to consult more specialized documents for guidance on measurement (see the suggested readings in Appendix C).

Measuring costs or benefits and valuing them in dollars requires many different skills. For example, consider the case of a project to clean up river pollution from a manufacturing facility. An industrial chemist is needed to calculate the incremental change in the amount of pollutants entering the river; a biologist is needed to determine the effect of this change on bacteria in the river; a health scientist is needed, in turn, to evaluate the effects of that change on the health and recreation opportunities of residents; and then a benefit-cost analyst is needed to estimate in dollars the value of these benefits to the community. Note that the benefit-cost analyst does only one of the calculations, and not necessarily the most difficult.

4.2 Some Important Concepts

Even when we know how to count in standard units, we still need to be careful about what we count. In particular, incrementality, transfers, opportunity cost, sunk cost and residual value are important concepts in benefit-cost analysis. Only incremental benefits and costs caused by the project should be compared, not those are merely associated with the project in some way. For example, if one did a benefit-cost analysis of a government grants program to encourage exporters, one would need to know not just the export sales made, but specifically what sales were made that would not have been made in the absence of the program.

To avoid double counting, the analysis must maintain a consistent point of view. However, it is not the only requirement. The analysis team also needs an in-depth understanding of the proposed investment to be able to identify a coherent set of costs and benefits without double counting. For example, suppose a new sewage-treatment plant is installed. The recreation value of the river improves, land values in the neighbourhood increase, and health problems decrease. However, if all these effects are counted as benefits there is probably double counting. The increase in land values is probably a measure of the other benefits, not an additional benefit.

4.2.1 Transfers compared with true benefits and costs

In benefit-cost analysis we count resources that are created or used up. Resources that are simply transferred from one pocket to another are not counted as costs or benefits. For example, income taxes are transfers from the point of view of the whole country. Taxes move resources around, but, apart from administrative and disincentive costs, nothing is used up.

‘Point of view’ establishes whether a transaction is a transfer or not. It determines whether resources are passed from one pocket to another (a transfer) or passed out of the group or used up (a cost). From the point of view of a private business, for example, income taxes are definitely a cost.

In certain circumstances, tariffs, grants, taxes, social-welfare payments and many other items can be considered transfers. What is important here is whether resources are gained by or lost to the stakeholder(s) from whose point of view the analysis is being done.

4.2.2 Opportunity cost and sunk cost

> In calculating the benefits of public projects, the proper valuation to use is the price consumers are willing to pay for the output, that is, producer’s price plus taxes minus subsidies. In evaluating costs, the correct approach is less clear cut. Consider, for example, taxes and subsidies on intermediate inputs. Taxes increase the cost of inputs to users above the value of real resources expended in producing them, while subsidies have the opposite effect. In evaluating these costs, the proper measure critically depends on whether the project’s demand for the inputs is met by new supplies, or by diversions from other uses. If the inputs come from new supplies, the correct measure is the value of real resources expended, which is equivalent to the price paid by other users minus taxes plus subsidies. If the inputs are obtained by depriving other uses, the correct measure is the value of the inputs in alternative use, or the producer’s price plus taxes minus subsidies.
>
> - Treasury Board, *Benefit-Cost Analysis Guide*, 1976

The opportunity cost is the true value of any resource foregone. It must be counted even if explicit cash transactions are not involved. For example, if I could sell my computer for $1000 but instead I use it on a project, the opportunity cost of the computer (to be counted against the project) is $1000, although there is no cash transaction involved.

A cost is ‘sunk’ if it is irretrievably made or committed. A sunk cost is not to be counted in a prospective benefit-cost analysis because it cannot be affected by the decision in question. For example, if I originally paid $3000 for my computer but its market value at the time of the analysis is $1000, then $1000 is the opportunity cost if I decide to use the computer on a proposed project rather than sell it, and the remaining $2000 is a sunk cost that is no longer relevant.

4.2.3 Externalities

> An attempt should be made to take into account all of the allocative effects in evaluations of the efficiency of government expenditures, some of which may be less obvious than others... Such implicit effects may be internal (to direct actors in the project) or external (to persons not directly acting in the project but included in the group whose point of view is being taken in the analysis). An example of internal implicit effects is foregone wages during education... External implicit effects (also referred to as spillovers, social effects, or third party effects) are commonly things like pollution or congestion... Ignoring implicit costs or benefits could lead to major errors in analysis.
>
> - Treasury Board, *Benefit-Cost Analysis Guide*, 1976

4.2.4 Residual value

A residual value is the value of an asset at the end of the investment horizon. For example, suppose you invest in a rental property. At the end of the investment horizon, the land is still a valuable asset. The residual value is a benefit to be counted when you appraise the project. In most cases, the residual value is the market value of the asset. However, governments often maintain ‘special-use facilities’ (research laboratories, for example) for which market value might not be a good measure. The value of a special-use facility may be as little as the market value of the land minus demolition costs to remove the buildings. On the other hand, the true value can be as high as the replacement costs of the buildings and the land.

In calculations of residual value for a benefit-cost analysis, the land and the buildings are often treated separately. The analyst uses an index to estimate the expected market value of the land. The analyst then estimates the economic life of the buildings and prorates the replacement value according to the percentage of economic life that will have passed by the end of the investment horizon. For example, suppose at t_0 a real property consists of land worth $1 million and buildings worth $2 million. By t_{10} (the investment horizon in this case), we expect the land value to have increased to $1.5 million (nominal dollars) and the replacement value of the buildings to have increased to $3.5 million (nominal dollars). Suppose also that 10 years is 50 per cent of the economic life of the building. The residual value of the real property at t_{10} would thus be approximately $1.5 million (land) plus $1.75 million (half the replacement value of the buildings).

Several problems can arise in treatments of residual values. One mistake is to count a residual value on an already-owned asset without counting the balancing opportunity cost at t_0. Whether the asset is already owned or not, its full value must be counted as a cost at t_0 if its residual value is to be counted as a benefit at t_n.

Another mistake is to make a conservative estimate of cost and a generous estimate of residual benefit. The way the cost is computed at t_0 must be comparable to the way the benefit is computed at t_n.

A third problem arises if the project itself is not defined correctly. Sometimes a non-essential component has a good residual value that masks a bad outcome of the essential components. When you are counting a residual value as a benefit, make sure that the asset in question really is an essential part of the project (see Section 3.3.5).

4.2.5 General administrative and overhead costs

When a large organization, like a government, analyzes many possible investments over time, it may have a problem deciding how to treat general costs that are not specific to a particular project. Such costs are sometimes called overhead costs or general and administrative costs. These are more or less **fixed costs.** One additional project will often make little difference. The standard practice in benefit-cost analysis is to take the **marginal** or incremental approach to counting costs and benefits, but this approach ignores most of the program and overhead costs. The problem with this as a standard practice, therefore, is that it is too generous to the investments and overstates the true returns. In the extreme, overhead costs never get counted anywhere in the organization's decision-making process.

If the organization only occasionally makes major investments, it may be reasonable to ignore program and overhead costs - in essence, letting them be borne by the run-of-the-mill operations of the organization. In this case, it is reasonable to take a marginal-cost approach. In contrast, if the organization makes many investments, it is preferable to include an 'average' allowance for overhead in the costs, although any single investment has little effect on overhead at the margin. If all investment options bear overhead equally, this factor is unlikely to influence the choice among them very much. Even so, it is preferable to have a realistic picture of investment returns, including overhead costs, than to have an unrealistically rosy picture.

4.2.6 Insurance and contingencies

Both insurance and contingencies are efforts to adjust for risk. Do not include the costs for either of these in the table of costs and benefits if intend to do a risk analysis through **simulation** (see Chapter 9). In a simulation, you take risk into account by using maximum-minimum ranges of all variables in the model and by assigning probabilities within these ranges. To include insurance or contingency costs as well would be to double count and to exaggerate the risk.

4.3 Valuing costs and benefits by market prices

In benefit-cost analysis, we normally consider **market prices** as being good measures of the costs and benefits of an investment. (When market prices do not exist in usable form, then the analyst has to construct them). Frequently, however, the market price is only an approximate measure of a cost or benefit. If I buy an apple for $1, for example, the benefit to me of the apple is at least $1 or I would not have purchased it. Clearly, though, the benefit could be higher. The apple might be worth $1.50 to me; that is, I might be willing to pay $1.50 for it if necessary. If I only have to pay $1, then I have a total benefit of $1.50, a cost of $1, and a ‘surplus’ of $0.50. Therefore, when we use market prices as measures of benefits, we are ignoring the **consumer surplus**, which might be important in some cases.

4.4 Consumer surplus and producer surplus as components of value

The concepts of consumer surplus and **producer surplus** are basic to modern benefit-cost analysis. Jules Dupuit, a French engineer, first stated them clearly in 1844. He pointed out that the market price is the minimum social benefit produced by the output of a project. In fact, some consumers would be willing to pay more for the outputs than they actually have to pay.

Consider Figure 4.4.1, which shows two **demand** curves for apples (that is, lines showing the relationship between the price of apples and the quantity of apples demanded at each price). In both cases, the total possible benefit to the community from apple production is given by the area under the demand curve. This area is easy to calculate if the demand curves are approximated by straight lines, as in this example, because the area under the curve is a triangle formed by the two axes and the demand curve.

Figure 4.4.1: When does a standard price approximate the various values people ascribe to an Item?

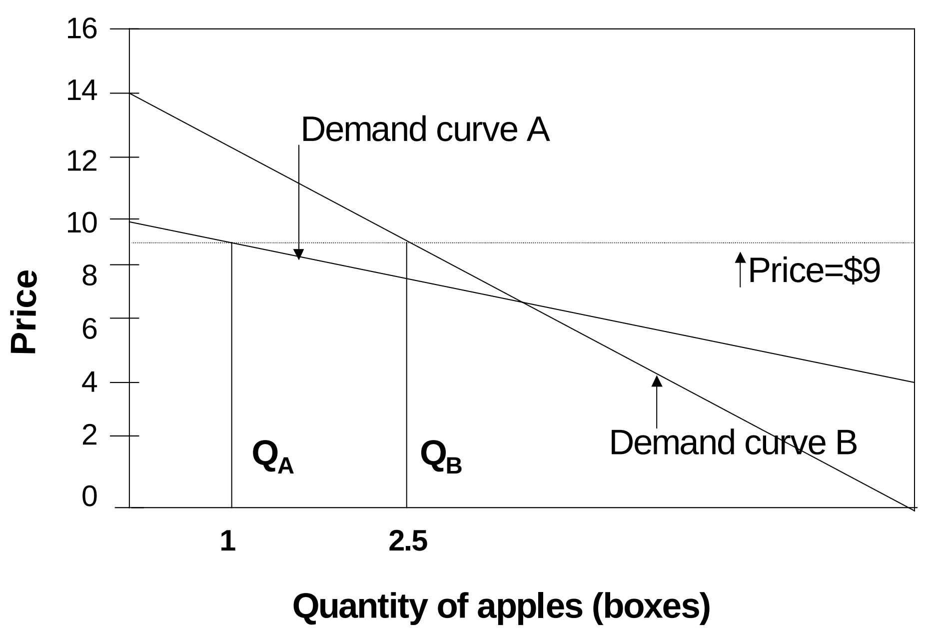

Of course, the community actually receives benefits only from the apples sold and consumed, and this quantity is limited by the price of apples. The area under the curve represents the total community benefit to the left of the point where the price line and the demand curve intersect. To the right of this price point, there is no effective demand, and so no benefits can accrue.

The area under demand curve *A*, to the left of Q_A, is very close to the ‘price × quantity’ rectangle that we normally assume to be the total benefit to the community (in this case, $9 × 1 box of apples). Only a small triangle above the price line is unaccounted for, and this represents the consumer surplus.

Demand curve *B*, on the other hand, has a large consumer-surplus triangle above the price line. In general, the flatter the demand curve (that is, the more that people have a standard value for an apple) the closer the price × quantity rectangle is to capturing the full benefit to the community. Conversely, the steeper the demand curve (that is, the more variety in the value that people give to apples), the less satisfactorily this price × quantity simplification represents the full benefit to the community, because a much larger consumer-surplus is ignored.

An analogous component is the producer surplus. If the producer is prepared to produce a certain quantity of apples at $6 per box, and the market price is $9, then obviously for this quantity of apples, at least, the producer will gain a windfall of $3 per box of apples. The total producer surplus is the area between the price line and the **supply** curve, to the left of their intersection (see Figure 4.4.2).

Figure 4.4.2: Consumer surplus and producer surplus

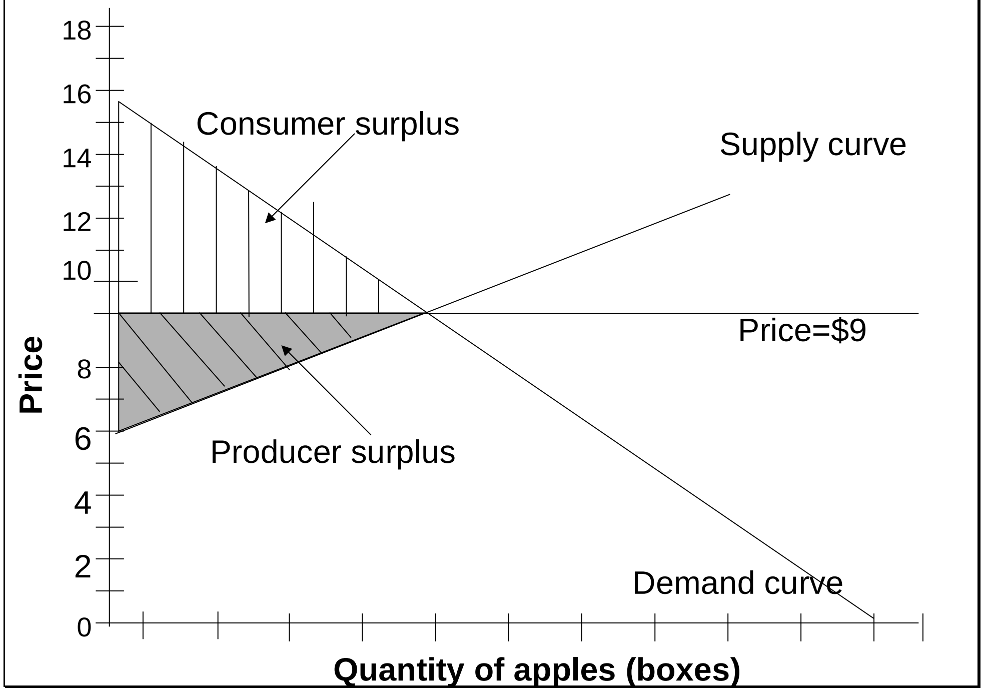

The benefit-cost analyst must decide whether price × quantity is a decent approximation of value; if the simplification is off the mark, more detailed calculations of value are needed. More detailed calculations of consumer and producer surplus are also needed when price × quantity is not an option because prices do not exist or are highly distorted.

4.4.1 Consumer surplus when a public investment changes the price of a good

Public investments in power, water, sanitation and telecommunications projects (and many others) may lower the price of the output. If so, valuing the benefits of the project at the new lower price understates the project's contribution to society's welfare. With a lower price, more consumers have access to the product or service; established consumers pay a lower price and consume more. A special case of this general rule is when supply is rationed at a controlled price below that which consumers would be willing to pay. This situation is rare in Canada (student places in medical schools would be an example), but it is common in some other countries. When it does occur, an increase in supply at the same controlled price involves a gain in consumer surplus over and above what consumers actually pay for the increased quantity of the good or service.

In some cases, part of the increased consumer surplus is offset by a decrease in revenues to the existing producers. For example, if a hydroelectric project reduces the average cost of producing electricity and increases the amount available, the market price of electricity might fall from P_1 to P_2, as shown in Figure 4.4.3. The established consumers save an amount equal to shaded area *A*, but this is offset, from the point of view of the whole economy, by a corresponding loss of revenues to the established producers. The net benefit of the change is therefore only the shaded area *B*. (Figure 4.4.3 is accurate but simplified - consumers' responses are complicated by substitutions among goods when relative prices are affected.)

Figure 4.4.3: Change in consumer surplus

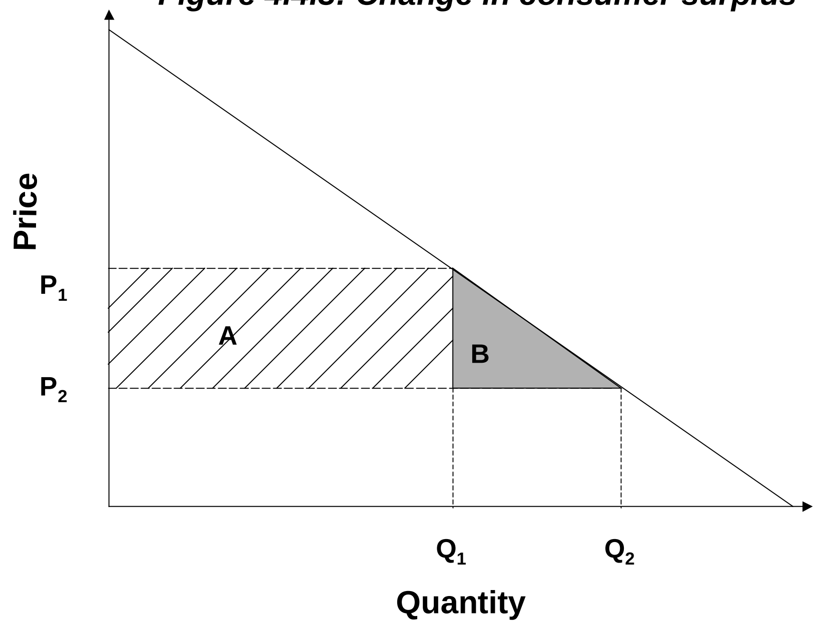

In contrast, if electricity was previously imported and the project is substituting domestic energy for imported energy, then the gain to consumers should theoretically be the whole change in consumer surplus (area *A* + area *B*). In reality, however, the outcome is more often different. The 'import-substitution' enthusiasts of the 1960s and 1970s were no doubt disappointed to see that many public investments produced outputs at prices above the international (border) prices. This, combined with protectionism to exclude the cheaper imports, led to marked reductions in consumer surplus without any offsetting benefits to domestic producers as a group.

4.4.2 Consumer surplus and loss of financial viability

If a public investment depends for its viability on estimates of consumer surplus, and is not viable on a strictly commercial basis, then the analyst must state clearly the amount of the financial shortfall and the source of funds to finance it. In addition, the analysis should explicitly address the benefits that the government subsidising the shortfall will derive from the arrangement, because this can be crucial to the project's sustainability.

If a subsidy is necessary to keep the project operating, then even if the project has a high economic NPV, there might be a significant risk of running out of funds for proper operations. In developing countries, where governments' fiscal positions are usually more precarious than in Canada, this has been a serious problem.

4.5 Valuing costs and benefits without good market prices

When market prices exist but are distorted for some reason, the analyst must estimate what prices would be in the absence of the distortions and then use these adjusted market prices (sometimes called **social prices** or true prices). When there is no market for the good or service in question, there are no market prices - distorted or undistorted. In this case, the analyst has to start from first principles, using the concepts of consumer surplus and producer surplus discussed earlier (see Section 4.4) to estimate the values for costs and benefits.

4.5.1 Estimating value when market prices are distorted

How important distortions in prices are depends on the point of view of the benefit-cost analysis (see Section 2.4). *True value* has meaning only when one knows the point of view. For example, when a private company faces market prices for its costs, those market prices are a good measure of true costs to the company. It does not matter to the company - though it might to an analyst taking the point of view of Canada as a whole - whether the market prices are distorted or not. From one point of view, the prices are good measures of true value; from another point of view, they are not.

In benefit-cost analyses conducted for the Government of Canada, the country as a whole is the most important point of view for the analyst. This requires the analyst to use social prices (sometimes called **shadow prices**) rather than market prices if the market prices are distorted. Such social prices may be substantially different from market prices in some situations, including the following:

- when the currency is misvalued because of foreign-exchange controls;
- when wage rates are kept artificially high by union rules or legislation, despite unemployment;
- when anti-competitive conditions, monopolies or ***monopsonies*** (only one buyer) exist;
- when taxes or tariffs are applied directly to the good or service, as in value-added taxes; and
- when the government regulates or otherwise controls or subsidises prices.

4.5.2 Estimating value when no market prices exist

The true values of resources used or generated by an investment may be difficult to obtain when there are no market prices at all or the market mechanisms are indirect and difficult to observe. The next section explains how values are estimated where no market exists. Examples are given for the following:

- the value of travel-time savings;
- the value of health and safety;
- the value of the environment;
- the value of jobs created;
- the value of foreign exchange;
- the residual value of special-use facilities; and
- heritage values.

4.6 Some examples of difficult-to-estimate values

4.6.1 The value of travel-time savings

Many benefit-cost analysis have been undertaken for transportation projects where the main benefit was timesaving for commuters, both business and leisure travellers.

Travel-time savings for business travellers are generally valued at the traveller's gross wage rate before tax (see Table 4.6.1). In calculating this, researchers have taken into account differences in the wage profiles of typical travellers by different modes of transportation; they have also taken into account time in transit that can be used for work purposes (some part of the 'cruise' time in train and air travel, for example). The definition of gross average wage generally includes an allowance for fringe benefits and overhead costs. Data are gathered through surveys such as the Canadian Travel Survey. Transport Canada undertakes benefit-cost analysis of transportation investments and has developed standards for business travel-time values.

Table 4.6.1: Average value of passenger travel-time savings

Mode	Business travel	Non-business travel
Auto	$27.30	$7.40
Air	$38.30	$7.40
Bus, highway and rail	$27.00	$7.40

Source: Culley and Donkor (1993)
Note: Values are dollars per hour. Base year is 1998 (January).

There is less consensus about how timesaving for leisure travellers should be valued. Because of the greater uncertainty, Transport Canada values travel-time savings for leisure travellers according to some broad principles rather than precise measurements. These principles include setting the value of non-work time savings by adult travellers at 50 per cent of the average national wage rate and making no distinction for differences in demographics (or wages) between modes of transportation.

The value of timesaving in transit of cargo can also be estimated. This is taken to be a function of the value of the cargo and the discount rate. For example, if cargo worth $1 million is in transit for one week less and the cost of capital is 0.002 per cent per week, then the savings would be $20. No methodology exists for distinguishing between business cargo and non-business cargo in this context.

All travel-time savings are counted, even five-minute ones. This is controversial, however. Some analysts believe that consumers do not value small time savings at the same average rate as they value large time savings - that is, a saving of 30 minutes on one trip might be worth more to a traveller than saving one minute on each of 30 trips. Indeed, some argue that a saving of one minute per trip is not worth anything no matter how many people are involved. As well, analysts do not generally take into account distinctions between the values of travel-time savings of different groups of people (employed vs. unemployed; retired vs. working), although it is theoretically possible to do so.

4.6.2 The value of health and safety

For some time, theoreticians thinking about the costs and benefits of health and safety concerned themselves with the question of the ‘value of a life’ in the form of the question: How much is it worth to avoid death? However, it gradually became clear that that was not a sensible formulation of the question. Queen Elizabeth I is reported to have said on her deathbed, ‘All my possessions for an instant of time.’ A better way to think about health and safety is in terms of small increments of risk. The result of government investment in health and safety measures tends to be a small lessening of the risks encountered by broad segments of the population. It is this lessening of risk that can be valued in benefit-cost analysis.

Researchers use three methods to estimate the value of reductions in risk to health and safety. Method 1 is to observe people’s actual behaviour in paying to avoid risks or in accepting compensation to assume additional risk. Method 2 is to ask people to declare how much they value changes in the risks to which they are exposed. Both of these methods are based on the willingness-to-pay principle, and both of them assume that people have the information and skills needed to assess risk and to report their risk-and-reward preferences accurately. Unfortunately, this assumption is almost certainly false. As well, it has not been demonstrated that people have stable risk preferences, even when they do have clear information on the costs and risks and have the skills to assess that information.

Method 3 is to assess statistically the number and type of injuries expected on the basis of historical data. The researcher then counts the treatment costs and wage-loss costs and extrapolates these to the whole affected population. This is a rational approach because it ignores people’s preferences (which are subjective and may or may not be well informed and rational) in favour of a rigorous estimate of the treatment costs and wage-loss costs that would be avoided by the proposed investment. This method avoids the methodological difficulties of methods 1 and 2, but at a cost. It tends to underestimate the true benefit of an investment that reduces risks. When people avoid injury, their benefit is almost certainly greater than avoiding medical treatment costs and wage losses: they also avoid pain and suffering, as well as perhaps the additional costs of becoming disabled. So method 3 gives a minimum estimate of values, but by how much it underestimates the true value we do not know.

In the research literature, the empirical estimates of the value of a life, for example, vary from about $200,000 to about $3 million, with some outliers as high as $50 million. In recent (1994) benefit-cost studies, Transport Canada used the following values: for a fatality, $2,500,000; for a serious injury, $66,000; and for a minor injury, $25,000 (1986 dollars). The values of these parameters are uncertain, however.

4.6.3 The value of the environment

Environmental benefits and costs are often a composite of several factors: health, aesthetics, recreation, and respect for nature. The health and safety aspects are discussed above. Respect for nature (apart from the health and aesthetic components) is close to an absolute value and extremely difficult to quantify, let alone value in dollars. The aesthetic aspects are also difficult to deal with in benefit-cost analysis: first, it is difficult to quantify the aesthetics of a situation at all; and, second, even if quantified, there is no market for aesthetic environmental benefits, or at least no direct market.

Although valuation of environmental goods presents problems, economists have developed some ingenious techniques to estimate the value that people place on such things as water quality and environmental protection. Most of these techniques have specialized applications - they work well in some situations but poorly or not at all in others. A general technique, which relies on the willingness-to-pay principle, does exist **(contingent valuation)**, but its use in the environmental area is controversial because the results it produces may not be as reliable as those produced by other techniques.

In this section we cannot present more than a brief synopsis of some of the major techniques. Reference to more complete works will be necessary [see, for example, Hanley and Splash (1993)].

Sometimes the value of a benefit can be inferred from the cost (for example, money or time) that a consumer is willing to spend to enjoy them. Two techniques that use this approach are the travel-cost (TC) method; and 'hedonic' pricing and land-valuation (LV) method. These methods are clearly limited to 'use values' (such as resource harvesting or direct enjoyment).

As mentioned above, there is a general evaluation technique that can be used for both use and non-use values (the latter include such things as the psychic value of preserving the ecosystem and the value of leaving a viable ecosystem for future generations). This technique is the contingent valuation (CV) method.

It is often appropriate to use a combination of techniques to measure all environmental consequences. This will be required, for example, if one is trying to measure both ecological and commercial losses associated with a resource, such as the loss to the fishery of a toxic spill. In such cases, care must be taken to avoid counting the same loss more than once.

4.6.3.1 The travel-cost (TC) method

The TC method uses the prices for market-traded goods to establish the value of untraded goods. The untraded goods are typically recreational in nature. The costs for traded goods are those that the recreationalist incurs to reach the destination and carry out the activity (the travel cost explains the name of this method). Outfitting costs (at home or on site), equipment rentals, access fees, and so on are usually included as well.

This defines a *minimum* value for the good that is firmly rooted in economic analysis. Whatever the true value of the experience itself, it cannot be less than what the recreationalist actually paid to participate. Application of the TC method is limited to certain use values of the environment, particularly in relation to site-specific activities.

The greatest limitation to the method is probably that it provides no information about non-participants and the value that they may attach to the site. The treatment of travel-time value is also controversial, because vacationers may in fact believe that 'getting there is half the fun'; thus, high (or relatively high) opportunity-cost values for time may be misplaced (see section 4.6.1).

4.6.3.2 Hedonic pricing and the land-valuation (LV) method

Hedonic pricing and the LV method is another technique that links market values to the enjoyment of untraded environmental goods. The objective is to find items that are alike except for one factor (for example, the value of waterfront cottages compared with similar cottages not on waterfront and then compare market values. However, neighbourhoods differ in a myriad of ways: proximity to shopping and cultural facilities, air quality, taxes, etc., not to mention quality of the housing stock (for example, size, layout, and amenities). Therefore there can be methodological problems related to omitted-variable bias or multicollinearity (with two closely related variables, it is impossible to determine which one 'explains' the results). As a result, it is often difficult to derive usable estimates in this fashion.

4.6.3.3 Contingent valuation (CV) method

The CV method essentially asks people what value they give to a resource (use and non-use values). You can ask people what they would be willing to pay to avoid some damaging action or, alternatively, how much compensation they would require to put up with it. The contingency part is a 'what if' question, describing some change in present circumstances, usually some new development. The greatest advantage to this approach is that it can be applied to any valuation problem, including those for which other methods exist.

The CV method is not problem-free, however: numbers put into the questions may bias the answers; respondents may get much of their information on the subject from the interviewer (interview bias); respondents may be influenced by their own perceived self-interest (strategic bias); and, the (usually) hypothetical nature of the exercise can induce laziness on the part of respondents. As a consequence, the results of these surveys are often contested.

Does this mean that the CV method is not worth the effort? One of its earliest and best-known uses was in connection with the proposed construction of a massive, coal-burning electrical generating station upwind from the Grand Canyon. The project went ahead despite strong ‘expressions of value’ on the part of respondents. Today, a clear day at the Grand Canyon is a rarity, and the quality of the experience for millions of visitors is significantly diminished. This example shows that the CV method *can* provide useful information.

4.6.3.4 Other valuation considerations

Sometimes problems arise when analysts estimate indirect market values from measurements of their effects. For example, an aesthetic improvement to the environment might be measurable through the higher property prices that result when the locality becomes more desirable, but the analyst should not double count. Sometimes analysts have listed an environmental benefit in addition to a benefit such as an increase in tourism, when in fact they are one benefit, not two.

4.6.4 The value of jobs created

> Whoever argues impact on jobs rather than impact on the consumer is not an economist but a politician.
>
> - Peter Drucker, *The New Realities*, 1989

The public and the Auditor General of Canada are justifiably suspicious of 'job-creation' claims. Therefore, the analyst should take pains to calculate the figures accurately and to substantiate them well. In particular, any job-creation claims should be on the basis of net effects. If one area of Canada is taxed to make investments in another area, the net effects on jobs nation-wide might be neutral or even negative once the taxation itself is taken into account as a cost. Unfortunately, the benefit of jobs created by expenditures in the public sector has often been cited as a benefit without recognition of the parallel cost (jobs lost) because of the taxation needed to provide the public-investment capital. A local point of view might take into account only the jobs created locally and ignore the jobs lost or displaced elsewhere in Canada, but this is not an appropriate point of view for the Government of Canada.

> This Guide recommends that, as a general practice, the project analyst should adopt the assumption that resources used in the project would otherwise be fully employed. This is especially true of skilled labour which is relatively mobile ... if the project analyst is of the opinion that special circumstances warrant the assignment of shadow prices to the use of otherwise-unemployed resources, the rationale for making such adjustments must be carefully outlined and defended.
>
> - Treasury Board, *Benefit-Cost Analysis Guide*, 1976

In general, job-creation claims are legitimate from the national point of view only when an investment can be made equally efficiently in two regions that have markedly different unemployment rates. In this situation, it may be legitimate to count a lower-than-market cost of labour in the area that has severe unemployment. The idea is that part of the labour force would otherwise be unemployed and therefore it has a low opportunity cost from the point of view of the economy as a whole.

To ascertain how labour should be priced, you need to ask what the workers would be doing in the absence of the project. If they would otherwise be unemployed but are giving up only leisure time to participate in the project, then the shadow wage rate might be low. The previous discussion about travel-time values (see Section 4.6.1) gives a rough order of magnitude of the expected price difference between paid-labour time and leisure time - about 4:1. From a social point of view, then, labourers who would otherwise be unemployed might 'cost' only 20 per cent of their visible wage rate.

Before making such radical adjustments to the shadow wage rate, though, analysts must be confident that the workers would truly have been unemployed. In some cases, analysts have incorrectly taken the overall unemployment rate in a region as the basis on which to calculate the shadow wage rate. This is particularly misleading in industrial regional-development projects, which often require skilled labour that is not available in the region at all, let alone among the region's unemployed. If skilled labour has to be imported for the project, the full wage rate may not be the only cost; there may be other associated costs, such as those for travel and new housing and services.

The labour situations in developing countries can be even more complex than the Canadian situation of dual-labour markets for skilled and unskilled workers. In some countries, disguised unemployment in the form of underemployment is common, especially in rural areas. The withdrawal of surplus workers from a rural area would be of little consequence for agricultural output, even if they had been previously employed. The opportunity cost of unskilled labour, then, might be zero or, at most, little more than subsistence costs, provided that travel, re-housing and similar costs are accounted for elsewhere in the analysis.

Another legitimate use for lower-than-market shadow prices for labour is in an analysis that takes the narrow fiscal point of view of the government. From this point of view, savings from avoiding welfare or unemployment payments are true savings (not just transfers). The savings can be treated as separate items or reflected in lower shadow prices for labour.

4.6.5 The value of foreign exchange

In some cases, a substantial portion of the benefits is generated because the project earns foreign exchange from exports. For example, the National Energy Board regulates the export of natural gas from Canada, and the value of foreign exchange is an important consideration. The key idea is that the Canadian dollar might be undervalued or overvalued in terms of foreign currencies. These distortions can arise if there is not a free and open market for the country's currency or if there are major distortions in the domestic economy because of taxes, subsidies or regulations.

In Canada, the Department of Finance has at times allowed various levels of premium on net export earnings, depending on circumstances, to reflect the true value of foreign exchange to the country. In 1995, Industry Canada calculated the premium to be between 3.5 and 4.5 per cent. In countries with more closed economies, the difference between the market price of foreign exchange and the shadow or true price can be much greater than this. In some cases, the value of the domestic currency is so distorted that prices expressed in it are virtually useless as measures of the true value of resources to the country as a whole. The calculation of the shadow price of foreign exchange is, needless to say, a job for an expert.

4.6.6 The residual value of special-use facilities

Many of the projects of the Government of Canada involve special-use facilities, such as laboratories or training facilities (see Section 4.2.4). At the end of the investment horizon, these facilities have a residual value that may be positive or negative. At a minimum, their value might be land value minus demolition and clean-up costs. At the other extreme, the residual value might be a substantial positive number that reflects a stream of benefits from the ongoing operations of the facility. The Chief Appraiser for Canada has set out procedures for valuing special-use facilities.

4.6.7 Heritage values

Federal Heritage Buildings Policy guides decisions about properties with heritage value. Public Works and Government Services Canada has published a Real Property Best Practice (dated June 1, 1993) that describes applications of this policy.

4.7 Misuse of benefit multipliers

When new resources are generated (or consumed) in a community, the total effect may be larger than the initial transaction would indicate. For example, suppose Josh Brown in the town of Bin inherits $1,000 from a distant aunt. Josh's net income is now $1,000 higher. Josh saves $600 of this inheritance and buys a new suit for $400 from Henry Smith. Henry's net income has increased by the profit on the suit, say $100. Henry saves $70 and spends $30. Obviously, by the time this chain of saving and spending peters out, the total increase in income of the whole community of Bin is larger than the original $1,000 windfall.

The proportion by which the total effect is larger than the initial income gains (or smaller than the initial income loss - the process works symmetrically upwards and downwards) is called the **multiplier** or the ***income multiplier***. The size of the multiplier varies from one community to another. As the example above shows the lower the savings rate and the quicker the rate of new transactions, the higher the multiplier will be. The plausible range of multipliers in the regions of Canada is from about 1.1 to about 2.5.

Unfortunately, multipliers have been misused more often than not. For example, some analysts have applied multipliers to the benefits of a project without any consideration of the equivalent multipliers that should be applied to the (opportunity) costs. This is legitimate only if the analysis is being undertaken from a *local* point of view and some outside agency such as the federal government is paying all the costs. Except in this special case, multipliers must be applied even-handedly to both costs and benefits.

If a project is being analyzed from a social (nation-wide) point of view, the correct application of multipliers is more likely to work against the project's viability than for it. If the investment capital is raised by taxing prosperous areas (where the income multiplier tends to be high) and the investment is made in remote areas (where the multiplier tends to be low because many goods and services are purchased outside the community), then taking the multipliers into account will make the project look worse than it otherwise would.

The use of multipliers in federal government analysis is seldom appropriate. Therefore, multipliers should *not* be included unless there is clear justification for their use.

Best practice - measurement

- The benefit-cost framework can be learned in a short time. In contrast, measurement of costs and benefits is a limitless topic. Other specialists, in addition to the benefit-cost analyst, are generally needed as part of the team.

- When market prices are distorted or do not exist, the main methods for estimating the value of costs and benefits are based on willingness-to-pay.

- Income multipliers should generally be avoided but, when used, must be applied even-handedly to costs as well as benefits.

- The literature can sometimes provide approximate values for such difficult-to-measure items as the value of a clean and natural environment, the value of timesaving for commuters, the value of jobs created, and the value of foreign exchange. Government of Canada standard parameters and benchmarks should be used whenever possible.

5. Time values

5.1 Why time matters

Remember that time is money.

- Benjamin Franklin, “Advice to a Young Tradesman,” 1743

The fact that costs and benefits are spread over time matters for two reasons. First, people prefer to make payments later and receive benefits sooner. Our financial system is built on this basic ***time preference***. There is a loss of earning power if income is postponed until a future date or costs are paid early on. Second, the value of the unit of measurement itself changes over time because of inflation leading to loss of the purchasing power of the currency.

These two factors, inflation and time preference, are independent. Even if there were no change in the purchasing power of a dollar, one would still prefer benefits earlier and costs later. The benefit-cost analyst should, therefore, make two separate adjustments to cash flow figures across time to convert them to standards units of value that can be added or subtracted. The first adjustment is for changes in the purchasing power of the dollar, and the second adjustment involves discounting to reflect time preference.

5.2 Inflation, nominal dollars and constant dollars

The costs and benefits across all periods should be tabulated initially in nominal dollars for three reasons. First, this is the form in which financial data are usually available. Second, adjustments, such as tax adjustments, are accurately and easily made in nominal dollars. Finally, working in nominal dollar enables the analyst to construct a realistic picture over time, taking into account changes in relative prices.

Nominal dollars do not have standard purchasing power. They are sometimes called **budget-year dollars** or current dollars. They are simply the face value of the currency that is paid or received in that period. They cannot be aggregated if they occur at different times because they are not in standard units of purchasing power. Theoretically, nominal dollars can only be added or subtracted if they occur at the same instant. In practice, it is acceptable to add and subtract nominal dollars occurring within the same period as long as the period is short (commonly one year), but it is not acceptable to add and subtract them across periods.

As soon as you are confident that the tables of nominal-dollar costs and benefits are complete and accurate, it is a good idea to convert all figures, or at least the net cash-flow line, to constant dollars before proceeding with calculations (constant dollars have constant purchasing power). To do this, you must select a base point in time at which to express the constant-dollar values. This can be any point in time, but it is often convenient to use t_0, which is the start of the investment period. Selecting the same point in time for constant-dollar conversions, and for present values, is best but not essential. By the way, conversion to constant dollars should not cause you to lose sight of the nominal-dollar tables. They need to be kept visible in the benefit-cost report at all times.

If the benefit-cost analysis is retrospective, the conversion from nominal to constant dollars is simple and accurate because the actual rate of inflation is known. If the benefit-cost analysis is prospective, then you will need projections of inflation. These projections are not easy to come by, and they tend rapidly to become more uncertain the further into the future they project. In this case, as in other cases where data values are uncertain, sensitivity and risk analysis become important tools.

In deciding which index of inflation to use for the conversion to ‘constant-purchasing-power’ dollars ask yourself this: Constant purchasing power for whom? Every index of inflation is based on price changes for a specific basket of goods and services, and the basket is normally defined by the customary purchases of a particular group of people. The index you choose should be as broad as possible. In Canada, the closest thing to a general **price index** covering everything is the Implicit Price Index (IPI). But even the IPI is not the final word on purchasing power because Canadians purchase many goods and services from foreign sources. Benefit-cost analysts often use Statistics Canada’s Consumer Price Index to convert nominal dollars to constant dollars. This is satisfactory if the appropriate reference group is consumers in general. In some cases, it will *not* be. An organization doing a fiscal analysis of a potential investment will need an index of inflation that reflects its own typical purchases. National Defence, for example, constructs its own index for constant-dollar conversions because the unusual mix of its purchases is not reflected in any standard index. Once you have chosen a suitable index of inflation, you are ready to calculate constant dollars. This part is simple.

In prospective analysis, analysts often assume a constant rate of inflation (approximated by using an average rate). This is not a good idea for the early years of the investment, when inflation can be predicted reasonably well. For a longer period, however, inflation projections become largely meaningless (considering the volatility in inflation rates we have witnessed over the past three decades), so the assumption of an average future rate of inflation is the best option. Developing estimates of an average future rate of inflation, however, is a job for a specialist.

The mechanics of adjusting future values to present values, and vice versa, is simple. These values are linked by compound interest. Interest is compounded when the interest earned on an initial principal becomes part of the principal at the beginning of the second compounding period. For example, if Joe Smith invests \$100 at 9 per cent interest compounded annually, then at the end of one year the investment will be worth $\$100 \times (1 + 0.09) = \$109$. Similarly, at the end of two years, t_2, he will have $\$100 \times (1 + 0.09) \times (1 + 0.09)$. Notice that the term in parentheses is repeated once for each period that elapses. In general form, at the end of n years, t_n, the investment will be worth $\$100 \times (1 + 0.09)^n$.

The relationship between constant dollars and nominal dollars is the same. If we start with a constant-dollar amount at t_0 and want to calculate the equivalent nominal-dollar amount at t_n, then we use the formula:

$$N = C(1 + i)^n \qquad [1]$$

where N is the amount in nominal dollars (\$); C is the same amount in constant dollars (\$); i is the annual rate of inflation (%); and n is the number of periods between t_0 and the actual occurrence of the cost or benefit at t_n. In benefit-cost analysis, however, we often find ourselves working in the other direction - that is, we know the nominal-dollar amount for some cost or benefit that will occur at some time in the future, so we need to calculate the equivalent constant-dollar amount for an earlier point in time, such as t_0. In that case, we use this formula:

$$C = N/(1+i)^n \qquad [2]$$

where C is the amount in constant dollars (\$); N is the same amount in nominal dollars (\$); i is the annual rate of inflation (%); and n is the number of periods between t_0 and the actual occurrence of the cost or benefit at t_n. Of course, interest rates are not always quoted conveniently in terms of the period with which you are working. For example, the interest might compound on a daily, monthly or quarterly basis. To find the effective annual interest rate, given a rate quoted for a shorter period, we use the following formula:

$$\text{EAR} = (1 + r/m)^m - 1 \qquad [3]$$

where EAR is the effective annual interest rate (%); r is the quoted interest rate (%); and m is the number of times per annum the interest is compounded.

EAR is also available from standard tables and as a function in many spreadsheet software programs.

5.3 Changes in relative prices

Unfortunately, it has been common in benefit-cost analysis to neglect changes in relative prices over the life of the project. An extreme example is when the analyst constructs a ‘typical year’ of cash flows and simply inflates all input and output values by some standard percentage for each of the following years. This shortcut ignores not only changes in relative prices, but also changes in the composition of inputs and outputs year by year. In general, it is not an acceptable procedure.

A consistent treatment of inflation and relative price changes is as follows:

1. Estimate the future relative price changes for each input and output for each period during the life of the project.

2. Estimate the shadow price of foreign exchange if imports and exports are involved.

3. Obtain estimates of the expected annual changes in the general price level (commonly called inflation).

4. Using these two estimates, calculate the **nominal price** for each input and output for each year of the project.

5. Using the prices estimated above, construct the first complete table of costs and benefits in nominal dollars.

6. Make any adjustments to the cash flows that need to be calculated in nominal dollars (such as adjustments for taxes or loan payments and adjustments in the stock of cash, sometimes called working capital). This gives the pro forma cash-flow table.

7. Deflate all items in the pro forma cash-flow statement for each year by the price index. This gives the constant-dollar table of costs and benefits that is the basis for all further analysis.

5.4 Future and present values

Even when the table of costs and benefits is in constant dollars, the figures are not yet in a standard unit. Constant dollars have standard purchasing power, but it makes a difference whether this is current purchasing power or future purchasing power. To make costs and benefits fully comparable, you must convert their values at various times to values at a single point in time. Present values are dollar values that are not only standardized for constant purchasing power, but are also standardized for the time of occurrence.

To make the conversion to present values, you need a discount rate that reflects the time preference of the reference group. How much is it worth to receive a benefit now rather than at some future time? In federal government benefit-cost analysis, the choice of discount rate has been contentious. Advocates of a project have tended to argue against high discount rates because they make projects look bad (benefits tend to occur later than costs; therefore, high discount rates tend to decrease the benefits more than the costs).

Once the discount rate is selected, calculating present values from future values and vice versa is straightforward. The formula is similar to equation [2] for the adjustment for inflation:

$$\mathrm{PV} = \mathrm{FV}/(1 + k)^n \qquad [4]$$

where PV is the present value at t_0 ($); FV is the future value at t_n ($); k is the discount rate (%); and n is the number of periods between t_o and t_n.

5.5 Discount rates

It is important to understand that the appropriate discount rate depends entirely on the point of view taken in the analysis and that this point of view must be stated explicitly. If, for example, the point of view is that of a particular group of people, then the appropriate discount rate would be one that reflects the time preference of the members of that group. Research shows that if the members of the reference group are poor, the discount rate that reflects their time preference is likely to be high - they will highly value immediate benefits because they have basic needs that are unmet. The cost of borrowing might not approximate their discount rate (unlike the case of a business corporation) if their access to credit is limited or distorted.

Unlike most individuals and organizations, governments frequently take two different points of view in assessing investments - the fiscal point of view (is the project a good one from the government's narrow fiscal perspective?) and the social point of view (is the project a good one for the country?). The discount rates can be quite different from these two perspectives.

5.5.1 The fiscal discount rate

(Narrow fiscal point of view of government)

The fiscal discount rate is the government's cost of borrowing. It is appropriate to use the actual cost of borrowing when the analysis is from the narrow fiscal point of view of the government and the marginal funds for the investment come from borrowing rather than from increased taxes. The fiscal discount rate tends to be low because governments are generally preferred borrowers (taxation is in the background as a guarantee of repayment).

The use of the fiscal point of view and thus of the fiscal discount rate is only appropriate when the proposed investment has few, if any, social implications. Examples are decisions to purchase computers or lease minor accommodation. If the project is large enough to matter to the general economy or if it has aspects that are of interest to the public, then the narrow fiscal point of view is probably inappropriate.

5.5.2 The social discount rate

(Broad social point of view of government)

The social discount rate is roughly equal to the **opportunity cost of capital, weighted** according to the source of investment capital. For the Government of Canada, this is foreign borrowing, foregone investment in the private sector, or foregone consumption. If you know what the government's investment is displacing and what the rates of return would have been for the displaced uses, then you can calculate the opportunity cost. Essentially, the argument is that the government must achieve a return on investment at least equivalent to what the money would earn if left in the private section to justify taxing the private economy to undertake public-sector investments. If the government cannot achieve this it would be better for Canada if the money is left untaxed in the private sector.

Since 1976, Treasury Board has required that benefit-cost analysts use a social discount rate of 10 per cent 'real' per annum - that is, a 10 per cent discount rate applied to **real dollars** (constant, inflation-adjusted dollars). This rate is a stable one because it reflects an opportunity cost in the private sector where the average rate of return to investment (over the whole economy) changes very slowly over the years, if at all. The government's estimate of the social discount rate has been robust, despite some challenges over the years. Social discount rates as low as 7.5 per cent real and as high as 12 per cent real have been proposed and supported by various economists. Estimates by the Department of Finance, however, have consistently supported the 10 per cent real estimate of the social discount rate.

Currently, the only serious challenge to the 10 per cent social discount rate is from those who argue that high discount rates unfairly devalue benefits to future generations, who have as much right to such basics as clean water and clean air as the current generation does. This argument for low discount rates in the public sector is not well based, however. A project with a high rate of return when all its costs and benefits are counted is better for the present generation and, through reinvestment, better for future generations as well. Only when benefits are non-renewable and consumed rather than reinvested is there conflict across generations, with one generation paying and another benefiting. Manipulating the discount rate does not lessen this conflict. It has to be addressed directly by intergenerational consumption analysis.

5.5.3 The rate-of-time preference for consumption

(Point of view of consumers)

Considerable confusion in benefit-cost analysis has been caused by analysts using different numeraires (the units of value). To avoid confusion, one should generally use a 'dollar of investment' as the numeraire and 10 per cent per annum real as the social discount rate. This common approach to investment and rates of return is familiar to economists and non-economists alike.

On the other hand, it is possible (and perhaps theoretically more precise) to use a dollar of consumption as the numeraire. After all, investment is not a final value in the way consumption is. The social rate-of-time preference for consumption is normally taken to be about 4 per cent. This is obviously a much lower discount rate than 10 per cent and on the surface may seem more attractive to those who think that benefits in the distant future (say, general environmental benefits) should not be discounted too heavily.

However if you use a dollar of consumption as the numeraire and a social rate-of-time preference for consumption as the discount rate, then (to make the analysis fair and consistent) you must calculate *shadow prices for the investment funds* in terms of a stream of consumption foregone. Economists who have calculated the shadow price of a dollar of investment funds in Canada and the United States have found that it is about $2.50 in ‘consumption dollars’.

The important point is that the rate used in the first approach (10 per cent discounting of costs and benefits, expressed in an investment-dollar numeraire) and the rate used in the second approach (4 per cent discounting of costs and benefits, expressed in a consumption-dollar numeraire, with a dollar of investment funds shadow priced at $2.50) give the same result when properly applied. Because the outcome of either approach, properly done, is the same, it makes sense to stay with the more easily understood concept of an investment-dollar numeraire and a 10 per cent discount rate (on which everyday thinking about investment and rate of return is based). What is not acceptable is to confuse the two approaches. To use a 4 per cent discount rate without shadow pricing the investment funds is incorrect.

5.6 Strategic effects of high and low discount rates

The choice of a discount rate is extremely important. It has a strong (although hidden) influence on the direction of an organization.

A low discount rate is favourable for the following:

- an active investment program, because capital seems inexpensive;
- outright purchase of assets;
- many and larger projects and programs; and
- projects whose benefits may be long-term.

A high discount rate is favourable for the following:

- a cautious capital investment program, because capital seems expensive;
- leasing and other kinds of deferred-payment options;
- short-term, flexible planning; and
- labour-intensive rather than capital-intensive solutions.

5.7 The discount rate as a risk variable

The 1976 Treasury Board *Benefit-Cost Analysis Guide* recommended a social discount rate of 10 per cent real, and 5 to 15 per cent real per annum in sensitivity analysis. Experience has shown, however, that this range was too broad. Most projects look good at a 5 per cent discount rate and poor at a 15 per cent discount rate. A credible and more useful range for the social discount rate is normally about 8-12 per cent real per annum (for risk analysis), with a most likely value of 10 per cent real per annum.

Because there is some uncertainty about the correct value of the discount rate, you should include it as a **risk variable** in the parameter table of a benefit-cost and risk analysis using simulation. This makes it less important to fix on a precise value of the discount rate and places more emphasis on identifying the likely range of values of the discount rate and on interpreting the results of the financial simulation (see Chapter 9).

Best practice - inflation adjustments and discounting

- To ensure that changes in relative prices are properly recognized, tables of costs and benefits should be first constructed in nominal dollars, and cash flows should be set out for each period to the investment horizon. Conversions to constant dollars or to present value dollars should wait until all costs and benefits over time are worked out in nominal dollars.
- Adjusting for inflation is not the same thing as discounting to present values, so each should be done independently.
- The appropriate discount rate depends on the point of view of the analysis and also on the choice of numeraire.
- The Government of Canada uses a fiscal discount rate (based on a narrow 'internal' point of view that is appropriate mostly for small projects) and a social discount rate (based on a nation-wide point of view). With the normal dollar of investment as the numeraire, the appropriate social discount rate (as measured by the Department of Finance and Treasury Board of Canada Secretariat) is about 10 per cent real per annum. The plausible range for risk analysis is 8-12 per cent.

6. Decision rules

There is nothing so practical as a good theory.

- Kurt Lewin, *Field Theory in Social Science: Selected Theoretical Papers*, 1951

A decision rule tells us whether an investment is worthwhile and whether one investment is better than another is. In this chapter we consider how decision rules are used with deterministic data - that is, we ignore uncertainty in the data for the moment (Section 9.7 shows how the same decision rules can be adapted for uncertain data).

6.1 Net present value

NPV is the present value of all benefits, discounted at the appropriate discount rate, minus the present value of all costs discounted at the same rate. An NPV is always specific to a particular point in time, generally t_a, the time of the analysis, or t_0 the start of the project.

The formula for the calculation of net present value is as follows:

NPV = initial investment costs + the sum of the present values of costs and benefits for each period within the investment horizon. [5]

The NPV can be calculated in several different ways. Obviously, you could calculate the NPV of benefits and the NPV of costs separately and then subtract them. More often, the analyst subtracts costs from benefits in each period, giving a single line of net cash flow, and then discounts the net cash flow to give the NPV. The arithmetic of this latter procedure is a little simpler, but, more important, the net cash flow is itself useful information for managers. Many projects and enterprises with a positive NPV have failed because of cash-flow problems.

For example, if the initial investment were $100 and there were $70 in benefits and $25 in costs for each of 3 years, and the discount rate were 10 per cent per annum, then the NPV would be:

$$
\begin{aligned}
\text{NPV} &= -\$100 + (\$70 - \$25)/(1 + 0.1)^1 + (\$70 - \$25)/(1 + 0.1)^2 + (\$70 - \$25)/(1 + 0.1)^3 \\
&= -\$100 + \$40.91 + \$37.19 + \$33.81 \\
&= \$11.91
\end{aligned}
$$

This formula follows the accounting convention discussed in Chapter 2; that is, all costs and benefits are assumed to occur at the end of their period, except for large initial expenditures, which occur at t_0 and are not discounted.

6.1.1 Net present value and break even

An NPV of zero does not mean ‘break even’ in the normal sense of costs equalling benefits. NPV is more like excess profit than it is like profit. If a project has an NPV of zero, the project earns the normal rate of return (which is, of course, equal to the discount rate). For example, if a project earns 10 per cent per annum and its cash flows are discounted by 10 per cent per annum, the result will be an NPV of zero.

We value NPV not because it tells us whether the project breaks even, but because it tells us whether it is worth doing the project instead of leaving the money in the normal alternative investment (which earns 10% per annum).

6.2 Two essential decision rules

Many projects have complex patterns of costs and benefits over time, and you cannot use the ‘eyeball’ method to determine which project is preferable. We need decision rules to guide us. Many decision rules have been proposed.

Some work well only in particular situations; others are prone to error. Only two rules are consistently accurate and reliable. These are given below.

Case 1: Single project, unconstrained budget, 'go' or 'no go' decision

Decision rule 1: Do not undertake projects whose NPV is less than zero, unless you are willing to 'lose money' to achieve a non-economic objective.

Example 6.2.1

	NPV	Decision
Project A	+$3	Accept
Project B	+$0	Indifferent
Project C	- $1	Reject

Case 2: Alternative projects, constrained budget, a 'best set' decision

Decision rule 2: Given a choice among alternative projects, maximize the total NPV.

6.2.1 Problem of independence from the scale of investment

People are generally comfortable with the idea that a project with an NPV of -$27 is unacceptable, but they are less comfortable with the idea that project B, whose NPV is +$3, is always preferable to project A, whose NPV is +$2, no matter how much investment went into each. The reason harks back to our previous discussion about NPV being like excess profit rather than profit.

A simple example for one period should make clear why you should always prefer the larger NPV. The key is to realize what happens when we standardize the level of investment - that is, when we take into consideration what happens to the capital we have left over if we choose the smaller investment.

Example 6.2.2

Suppose project A requires an investment of $100 and project B an investment of $150. If you invest in project A instead of project B, then you have an unused residual of $50, which earns the normal rate of return invested somewhere else. This residual, however, has an NPV of zero (it earns the same rate of return, say 10 per cent, as the rate used to discount it to a present value). Therefore, if you choose project A, you have a total of $100 earning 10 per cent plus an NPV of $2, and you also have $50 earning the normal rate of return. If you choose project B, similarly, you have $150 earning the normal rate of return plus an NPV of $3.

Project A earns (10% of $100 + 10% of $50) + $2
Project B earns (10% of $150) + $3

The figure in parentheses will be the same for project A and project B whatever their scale of investment, so it is only in the NPV that differences will show. The amount of investment involved in two alternatives is irrelevant to our decision once we know the NPVs. Simply choosing the better NPV will always be correct. Another way of looking at the same calculation follows.

Example 6.2.3

	Basic return	**NPV**	**Total return**
Project A			
Investment of $100	$10	$2	$12
Residual of $50	$5	$0	$5
			$17
Project B			
Investment of $150	$15	$3	$18

It is important that the decision-maker grasp this concept if NPVs are to be a useful guide. An NPV of zero does not mean ‘break even.’ It means the project earns the normal rate of return, say 10 per cent. A negative NPV of, for example, -$300 does not necessarily mean the project makes a loss in the colloquial sense. It means that it makes the normal 10 per cent return, less $300.

You can rank projects by their NPVs without worrying about the scale of the project. In contrast, you cannot rank projects by their internal rates of return unless you consider their scale as well (see Section 6.3.1). This is counter-intuitive for many managers.

6.3 Unreliable decision rules

6.3.1 The internal rate of return

The **Internal Rate of Return (IRR)** is the discount rate that makes the NPV of the project zero. An IRR higher than the standard discount rate indicates that you should go ahead with the project, and when you are choosing among alternative projects, a higher IRR is preferred. If project A earns an IRR of 15 per cent, for example, whereas the ordinary project earns 10 per cent, then project A is an attractive investment.

The IRR has three important limitations (see boxes below) that make it a poor substitute for NPV as a decision rule. Nevertheless, many managers find the IRR intuitively appealing in a way that the NPV is not. They tend to think that the meaning of an IRR is transparent, but it is not. When you calculate the IRR, you need to interpret it with care.

The underlying formula for the IRR is the same as for the NPV. If you know the discount rate, you can calculate the NPV and vice versa. The mathematics of the IRR calculation, however, is not based on a proof and a formula. In practice, the analyst uses a computer to calculate the IRR by trial and error iterations. Given a guess at the likely IRR, the computer enters higher and lower values for i in the formula until it results in an NPV of zero.

Most spreadsheets in common use have a limit on the number of iterations the computer will try. If the computer does not find a discount rate that gives an NPV of zero within this limited number of iterations, it gives an error message. The analyst then has to start the process again with a different guess at the value of the IRR. In addition to this procedural awkwardness, the IRR has two other limitations that make its use doubtful. These are given below.

Limitation 1: Simple comparisons between IRRs may be misleading if the projects are not the same size. A project with an IRR of 7 per cent is not necessarily a better choice than one with an IRR of 6 per cent. The size of each project and the discount rate can influence which project is best.

Example 6.3.1

	Project A	**Project B**
Total cost	$100	$10,000
IRR	7%	6%
Discount rate	5%	5%

If you choose project A, you will have $100 earning 7 per cent plus the residual $9,900 earning 5 per cent (total return = $7 + $495 = $502). If you choose project B, you will have the whole $10,000 earning 6 per cent ($600). Project B is better, even though it has a lower IRR than project A.

Limitation 2: In many cases, more than one value of the IRR will solve the equation, and it may not be apparent to the analyst that other equally good values exists because the computer typically stops when it finds any acceptable value of the IRR.

Multiple values of the IRR (some negative, some positive) are especially likely if the annual net cash flow of the project alternates between positive and negative figures, a common event because of the cyclical re-capitalisation requirements of projects and/or fluctuations in the prices of inputs and outputs. In some cases, analysts 'bend' the accounting rules to obtain a cash-flow pattern that gives a single value for the IRR, but this is not a satisfactory solution. At best, the possible existence of multiple values of the IRR throws a shadow over its use; at worst, it may lead to incorrect choices among projects.

6.3.2 The benefit-cost ratio, payback period, and present value of costs

Decision rules other than the NPV ones given in Section 6.2 are sometimes used correctly, but none of them are satisfactory as a general rule. The three most common involve **benefit-cost ratios, payback period** and the present value of costs.

Benefit-cost ratios

A benefit-cost ratio is the ratio of the present value of benefits to the present value of costs. The decision rule here is that you should reject any project with a benefit-cost ratio of less than 1, and you should rank projects in order of their benefit-cost ratios. The first part of this rule works. The second part, however, may not. This is because it is possible to change the benefit-cost ratio substantially by artificial changes in the accounting for benefits and costs (although it is not possible to change a ratio of less than one to a ratio of greater than one or vice versa — try it!).

Remember that a positive benefit is equivalent to a negative cost. Almost any cost or benefit could serve as an example. Consider expenditures on an access road to a new park. These could be added to the costs of the park or subtracted from the benefits. Either choice is correct. However, the benefit-cost ratio would be increased or decreased artificially, depending on this arbitrary accounting decision.

Example 6.3.2

	Benefits	**Costs**	**Benefit-cost ratio**
Project A	$100	$60	100/60 = 1.66
Project A (same project, but netting $30 out of the benefits rather then listing it as a cost)	$70	$30	70/30 = 2.33

Payback period

The payback period is the time it takes for the cumulative present value of benefits to become equal to the cumulative present value of costs. In general, shorter payback periods are better. However, this can be a misleading decision rule because it ignores everything that happens after the payback point. It is quite possible for a project to have a higher NPV and a longer payback period (see Figure 6.4.1). ‘A’ has a quicker payback, but ‘B’ reaches a higher NPV.

Figure 6.4.1: Two ‘paybacks’ compared

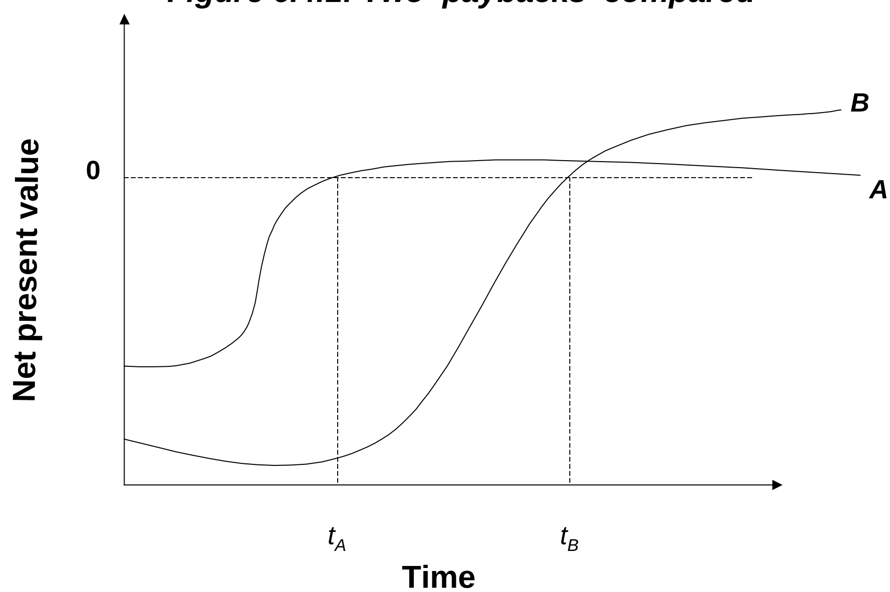

Present value of costs

When the benefits of two alternatives are exactly the same, you may choose between them on the basis of the lowest present value of costs. This is not a reliable decision rule, however, because you cannot tell from the present value of costs whether the project should be done at all. As well, the premise that benefits are constant is generally a simplification and often may not be valid. For example, a constructed facility will seldom be exactly the same, qualitatively and quantitatively, as leased accommodation. Some of the differences might be quite important. Furthermore, the present value of costs is subject to manipulation of the type described above in the discussion of benefit-cost ratios. That is, the accounting for costs and benefits (benefits can be counted as negative costs) can artificially change the apparent present value of costs.

Best practice - decision rules

NPV decision rules are best. Other decision rules should be used with extreme care.

The two basic decision rules are the following:

1. Do not undertake projects whose NPV is less than zero, unless you are willing to 'lose money' to achieve a non-financial objective.

2. Given a choice among alternative projects, maximize the total NPV.

7. Sensitivity analysis

The whale that wanders round the Pole
Is not a table fish.
You cannot bake or broil him whole.
Nor serve him in a dish.

- Hilaire Belloc, "The Whale," *The Bad Child's Book of Beasts*, 1896

7.1 What is sensitivity?

In benefit-cost analysis, the outcome is typically influenced by several uncertain factors. This is true in fields as diverse as health, education, employment, and economic development. It is important to know how 'sensitive' the outcome is to changes in those uncertain factors. It helps you to determine whether it is worthwhile spending money to obtain more precise data and whether you can act to limit uncertainty (for example, you could redesign the project components or simply keep a watchful eye when managing the project). As well, sensitivity analysis helps you to communicate to decision makers the extent of the uncertainty and risk in the program.

Nevertheless, sensitivity analysis is a limited tool. It treats variables one at a time, holding all else constant. Simultaneous actions and interactions among variables in the real world are ignored. It can be a mistake to take the results too seriously because a variable that appears to be key when considered in isolation might or might not be key when considered along with other variables that strengthen or weaken its effect on the outcome of the project. Only a risk analysis (Hertz and Thomas 1983, 1984) can accurately identify the influence of each variable.

Nevertheless, sensitivity analysis is a helpful (although limited) step in exploring the deterministic model. It is the second of three phases in the general analysis:

1. Build a deterministic model using single 'best' values (base values) for the input variables.

2. Explore the outcome's sensitivity to each input variable and then take action to reduce the risk of uncertainty where possible.

3. Make a full risk analysis using probabilities for many variables simultaneously.

Sensitivity analysis gives you a better understanding of the model. As this understanding develops, you can take action when appropriate. In some cases, the only action you can take is to obtain better data. For example, if you are deciding whether to purchase a heart-lung machine, the outcome is sensitive to 'probability of an influenza epidemic,' a variable that you cannot control. In other cases, you might be able to fix or constrain the value of the variable. For example, if the outcome is particularly sensitive to an operator's wage rate, then you could negotiate this rate beforehand. Fixing the wage rate would dramatically lower the sensitivity of the outcome to this variable. The more you can minimize the sensitivities, the more precise the estimate of the outcome will be.

7.2 Gross sensitivity

In its simplest form, which we might call gross sensitivity, sensitivity analysis involves calculating (one variable at a time) how much the NPV changes if the influencing variable changes by a standard percentage, say 10 per cent. Suppose you are deciding whether to purchase a heart-lung machine whose NPV is affected by four variables: insurance costs, operating costs, the price of the machine, and the usage rate. A quick glance at Table 7.2.1 indicates that the decision is quite sensitive to three of the four variables.

Table 7.2.1: An example of gross sensitivity of the NPV to input variables

Variable	Change in NPV in response to a 10% change in the variable

Insurance costs	15%
Operating costs	21%
Machine price	7%
Usage rate	19%

7.3 What determines sensitivity?

A superficial interpretation of Table 7.2.1 could be misleading. The 'effective sensitivity' of the outcome to a particular variable is determined by four factors:

- the *responsiveness* of the NPV to changes in the variable;
- the magnitude of the variable's *range of plausible values*;
- the *volatility* of the value of the variable (that is, the probability that the value of the variable will move within that range of plausible values); and
- the degree to which the range or volatility of the values of the variable can be *controlled*.

The first of these factors, the responsiveness of the NPV to changes in the variable, has two components. The first component is the direct influence of the variable on the NPV. The second component is the indirect influence of the variable, through its relationships with other variables that themselves are related to the NPV. Positive correlations with other influential variables will magnify the ultimate influence of both, and negative correlations will dampen their influence. These influences cannot be fully identified until you have set up a simulation model that is capable of dealing with the simultaneous interactions of many variables.

7.4 Sensitivity and decision making

We are most interested in the sensitivities that might change a positive decision on the project to a negative decision and vice versa. Four calculations help us estimate the likelihood of such a switch:

1. What is the *range of influence*? That is, how much does the NPV change when the variable changes from its lowest plausible value to its highest plausible value?
2. Does this range of influence contain an NPV of zero? If it does, then the variable has a **switching value** - that is, a value at which our **appraisal** of the project switches from positive to negative.
3. What is the *switching ratio* for the variable? That is, by what percentage does the variable have to change to hit a switching value?
4. What is the *switching probability*? That is, how likely is the variable to reach the switching value?

If we add this information to the gross-sensitivity calculation shown in Table 7.2.1, we begin to have a reasonably complete picture of the likely sensitivity to a variable, although, of course, within the limits imposed by considering one variable at a time (see Table 7.4.1).

Table 7.4.1: An example of several indicators of the sensitivity of NPV to input variables

Input variables

Indicators of sensitivity	Insurance costs	Operating costs	Machine price	Usage rate
Gross sensitivity	15%	21%	7%	19%
Range of influence	10%	17%	5%	35%
Switching value	No	Yes	No	Yes
Switching ratio	-	9%	-	63%
Switching probability	-	40%	-	42%

By scanning all the data in Table 7.4.1, you can easily see that the NPV is sensitive to neither insurance costs nor machine price. In particular, neither variable can move the NPV enough to hit a switching value. In contrast, the NPV is sensitive to both operating costs and usage rate. Of these two, usage rate obviously has a more volatile value - although it has to change by a much larger percentage to hit a switching value than the value of operating costs does, it is about equally likely to do so (42 per cent vs. 40 per cent). Although usage rate has to change more than operating costs to cause a crucial change in the NPV, its volatility makes it equally likely to do so. On this evidence, we would tentatively conclude that the two key variables are equally influential.

7.5 Two-variable sensitivity analysis

So far, we have analyzed sensitivity one variable at a time. Extending this to two variables is the next step toward true risk analysis. Scenarios defined by two interacting variables, although still not complete and realistic indicators of sensitivity, are at least an improvement on single-variable analysis. The joint influences of two input variables on NPV can be shown in a matrix-like Table 7.5.1.

Table 7.5.1: Influence of combinations of two input variables on net present value

Discount			Usage rate		
Rate	500	475	450	425	400
0.10	$16,814	$12,987	$9,161	$5,335	$1,509
0.11	$14,923	$11,200	$7,476	$3,753	$29
0.12	$13,082	$9,459	$5,835	$2,212	***($1,411)***
0.13	$11,288	$7,763	$4,238	$713	***($2,812)***
0.14	$9,541	$6,112	$2,683	***($746)***	***($4,175)***
0.15	$7,840	$4,505	$1,170	***($2,165)***	***($5,500)***
0.16	$6,185	$2,941	***($302)***	***($3,545)***	***($6,788)***
0.17	$4,573	$1,420	***($1,734)***	***($4,887)***	***($8,040)***
0.18	$3,004	***($61)***	***($3,127)***	***($6,192)***	***($9,257)***
0.19	$1,478	***($1,501)***	***($4,481)***	***($7,460)***	***($10,440)***

0.20	***($6)***	***($2,902)***	***($5,798)***	***($8,693)***	***($11,589)***

For each discount rate, you can see what usage rate has to be attained if the machine is to be economical. You could draw a diagonal line (from the lower left toward the upper right) that represents all the combinations of discount rate and usage rate that result in a NPV of zero. This diagonal divides the table into two ‘strategy regions’ (defined by combinations of discount rates and usage rates). In one of these strategy regions, all values of the NPV are negative (the bold numbers); in the other region, all are positive.

If most of the sensitivity in the model results from only two key variables, then this sort of analysis is very instructive. Even if there are more than two key variables, two-variable analysis takes us at least one step closer to understanding the workings of the model in a realistic setting.

7.6 Graphic analysis of sensitivity

Although we have been interpreting tables to ascertain sensitivity, in practice, you are more likely to use graphs because the variable x outcome interaction is visible over a reasonable range of values. Sensitivity analysis is exploratory, not definitive, so making the patterns in the data visible is the first priority.

7.6.1 Sensitivity curves

A graph that shows changes in NPVs against changes in the risk variable is simple and useful. You can easily read the switching values to see how sensitive the outcome is to changes in the variable. If the changes in the variable are presented on the graph in percentages (and thereby standardized), it becomes possible to put the curves for two or more variables (calculated one at a time, of course) on the same graph. This is useful because the slopes of the curves indicate the relative sensitivity of each variable. The more the NPV changes for a given change in a variable, the more sensitive it is to that variable, volatility being equal.

If the percentage change in NPV is on the *x*-axis and the percentage change in the risk variable is on the *y*-axis, then any flat curves indicate a strong sensitivity. As you can see from Figure 7.6.1 the NPV of the example project is more sensitive to variable ‘B’ than it is to variable ‘A’. A 10 per cent change in risk variable ‘B’ results in a much larger change in NPV than the same change in variable ‘A’.

Figure 7.6.1: Sensitivity of net present value to a benefit variable

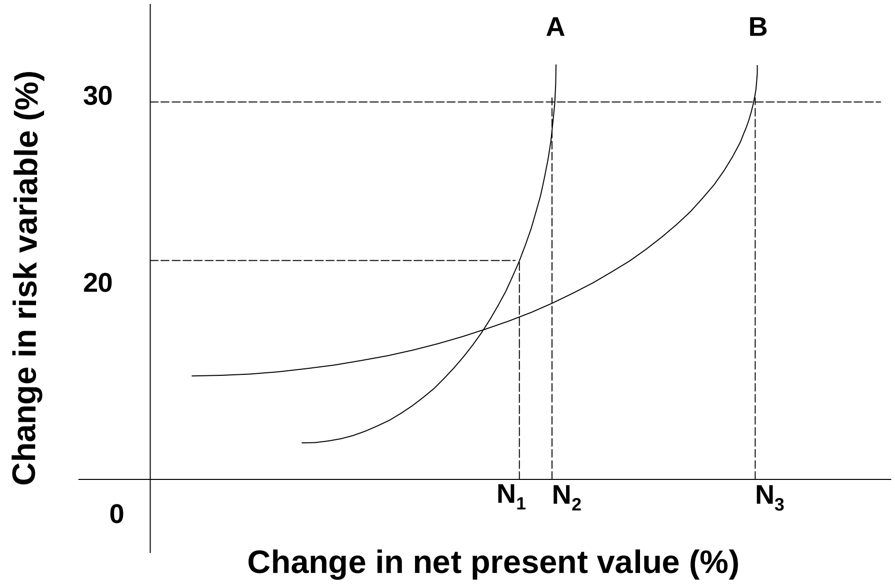

7.6.2 Spider plots

Sensitivity plots can be consolidated to show many input variables on one chart. This type of consolidated chart is called a *spider plot* (see Figure 7.6.2 for an example). The centre of the spider plot is the NPV when all the variables are at their baseline values. The curves on the spider chart show how the NPV changes as the values of each variable change, all others being held equal.

The lengths of the spider lines vary because each variable has its own plausible range within which it can change. The values of one variable might vary by only 10 per cent up or down from its baseline value; another variable might be highly uncertain, varying between +170 per cent and -60 per cent.

Figure 7.6.2: A spider plot

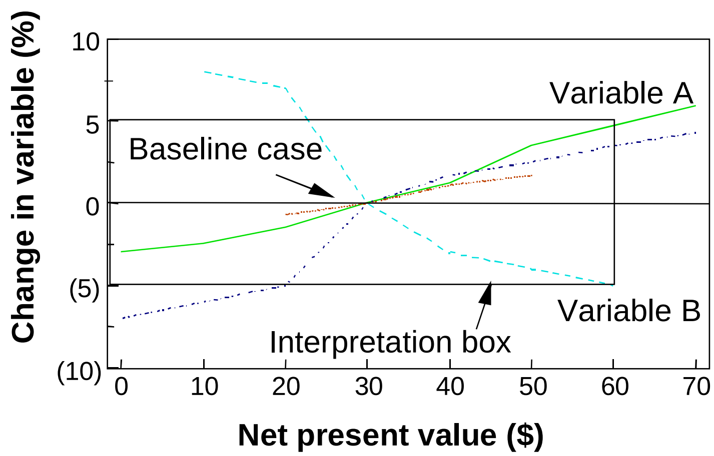

The interpretation box superimposed on the spider plot is an aid to interpretation. The top and bottom of the box indicate +5 and -5 per cent change, respectively, in the variables. Thus, wherever a 'tentacle' of the spider plot crosses the top or the bottom of the box, it shows what the NPV would be if there is a 5 per cent change in the particular input variable.

The NPV at the left side of the interpretation box is zero, and therefore the tentacles cross this side at the switching value of the variables. The right side of the box is set at an NPV of 60 for symmetry (the NPV of the baseline case is 30). These dimensions of the box are reasonable for this particular analysis, but they will vary in other analyses (the box is an aid to reading the spider plot, nothing more). The spider plot shows how much each variable would have to change, other things held constant, to produce an NPV of zero (the switching ratio).

7.6.3 Tornado charts

Tornado charts give us another quick, although partial, picture of relative sensitivity. Each bar in the tornado chart shows the range of the NPV when each variable is allowed to change (one at a time) from its highest to its lowest value. You can see from a quick glance at the shape of Figure 7.6.3, (which has variables arranged in descending order of influence from top to bottom), why it is called a tornado chart. Of course, the magnitude of the variable's range of plausible values is not the only factor that determines sensitivity. The volatility - the probability that the value of the variable will move within that range - is also important to sensitivity, but you cannot see volatility on a tornado chart.

Figure 7.6.3: A tornado chart showing in the change in NPV if an input variable is changed from its lowest to highest value

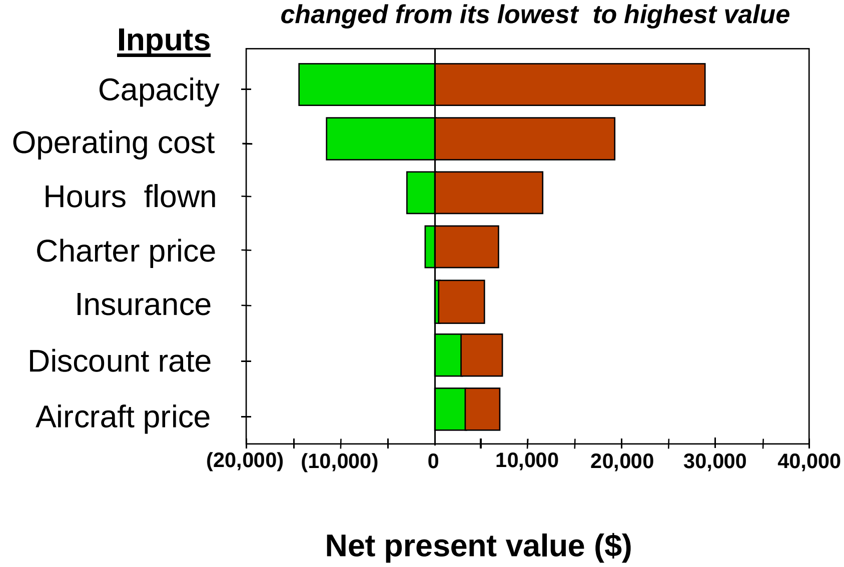

The length of each bar is a measure how much the variable can influence the NPV. The shading within the bar changes at the NPV, which corresponds, to the deterministic value of the variable.

7.7 Action on sensitivities

Once you have identified the key sensitivities among the risk variables, one by one (everything else held constant), you can start to think about managing risk.

- Are there input variables in the model that are correlated and therefore dampen or enhance the influence each might have in isolation?
- Can diversification help? Are there other investments that could be made at the same time where the same variable works in the opposite direction?
- Could you identify the value of the key variable with more certainty by gathering more information, and if so, is the information worth the cost to gather?

Once you have answered these questions, you can formulate an action plan to minimize uncertainty and thereby limit risk.

Best practice - sensitivity analysis

- Sensitivity analysis is a useful technique for finding out how important each variable in the benefit-cost model is.

- Sensitivity analysis cannot deal with more than two variables at a time, so it does not tell us much about the project's level of risk. Until all variables are allowed to vary simultaneously, we do not know whether their individual effects on risks are magnified or cancelled out by each other.

- Four factors contribute to sensitivity: the responsiveness of the NPV to changes in the variable; the magnitude of the variable's range of plausible values; the volatility of the value of the variable (that is, the probability that the value of the variable will move within that range of plausible values); and degree to which the range or volatility of the value of the variable can be controlled.

- Graphic analysis, including the use of sensitivity curves, spider plots, and tornado diagrams, is often useful.

- The switching value of a risk variable can be an important consideration in an investment decision. It can help the decision maker weigh the risk.

8. General approaches to uncertainty and risk

Making risk characterisations more complete, subtler and more data rich should help decision-makers make more balanced, subtler and better-informed decisions.

- National Academy of Public Administration, *Setting Priorities, Getting Results*, 1995

Dealing with uncertain data is a very large part of practical benefit-cost analysis. Chapters 3 and 4 discussed uncertainty in the incremental effects of a program or project. In this chapter, we look at ways to cope with uncertainty in the overall financial and economic analysis. Many of the tools are similar.

8.1 Approaches to quantifying uncertainty-related risk

There are three approaches to dealing with financial and economic risk in benefit-cost analysis:

- expected values (certainty equivalents) of scenarios;
- risk-adjusted discount rates; and
- risk analysis through simulation.

Given the present state of the art, the first two approaches have limited applicability. Only the third method, simulation, offers a practical technology for analyzing the overall risk of a project.

8.2 Expected values of scenarios

If an investment has two possible outcomes - $10 and $100 - and their probabilities are 30 and 70 per cent, respectively, then the **expected value** or **certainty equivalent** of the investment is (0.3 x $10) + (0.7 x $100) = $3 + $70 = $73. If you have a completely rational attitude to risk, then it shouldn't matter to you whether you make the investment or accept the $73 instead.

Few benefit-cost analysts take this scenario approach because in most cases there are so many possible outcomes that it is too difficult to think clearly about the probability of each separately. On occasion, however, scenarios can provide useful information about risk. For example, an oil company is trying to decide whether to build a new pipeline across an iceberg-infested strait. The pipeline costs $100 million (present value at t_0). The chief executive officer (CEO) foresees three possible scenarios; and each of these has a predictable outcome for the company.

Scenario 1: No iceberg hits the pipe.

Outcome: Company revenue from the pipeline: $135 million, t_0.

Scenario 2: An iceberg hits the pipe, but the pipe can be repaired.

Outcome: Company revenue from the pipeline: $93 million, t_0.

Scenario 3: An iceberg hits the pipe, and the pipe cannot be repaired.

Outcome: Company revenue from the pipeline: $9 million, t_0.

Now comes the difficult part. Suppose the CEO commissions a study by iceberg-risk consultants and is told that there is a 60 per cent chance that no iceberg will hit the pipe, a 30 per cent chance that an iceberg will hit the pipe but the pipe will be repairable, and a 10 per cent chance that an iceberg will hit the pipe and the pipe will not be repairable. Therefore, the expected value of the $100 million investment is (0.6 x $135 million) + (0.3 x $93 million) + (0.1 x $9 million) = ($81 + $27.9 + $0.9 million) = $109.8 million. The CEO decides to go ahead with the pipeline since the expected benefit ($109.8 million) is greater than the cost ($100 million). An iceberg hits the pipe, but the pipe is reparable. The company loses $7 million. The CEO, however, had made the right decision, given the information he or she had to work with.

The main difference between this procedure and the recommended simulation procedure is in the reliability of the estimates of probability they make. Did the iceberg-risk consultants really have the expertise to assign probabilities to the likelihood of an iceberg collision? There were no data. Assigning subjective probabilities to ‘big-picture’ scenarios is essentially a guessing game. In contrast, it is plausible that an apple-pricing expert can forecast apple prices within a reasonable range a year ahead, based on historical price data, demand trends, and consideration of factors that might intrude. There is a subjective or judgmental element in forecasting apple prices, too, but there are data available. Legitimate experts have developed good judgement in the matter, so they are able to express expected apple prices as a range and specify a **probability distribution** with reasonable confidence. Risk analysis is part science and part art; and part of the art is knowing when and where in the benefit-cost model to use probabilistic data.

8.3 Risk-adjusted discount rates

> The practice of changing the discount rate to allow for uncertainty in project evaluation has limited validity. In effect, it implies that uncertainty compounds itself at a fixed rate over time. This is unlikely to be the case. Where different degrees of uncertainty can be ascribed to future values of variables, it is preferable to let estimates of future annual benefits and costs reflect these different degrees of uncertainty, and to aggregate present values using an interest rate that has not been adjusted for risk.
>
> - Treasury Board, *Benefit-Cost Analysis Guide*, 1976

Another approach that purports to deal analytically with risk is risk-adjusted discount rates. The basic idea is that, in a perfect market, all investments earn the same rate of return. Otherwise, capital would flow to the high-return areas pushing average returns down until the rates equalised. Therefore, visibly different rates of return must incorporate the same basic rate plus a premium for risk so that, in the long run, only the basic return is gathered by the investor. If this is so, then the appropriate discount rate (cost of capital) is the basic rate plus a premium for risk. This combination is the risk-adjusted discount rate.

This approach has several flaws. One is that we do not operate in perfect markets, so the differentials observed in rates of return might be due to other systematic or random factors, rather than to project risk. Another objection is that the argument confuses the lender’s risk in lending capital to the investor with the risk inherent in the proposed project. They are not the same things.

Furthermore, the use of a risk premium on the discount rate (which is a compound-interest rate) can lead to odd results. A risk premium on the discount rate means that the absolute value of the risk premium in dollars increases as time goes on, and this does not make sense for many investments that are uncertain for an initial period but tend to settle down to a known pattern of costs and revenues and therefore low risk (real estate development, for example).

These are somewhat abstract objections to an abstract theory, however. The practical matter is that there is no known way to calculate the risk premium in a specific case. The closest one can come to such a calculation is where there are data for a large number of transactions that can be analyzed statistically, such as in the stock market. Here, the variability of a stock price is a reasonable surrogate for one type of risk, but no one has demonstrated a way to use this concept of ‘volatility risk’ to adjust the discount rate for a single-project investment. Except in a limited number of special cases, risk is more than volatility (see Section 9.8).

Making subjective estimates of the risk premium is not a good idea for two reasons: first, as pointed out earlier, there are generally no data or clear expertise in the matter; and, second, putting a risk premium on the discount rate obscures the outcome of the analysis. The influence of risk estimates working through a premium on the discount rate is too complex for intuitive understanding. All in all, discount-rate adjustment is not a good way to come to grips with probabilities and risk.

8.4 Risk analysis through simulation

> "In my experience, those who finance benefit-cost analyses have scant patience with (paying for a study of all for the small-probability risks involved with a project, even if some of the uncertainties involve catastrophic losses if they actually occur). The unfortunate conclusion is that if allowance is to be made for unfavourable contingencies, it will almost certainly have to be done quickly, roughly and cheaply. Thinking about practical ways to bring real-world project profiles closer to the goal of representing the expected values, I find myself gravitating more and more to a generalised advocacy of the use of simulation techniques (Monte Carlo) … while arguing in favour of maintaining a high degree of economy and simplicity in their application"
>
> Arnold Harberger, 1997

The only practical approach to financial and economic risk analysis is to use simulation. Simulation predicts the possible outcomes of the benefit-cost model, given the variables that influence those outcomes. It enables the analyst to give more comprehensive and realistic advice to the decision-maker. In the older deterministic method of benefit-cost analysis, the analyst offered a single figure for NPV, but it was always unclear what the probability of this single outcome was. The decision-maker did not know how much confidence to place in the figure (especially given the rather esoteric calculations that produced it) and therefore tended to make a subjective judgement.

Simulation shows the range of NPVs possible, given the factors that can vary, and provides an overview of the probabilities within that range. Decision-makers know there is risk in every decision. There are no guarantees. Sometimes the right decision doesn't turn out well because the changeable factors turn unfavourable. The decision-maker relies on the analyst to give as full and accurate a picture of the possible risks and rewards as possible. The simulation tools available for risk analyses are discussed in Chapter 9.

Best practice – analysing uncertainty

- Risk arises from uncertainty in the data. The analysis of incremental effects, and the economic analyses of an investment both contribute to assessing risk.
- There are three approaches to analysing financial and economic risk; expected values of scenarios, risk-adjusted discount rates, and risk analysis through simulation. Of these three, only simulation offers a reliable methodology for assessing overall risk.

9. Risk analysis

> If a man will begin with certainties, he shall end in doubts; but if he will be content to begin with doubts, he shall end in certainties.
>
> - Sir Francis Bacon, *The Advancement of Learning*, 1605

9.1 Introduction

Financial and economic risk analysis is a technique that enables us to determine how much risk there is in accepting or rejecting a particular project. We can also use it to compare the likely outcomes of two or more alternative projects. It is an important technique because it allows us to use data that are uncertain to obtain results that are a good picture of the likely outcomes. The technique takes into account possible variations in the costs and benefits that we may be aware of but that we ignore when we use single best-guess numbers in an everything-goes-according-to-plan analysis.

Each benefit-cost model generally has several variables that are subject to uncertainty. To use this model for risk analysis, you need a computer program for simulations. The computer runs the model over and over again, each time selecting a value for each variable. The program is not difficult to use - you simply instruct the computer to choose a value within a certain range and to select that value according to the stated probabilities. For example, suppose there are three possible values for a given variable in the range - 3 (50 per cent probability), 4 (30 per cent probability) and 5 (20 per cent probability) - and the computer runs the model 1,000 times. For about 50 per cent of the runs, the computer will choose 3 for the value of the variable; for 30 per cent of the runs, it will choose 4; and for 20 per cent, it will choose 5.

A major advantage of a computer simulation is that it can consider a number of uncertain variables (risks) simultaneously, choosing values according to the ranges and probabilities of each. This enables you to model the likely outcome of the benefit-cost analysis more or less realistically. In Chapter 7, we pointed out that sensitivity analysis is a limited technique because it can handle only one or two variables at a time and has to hold all the others constant. Risk analysis goes beyond this limitation by allowing all of the variables to vary at the same time. Their influence and interactions are then simultaneous, just as they are in the real world.

9.2 The steps of risk analysis

Benefit-cost analysis is best approached as a risk analysis because there is always some uncertainty in the data. The steps in risk analysis are the following:

1. Set up the basic model that will calculate NPV. This model is sometimes called the deterministic model because it uses a single deterministic value for each variable (see Chapters 2 and 6).

2. Link the uncertain variables in the model to information about their maximum and minimum values (range) and about the probabilities of various values within those ranges.

3. Run the model many times to obtain a large number of NPVs (to see what all the possibilities are) - that is, construct an investment results table (see Section 2.5).

4. Determine the frequency with which various NPVs occur in the results, and, on this basis, predict the likely range of the NPV and the probabilities of various NPVs within that range.

5. Using the decision rules, interpret this information to identify the best alternative investment or, if there is only one, to decide whether it is likely to be a good investment (see Section 9.7).

9.3 The mechanics of risk analysis

We cannot emphasize enough that setting up a good deterministic model before you think about risk is extremely important. Risk analysis is not a substitute for careful and detailed development of tables of costs, benefits and parameters.

Risk-analysis software builds on the underlying benefit-cost model. Once the deterministic model is working adequately, you use the software for two additional steps:

- selecting sets of values for the uncertain variables, according to specified probabilities, for each run of the benefit-cost model;
- using these sets of values to calculate the possible outcomes and analyze the results.

The first step, selecting sets of values for the uncertain variables, is based on sampling. Most risk-analysis programs use the *Monte Carlo* method (simple random sampling according to a specified probability distribution) or the *Latin Hypercube* method (stratified sampling); some use both. Generally, Latin Hypercube can accurately re-create the specified probability distributions in fewer iterations than Monte Carlo can and is therefore the best choice if your software can use either one. Each run of the program completely samples for all risk variables and recalculates the worksheet. This is called an *iteration***.** The whole procedure of many iterations is a simulation. The program is simulating the range and probabilities of the investment's outcomes in the real world.

Different software packages have somewhat different procedures. However, the basic operations that demand attention are the same:

- You must link the uncertain variables in the benefit-cost model to their range and probability data.
- You must specify the 'bottom line' of the benefit-cost model (the spreadsheet cell location of the NPV) so that the risk-analysis software can link it to the results table.

All the other operations are automatic. That is, the risk-analysis software calculates the NPV a large number of times (to see what all the possibilities are) and analyzes those results statistically and graphically.

9.4 Adjusting for the covariance of related risk variables

You must keep in mind that some risk variables might be correlated. For example, if NPV from a particular run of the benefit-cost model is based on a high value for 'total fish catch' and a high value for 'average price of fish,' then the NPV may be outside the plausible range in the real world. What normally happens is that a high fish catch goes with a low fish price and vice versa. If the outcome of your analysis is to be realistic, you must take these correlations into account.

You do this by instructing the computer to select values for the uncertain variables in each run that respect the specified correlations between them. If the computer selects a high value for total fish catch, for example, then it must select a value for fish price in accordance with the stated correlation. The computer selects an appropriate value for each variable sequentially. The value it selects for fish price is guided by not only the probable range of fish prices and the shape of probabilities within that range, but also the value for total fish catch already selected during this run and the correlation between the two variables.

Failure to take covariances into account can lead to large errors in judging risk. For example, in his pioneering study of 'Risk Analysis in Project Appraisal,' Pouliquen (1975) cited a project for which the risk of failure was 15 per cent when labour productivity and port capacity were treated as independent variables but about 40 per cent when their positive correlation was taken into account.

Software programs take different approaches to using information about covariances among the risk variables. Sometimes measuring the correlation coefficients is a job for an expert (see Chatterjee 1994). It is often unnecessary, however, to resort to comprehensive descriptions of statistical dependence in applied project work. Pragmatic methods of specifying approximate rank-correlation coefficients for the key pairs of risk variables are often adequate.

9.5 How many times does the model need to run?

Each time the benefit-cost model runs, it generates an NPV; eventually, there will be enough results to give a full and accurate picture of the likely outcome of the investment. The number of runs needed depends on how wide the ranges of the variables in the model are, and how predictable the values are within those ranges (how much they cluster around a central value).

After a certain number of runs, the results table will stabilise, and more runs will not significantly alter the distribution of the NPVs. Some risk-analysis programs have a built-in regulator. Runs continue until new runs alter the result by less than 1 per cent, for example. If the program does not have such a regulator then, after a reasonable number of runs, you have to check that the probability distribution of results has no obvious gaps and has settled down to a stable range and shape. This is quite important because it is certainly possible to get an incorrect picture of the investment's outcomes from too few runs.

9.6 Interpreting the results of the risk analysis

Risk analysis produces a list of NPVs, one for each run of the benefit-cost model. You can then analyze these NPVs statistically and graphically to see what the probabilities of various outcomes are. There are two types of graphs that show the probability distribution of the NPV. The first type is a *probability-density graph*, which shows the individual probability of each NPV (see Figure 9.6.1). The second type is a *cumulative-distribution graph***,** which shows how probable it is that the NPV will be lower than a particular value (see Figure 9.6.2). Both types of graphs are useful for communicating with the decision-maker.

As well as constructing the graphs to show the distribution of results, most simulation software calculates some useful numbers, including the likely range of the NPV (minimum to maximum), the key probabilities (such as the probability that the NPV will be greater than zero), and the expected value of the investment. Together, these factors guide the investment decision.

Figure 9.6.1: Probability-density distribution of net present values approximated by a normal curve

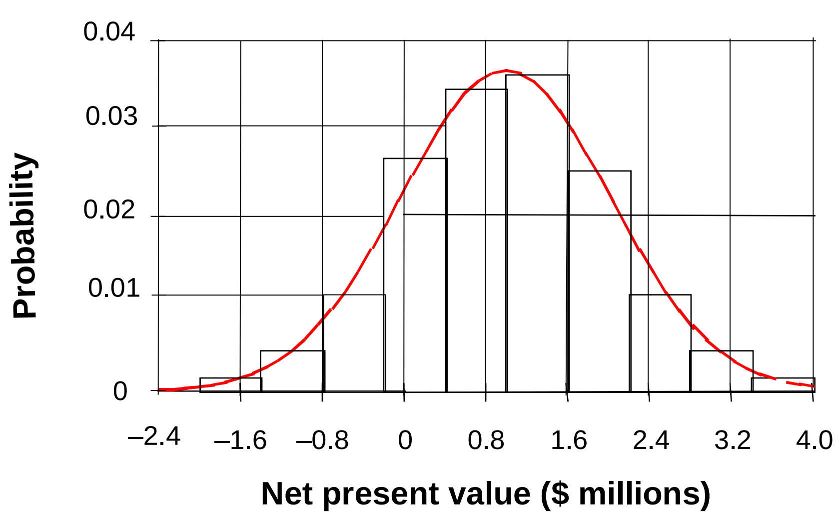

Figure 9.6.2: Cumulative-probability distribution of net present values approximated by a curve

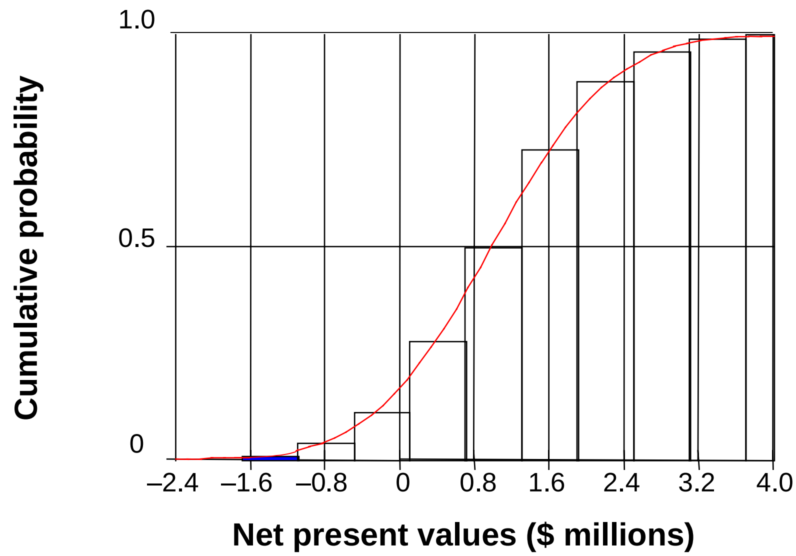

In the example shown in Figure 9.6.1 and 9.6.2, the range of possible NPVs is from about minus $2.4 million to plus $4 million. The most likely outcome (the mode of the distribution) is plus $1.3 million (most easily read from the probability-density graph, Figure 9.6.1). The probability of a loss (NPV < 0] is about 20 per cent (most easily read from the cumulative-distribution curve, Figure 9.6.2), and the expected value (the sum of all outcomes multiplied by their probabilities) is plus $0.997 million. Note that the expected value is the key figure for the decision maker, not the most likely value in the sense of the single value that is most likely to occur. These are quite different figures. For example consider a distribution of these values 6 (probability 0.4), 7 (probability 0.3) and 8 (probability 0.3). The most likely value is '6'. The expected value is (6 x 0.4) + (7 x 0.3) + (8 x 0.3) = 6.9.

9.7 Decision rules adapted to uncertainty

In Chapter 5 we considered decision rules in the context of deterministic benefit-cost models. Once we recognise uncertainty in the data, making decisions becomes less clear cut, although the principles are the same. The general rule is to choose the project with the highest ENPV. At the same time it is important to make risk transparent to the decision-maker. For example, two investments of public funds may have the same ENPV, but very different risk profiles. Project A may have possible high gains and possible high losses. Project B may have less spread in its possible outcomes. As long as the ENPVs are the same there is no immediate reason to choose one project over the other. On the other hand, if one is holding a portfolio of investments, then either project A or project B might have advantages in terms of improving the portfolio as a whole. Portfolio theory is beyond the scope of this guide. However, little is lost for the Federal government by this omission. It has such a large portfolio of investments in projects and programs that its best strategy in choosing between project A and B is rational risk neutrality. That is, simply, choosing the best outcome over a large portfolio.

Rules for rational decision making are illustrated by the following cases. In each case, both the cumulative-distribution graph and the probability-density graph are provided for comparison. The cumulative-distribution graph of the NPV is more useful for decisions involving alternative projects, whereas the probability-density graph of the NPV is better for indicating the mode of the outcomes of a single proposed project and for understanding concepts related to ENPV.

Decision rule 1: If the lowest possible net present value is greater than zero (see Figure 9.7.1) accept

the project.

Figure 9.7.1: Probability-distribution curves for a single project (positive net present value)

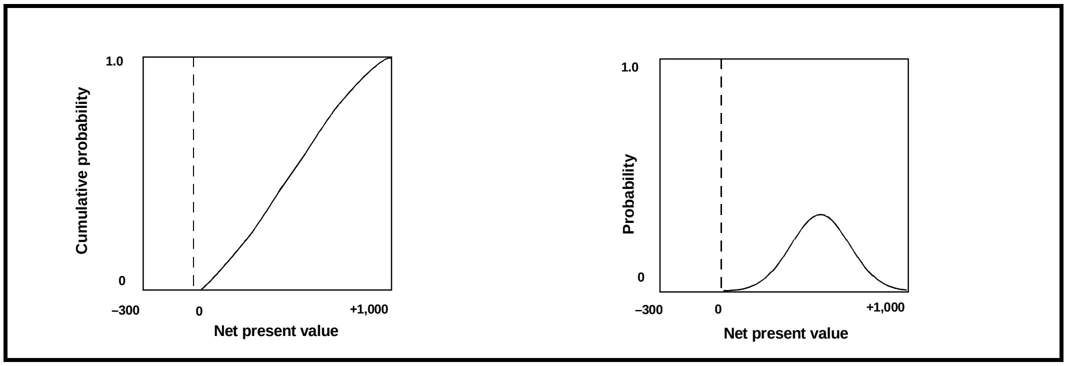

In Figure 9.7.1, the project shows a positive NPV even in the worst case. There is no probability of a negative return; therefore, the project is clearly acceptable.

Decision rule 2: If the highest possible net present value is less than zero (see Figure 9.7.2) reject the project.

Figure 9.7.2: Probability-distribution curves for a single project (positive net present value)

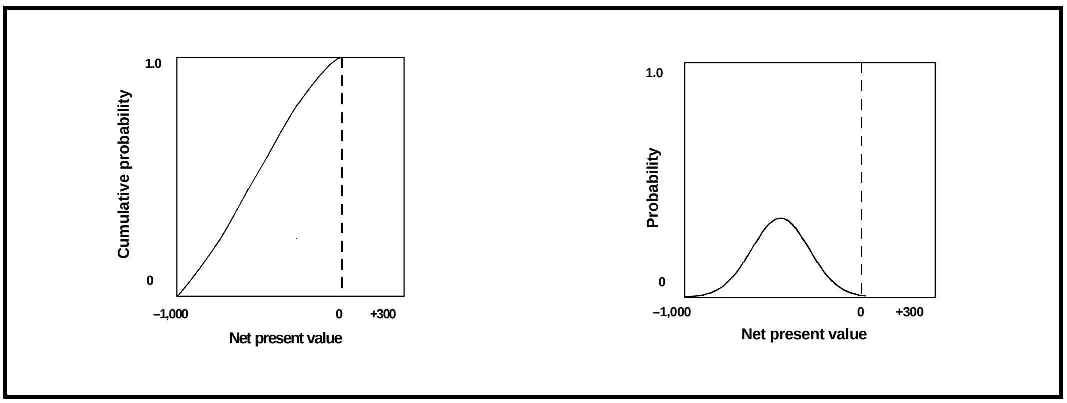

In Figure 9.7.2, the project shows a negative NPV even in the best case. Clearly, the project should be rejected.

Decision rule 3: If the maximum net present value is higher than zero and the minimum is lower (see Figure 9.7.3) accept the project if the **expected net present value** ***(the sum of all possible outcomes, each multiplied by its probability) is greater than zero. (Keep an eye on the risk of loss.)***

Figure 9.7.3: Probability distribution curves for a single project
(containing positive and negative net present values)

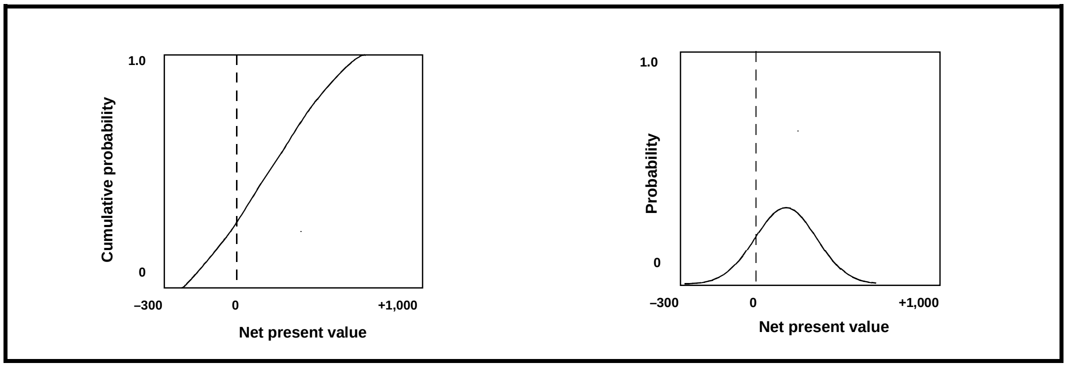

In Figure 9.7.3, the curves show some probability of a gain as well as some probability of a loss. The decision, therefore, depends on the ENPV and the risk tolerance of the investor. The rational decision-maker (neither risk-averse nor risk-loving), should accept the project if the ENPV is positive and reject it if the ENPV is negative.

Decision rule 4: If the cumulative-probability-distribution curves for two mutually exclusive projects do not intersect (Figure 9.7.4, left side), choose the option whose probability distribution is farther to the right.

Figure 9.7.4: Probability distribution curves for the NPVs two projects

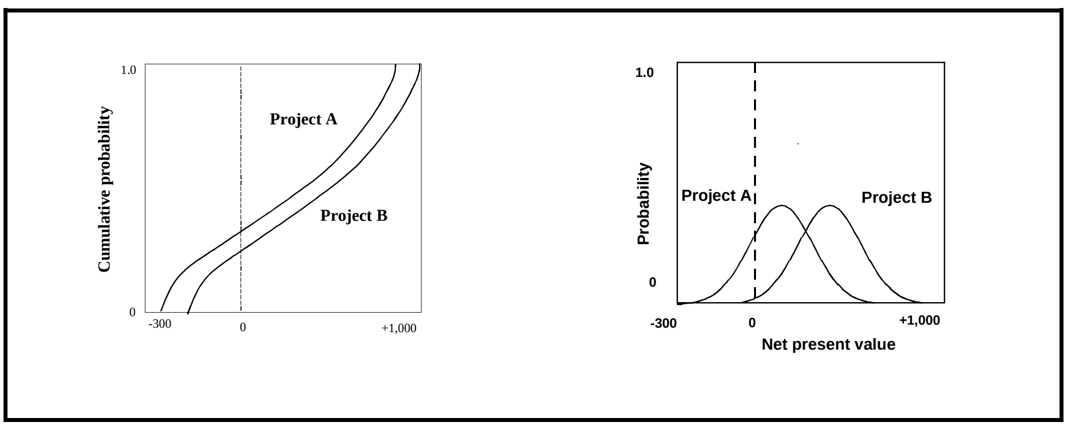

In Figure 9.7.4, the probability that any specified positive outcome will be exceeded is always higher for project B than it is for project A. The decision maker should, therefore, always prefer project B over project A.

Decision rule 5: If the cumulative-probability-distribution curves for two mutually exclusive projects intersect (see Figure 9.7.5), be guided by the expected net present value. If the ENPVs are similar, consider the risk profile of each project.

Figure 9.7.5: Probability distribution curves for the NPVs of two projects, one of which has a broader range of possible NPVs

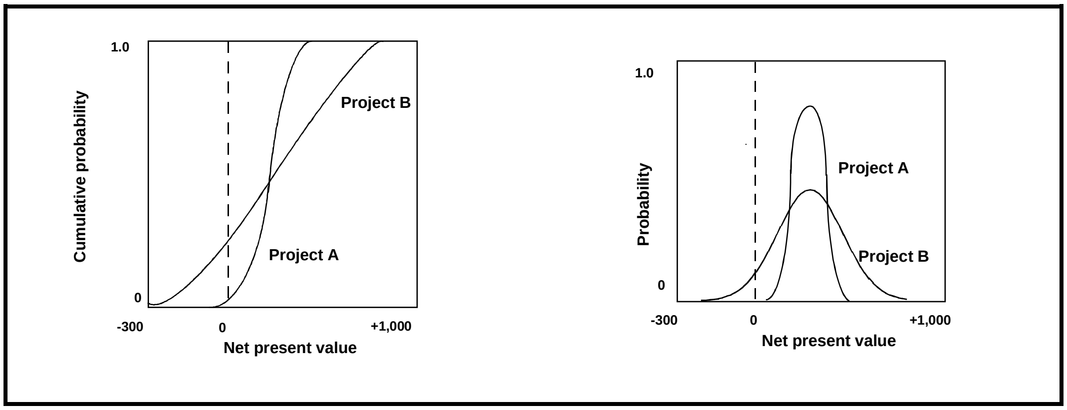

Risk-loving decision-makers might be attracted by the possibility of a higher return (despite the possibility of greater loss) and therefore might choose project B in Figure 9.7.5. Risk-averse decision-makers will be attracted by the possibility of lower loss and will therefore be inclined to choose project A.

9.8 Assessing overall risk

Two measures are particularly useful summaries of the overall level of risk in a public investment: the expected-loss ratio and the risk-exposure coefficient.

9.8.1 The expected-loss ratio

The *expected-loss ratio* is the absolute value of expected loss (all possible losses weighted by their probabilities) as a proportion of total expected value of all possible outcomes. Figure 9.8.1 shows all possible outcomes, each weighted by its probability. Area *A* under the curve of the NPV distribution represents this. Similarly, area *B* in Figure 9.8.2 represents the expected value of all losses. The risk can be thought of as the relationship between *B* and *A* (that is, area *B* divided by area *A*). If the area *B* is 15 units and *A* is 100 units, then the expected loss ratio is 0.15. If *B* is a large portion of *A*, the project is risky.

Figure 9.8.1: All possible outcomes for a project, each weighted by its probability

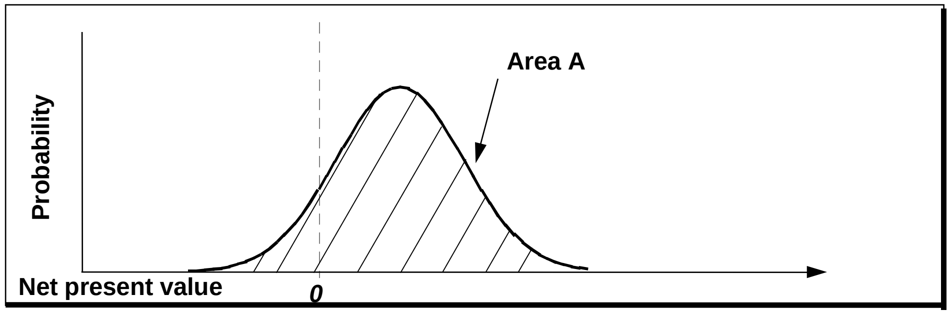

Figure 9.8.2: Expected value of the total possible losses for the project

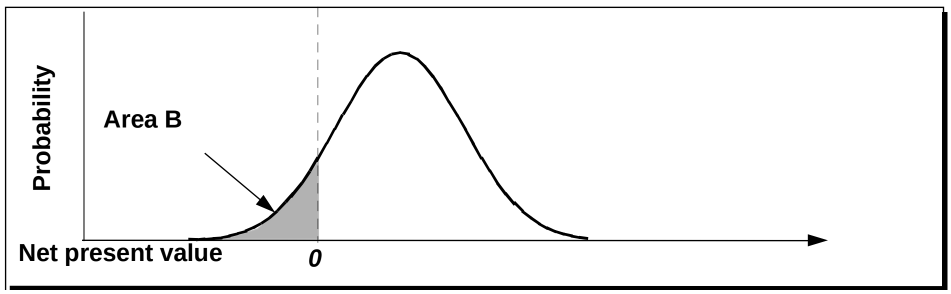

9.8.2 The risk-exposure coefficient

In many cases, the expected-loss ratio may be an adequate indicator of risk, but it does not capture all aspects of risk. Two projects can have the same expected-loss ratio but different levels of risk because one's outcomes are more spread out than the other's is or because more of the spread is in the negative-NPV area. Figures 9.8.3 and 9.8.4 show projects with the same expected-loss ratios but very different levels of risk according to other criteria.

Figure 9.8.3

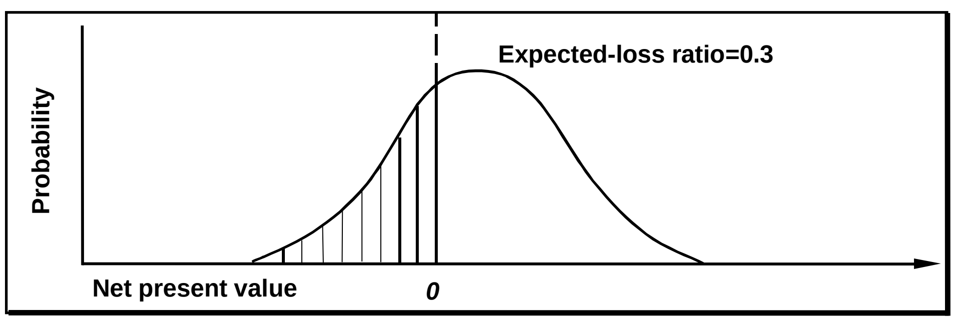

Figure 9.8.4

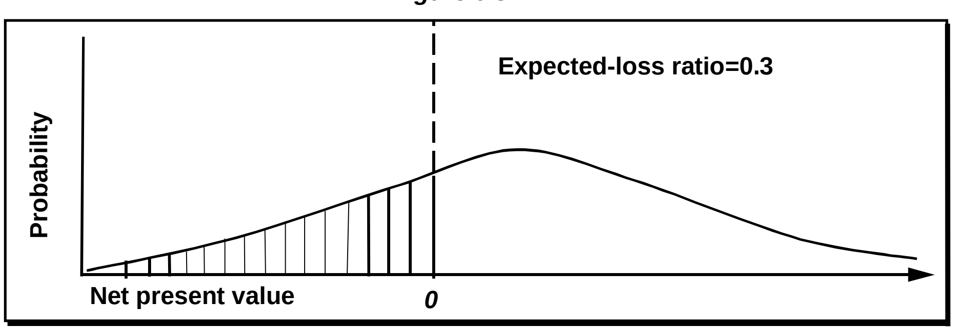

We need to look at two additional aspects of risk:

- How spread out (dispersed) are the possible outcomes (measured by **standard deviation**)?

- What proportion of the possible outcomes is on the loss side of the outcome distribution (that is, to the left of the NPV = 0)?

When we consider these two factors along with the expected-loss ratio, we obtain a risk-exposure coefficient (REC), a more complete measure of risk:

$$\mathrm{REC} = L_{\mathrm{E}}\,(\mathrm{SD})(D_{\mathrm{L}}/D) \qquad [6]$$

Where L_E is the expected loss ratio; SD is standard deviation of the outcome distribution; D_L is the distance on the NPV axis from the minimum value to zero; D is the distance on the NPV axis from the minimum to the maximum value. You may find the risk-exposure coefficient too mathematically complex to be intuitively appealing if you are dealing with a relatively simple ‘go’ or ‘no go’ decision on a single project. In that case, you might find the expected-loss ratio more useful. If you are comparing two or more alternatives and if those alternatives involve the investment of large resources, however, it is worth going the extra step to calculate the risk-exposure coefficient so that you can rank the projects according to risk.

9.9 The advantages and limitations of risk analysis

Some advantages of risk analysis are the following:

- It can rescue a deterministic benefit-cost analysis that has run into difficulties because of *unresolved uncertainties* in important variables.

- It can help *bridge the communications gap* between the analyst and the decision-maker. A range of possible outcomes, with probabilities attached, is inherently more plausible to a decision-maker than a single deterministic NPV. Risk analysis provides more and better information to guide the decision.

- It identifies where *action to decrease risk* might have the most effect.

- It aids the reformulation of projects to better suit the preferences of the investor, including *preferences for risk.*

- It induces careful thought about the risk variables and uses information that is available on ranges and probabilities to *enrich the benefit-cost data*. It facilitates the thorough use of experts.

The limitations of the risk analysis include the following:

- The *problem of correlated variables*, if not properly contained, can result in misleading conclusions.

- The use of ranges and probabilities in the input variables makes the *uncertainty* visible, thereby making some managers uncomfortable.

- If the deterministic benefit-cost model is not sound, a risk analysis might obscure this by adding a layer of probabilistic calculations, thereby creating a spurious impression of accuracy.

Best practice - financial and economic risk analysis

There are at least three meanings of *most likely outcome*: 'the deterministic value of the NPV' (the outcome when one assumes best-guess figures for each input); 'the mode of the probability distribution of NPVs;' and 'the expected value' (the sum of possible outcomes, each multiplied by its probability). The last is the best guide to the choice of investment option.

- Simulation techniques provide a realistic picture of overall risk in the project. The expected-loss ratio and the risk-exposure coefficient are useful measures of overall risk.
- Commercial software programs make risk analysis a relatively simple task, once the basic (deterministic) benefit-cost model has been constructed and information about variable ranges and probabilities has been collected.
- For situations where there is significant uncertainty, the following benefit-cost decision rules apply:

1. If the lowest possible NPV is greater than zero, accept the project.
2. If the highest possible NPV is less than zero, reject the project.
3. If the maximum NPV is higher than zero and the minimum is lower, calculate the ENPV. If the ENPV is greater than zero, accept the project. (Keep an eye on the risk of loss.)
4. If the cumulative-probability-distribution curves for two **mutually exclusive** projects do not intersect, choose the option whose probability distribution is farther to the right.
5. If the cumulative-probability-distribution curves for two mutually exclusive projects intersect, be guided by the ENPV. If the ENPVs are similar, consider the risk profile of each alternative.

10. Probability data

High arbiter chance governs all.

- John Milton, *Paradise Lost*, Book Two, 1667

In Chapter 4 we discussed some difficult-to-measure inputs to benefit-cost analysis. In this chapter, we extend that discussion and consider some general aspects of collecting data. Keep in mind that this guide provides the benefit-cost *framework*, but it can provide only a glimpse of what are needed in particular cases to measure the costs and benefits.

10.1 Types of risk variables

Three types of risk variables are used in benefit-cost analysis:

- *full-horizon variables* - Some variables are the same for each period of the analysis: once a value is selected, it is used throughout the benefit-cost model. For each run of the model, the risk-analysis computer program will select a different value within the plausible range, but only one value is used in each run. The social discount rate is an example. We know that it is stable over time.

- *single period variables* - Some variables have values that change over time within a known range, and the true value in each period is independent of the value in any other period. In this case, it is simplest to have a separate variable for each period of the analysis. An example is rainfall. We know that the annual rainfall in a particular location varies between 73 and 116 cm and that the probabilities of any particular value in this range are distributed roughly according to a normal curve. We also know that precipitation in any year is essentially independent of the precipitation in the previous year. In a benefit-cost model that has a 25-year investment horizon, therefore, we would have 25 values for the rainfall variable in the parameters table.

- *path variables* - Some variables change over time in a regular pattern. The value in one period is related in a systematic way to the value in the previous period. For example, the inflation rate in one year is likely to be within a certain range (up or down) from the rate in the previous year, and a trend, once established, tends to continue for some time. We know the starting rate of inflation - the rate in the current year. We would have 25 inflation rates in our benefit-cost model (as with the rainfall variable discussed above), but we would also have to program the model so that each time it runs it selects a different path of inflation rates for the investment period. The path selected must be in accordance with the rules of behaviour for this variable. This is not difficult programming, but it is beyond the scope of this guide to describe it in detail.

10.2 Using historical data

To do a risk analysis, you need to know the range of values each variable can take (minimum to maximum) and the probability distribution of values in that range. For example, if the price of a barrel of oil is between $12 and $32 and the probability of any value in that range is described by a uniform probability distribution (all values are equally probable), then the computer has all the information it needs to sample values of the price of oil to use in iterations of the benefit-cost model.

If historical data are available, you might use the maximum and minimum values that occurred in the past as an appropriate range for the values of the variable. Identifying the probability distribution within that range is more difficult. However, software can facilitate the task. Suppose you have monthly oil-price data for the past 10 years. The first step in identifying the probability distribution for oil prices is to group the raw data. How frequently do various oil prices occur? The raw data are grouped in a frequency histogram.

In some cases, it might be desirable to filter the raw data before exploring the frequency patterns. Software can filter out any data that fall outside a stated absolute or relative range - outside two standard deviations from the mean of the data, for example. This is useful when there are data outliers because of special circumstances that are unlikely to be repeated. We will not be surprised if the curve that best fits the data (see Figure 10.2.4) is a symmetrical one with a strong central mode - the normal, poisson or triangular distribution, for example.

Figure 10.2.1: A 'normal' curve fitted to a frequency histogram

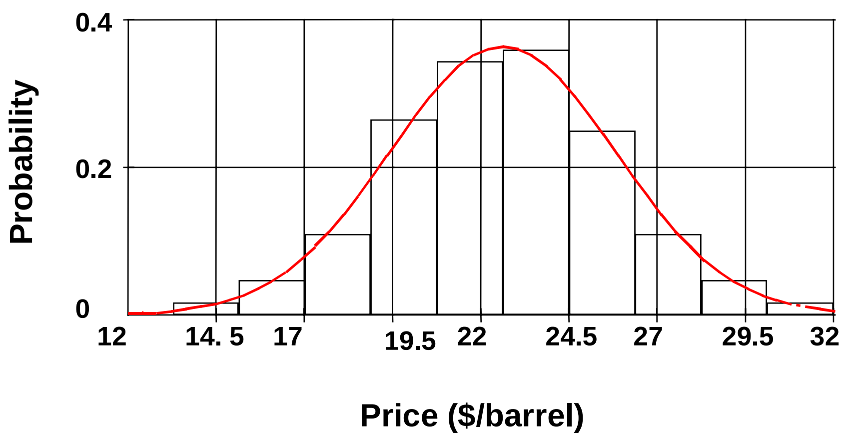

You can use your software to try all of the standard distributions in its repertoire and to list them in order of goodness of fit, using the chi-square test to rank the various curves by order of fit. You can then select the best curve to use for oil prices in your benefit-cost analysis. This selection is aided by a difference graph (Figure 10.2.2), which shows the differences between the expected oil price (based on the probability-distribution curve) and the actual historical price. The smaller these differences are, the better the curve fits the past data.

Figure 10.2.2: Differences between the expected price of oil (based on the probability-distribution curve) and the actual historical price

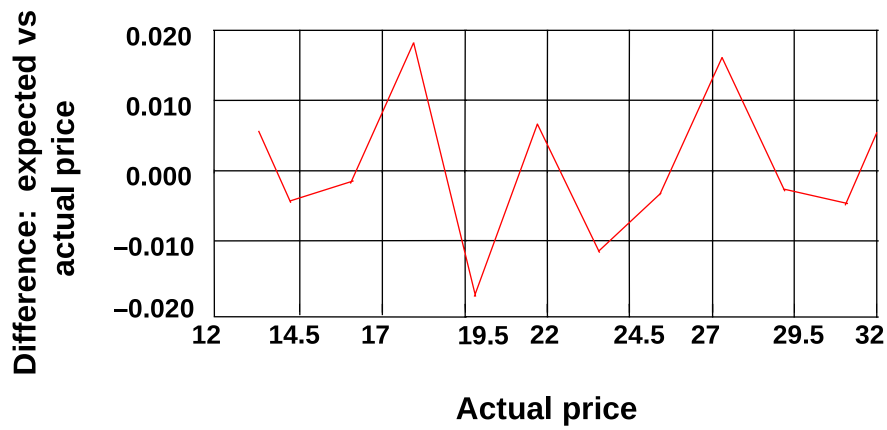

10.3 Expert judgement

If we do not have enough historical data to underpin an estimate of the range and probabilities of a particular variable, then we have to rely on expert (somewhat subjective) judgement. For example, the ‘take-up’ (participation) rate for a new program might be vital to its outcome, but there might be no direct experience to draw upon to forecast the likely rate. At one stage, the Government of Canada offered to match, dollar for dollar, all research funds raised by universities from the private sector. A benefit-cost analysis of this proposed program needed an estimate of the contributions likely to be gathered by the universities. There were no specific data available, but there were experts in the general field - people with a lot of experience in fundraising in general and in university/private-sector fundraising in particular.

A committee of these experts was asked to make informed estimates on the values of the variables to be used in the benefit-cost analysis. The committee used the **Delphi method.** The essence of this technique is that each member of the committee makes initial estimates based on the information provided. These initial estimates are then given to all the members and discussed, after which each member makes a second estimate. The discussion and re-estimation continue until the estimates converge.

10.4 Common probability distributions

Various levels of sophistication are possible in specifying the probability distributions for the input variables to the benefit-cost model. Most often, a fairly simple, straightforward approach is adequate. This means selecting probability-distribution shape in one of two ways:

- specifying a standard statistical shape, such as a flat, normal, triangular, or Poisson distribution; or

- specifying a step distribution, which notes the probabilities for each segment of the variable's range.

If you have historical data for the variable but are unsure which probability distribution they fit, you can use software that will analyze the data and show the best fit.

10.5 Risk preferences

In our discussion of expected values (see Section 9.7), we have assumed that the decision-maker is wholly rational and neither risk-loving nor risk-averse. For example, a 50/50 chance of gaining $10 will be valued at $5. This is a reasonable assumption for a decision-maker that feels no wealth constraint - the investment in question is small in comparison with his or her wealth.

Many decision-makers in the Government of Canada, however, do feel a wealth constraint. They will view a 50/50 chance of losing or gaining $1 billion with a lot more trepidation than a 50/50 chance of losing or gaining $5 million, although the NPV of each is the same. There are techniques to ascertain the ***utility function*** of the decision-makers that are risk-averse or risk-loving (rather than risk-neutral). A utility function simply gives mathematical form to the strength of the decision-maker's preferences for various possible outcomes and can be used to modify the NPV decision rules to more truly reflect these preferences.

In most cases, governments are risk-neutral. That is, they are rational decision-makers. A government has a large portfolio of projects and programs and therefore can act fully rationally with confidence that, on average, things will turn out well if the decision rules are strictly applied. In the case of a single large or politically sensitive project, however, the question of a wealth constraint might arise.

10.6 Common project risks

A variety of risks can affect an investment. Some common sources of risk are the following:

- *Investment lumpiness* - Can the waters be tested gradually, or is it all or nothing?

- *Timing* - What if the project is delayed? What if it takes longer than expected for the project to reach full production? Is there a best time to start the project?

- *Salvageability* - How much of the investment can be recouped if things go wrong?

- *Uncertain incremental effect* - What will the outputs of the project be?

- *Uncertain parameter values* - What discount rate and inflation rates are appropriate?

- *Volatile preferences* - Are the target beneficiary's needs or preferences unstable?

An investment decision is highly risky if a big investment has to be made without a chance to test the waters, delay is highly damaging, little can be salvaged if the project goes wrong, the likely output is uncertain, some of the key measurement parameters are uncertain, and the beneficiary's preferences are unstable.

Best practice - handling probability data

- Base the probability estimates on historical data where possible and on structured subjective estimates by experts where necessary.

- Treat the three types of risk variables (full-horizon, single period, path) appropriately in the benefit-cost model.

- Keep it simple. Don’t let the technique overwhelm the data.

11. Comparing options of different types against different criteria

> Some of the arguments in favour of confining benefit-cost analysis to efficiency considerations seem to rest on a misconception of the analyst's role in the policy-making process; such arguments appear to be based on the assumption that the economist is not competent to assign prices to distributional effects. One can agree entirely that the analyst should not attach his own valuation to such effects. It does not follow however that equity issues should be ignored in benefit-cost analysis or that the decision-maker cannot value distribution effects if given the information on which to base a judgement. Indeed, the analyst who avoids measuring or describing non-price effects, which may have an important influence on a decision, seems to be ducking their responsibilities. Moreover, the measurement of distributional effects by the analyst is not an easy task and may require as much work as measurement of efficiency effects.
>
> - Treasury Board, *Benefit-Cost Analysis Guide*, 1976

11.1 Issues of fairness

> Money is like muck, not good except it be spread.
>
> - Sir Francis Bacon, *The Essayes or Counsels, Civill or Morall*, 1625

Questions of fairness are among the most difficult issues in benefit-cost analysis. It is essentially a problem of multiple criteria and not a trivial one. Benefit-cost analysis generally assumes that everyone in the reference group takes the same point of view. This is fine when there is a single investor, but is shaky when the reference group is a whole economy.

The problem is that the Government of Canada has fairness objectives as well as efficiency objectives and they often clash. It wants to maximize Canadians' wealth and incomes, but it also wants the distribution of the net benefits of government projects and programs to be equitable. There is often no way to maximize both objectives at the same time.

Another factor is the possibility of a dead-weight loss in redistributing income or wealth. This loss arises partly from the administrative costs of collecting taxes and running programs and partly from disincentives caused by taxes that dampen economic activity. If we add to this taxation-induced loss the project-specific inefficiency losses that may arise when projects are justified by their distributional effects, the transfer of income might be expensive. Taking $100 away from person A so that we can improve person B's welfare by $25 is neither sensible nor equitable.

Nevertheless, most Canadians do not believe that a dollar of benefit to the rich should count the same as a dollar of benefit to the poor. In some sense, they value a dollar of benefit to the poor more. Taking this value into account in benefit-cost analysis, however, raises a host of difficulties.

For one thing, having a low income is not the same as being needy. For example, many students, retired people with substantial capital assets, and rural residents with low living expenses have relatively low incomes, but they may not be needy in the sense that justifies special transfers from the government. A straightforward income-per-household criterion is thus likely to be an unreliable guide to the need for social transfers. Governments have recognized this and have responded in various ways, such as giving everyone the right to certain basic services (free primary and secondary education, for example) and establishing a series of filters through which households must pass to achieve eligibility (access to public housing, for example). None of these responses, however, provides a clear and simple way to adjust benefit-cost calculations to take fairness into account.

11.1.1 Equity approach 1: Ignore distributional issues

One school of thought is that income redistribution is best accomplished through explicit transfers in cash and in kind, not through skewed capital investments that artificially favours particular groups of people. The idea is to make efficient investments to create the largest pie, which can then be divided as society wishes through instruments that do not involve high transaction costs or economic inefficiencies. In Canada, with its tradition of strong and stable government, there is a lot of force to this argument. In some other countries, though, capital investments may be one of the few plausible instruments for income redistribution.

In Canada, if a project or program is recommended on equity grounds, then its equity cost should be visible to decision-makers and to the public. The equity cost is the difference between the NPV of this project or program and the NPV of its most efficient alternative. Making this cost visible is the best safeguard against unreasonable demands by special-interest groups.

11.1.2 Equity approach 2: Use distributional weights

In the 1960s and 1970s, the use of distributional **weights** was in vogue among many benefit-cost analysts. The idea was quite simple. Both costs and benefits to low-income groups would be given extra weight: for example, $1 in benefits to a low-income group might be counted as $2.

The use of such weights foundered on two shoals: the weights were arbitrary, and they made the analysis opaque. Decision-makers could not trust the analyses because they did not know how its outcome had been affected by the inclusion of subjective weights to change the values of costs and benefits.

There were some attempts to infer weights from existing and past policy decisions, especially from income-tax scales. Income-tax scales, however, are set with an eye on efficiency and incentive effects, as well as equity and distributional concerns. Separating one factor from another proved impossible.

Nevertheless, some serious attempts were made to implement distributional weights. For example, the World Bank used them in its project appraisals for two decades, abandoning the approach as unworkable only as recently as 1995. The bank's analysts had often assumed that the marginal utility of income declines exponentially as income increases. This implies a very rapid shift in the weight accorded to a marginal dollar of benefit as one moves up the income or wealth scale.

One problem with this approach is that it implies that benefits should be shifted from those at the top of the income scale to those, say, three quarters of the way to the top, as well as shifted from those in the bottom half of the income range to those at the very bottom. This is not generally what people have in mind when they think about redistribution.

Most decision makers are more comfortable with a weighting scheme that gives a premium to benefits to people in a narrow range at the bottom of the income scale (for example, below a poverty line). For example, weights might start at 2.0 at zero income and move fairly rapidly to 1.0 at the 25th percentile, perhaps adjusted by family size.

11.1.3 Equity approach 3: Focus on basic needs

Concentrating redistribution on the poorest *quintile* (20 per cent) of income earners is an attractive option. Generally, however, the use of major capital investments that might help this group is impractical because the poorest of the poor tend to have few skills or to live in remote areas, where capital investment is inefficient. Siting projects where they otherwise would not go, local hiring quotas, and similar mechanisms are seldom effective.

Providing for the basic needs of such people is probably best done directly and in kind (free education and health care, for example), rather than indirectly through the side effects of capital investments by the Government of Canada. For a major project or program, however, it is possible to construct an index of family welfare in a region or community that would operate something like a system of weights. Such an index would have component indices of, say, access to education, access to medical care and sanitary facilities as defined by the UN Panel on Basic Needs, level of nutrition, and quality of housing. A level of 100 on the index might represent the national average.

Constructing an index is not technically difficult, but is nevertheless a job for an expert, not for the benefit-cost generalist.

Even when an index is available, one still faces the same question as with other kinds of weights - how much economic inefficiency is it worth to move the distribution index one point? Society's willingness to pay is unknown. The approach is probably best restricted to projects whose primary purpose is to alleviate poverty or to address one or more basic needs (education, health, nutrition, etc.).

Innovation in methods to handle distributional issues is a very high priority for benefit-cost analysis. An interesting example that is not applicable to Canada in a simple way, but shows the level of innovation possible, is the map of extreme poverty developed by Chile in the 1970s. The Chilean approach is to monitor neighbourhoods, rather than individual families. The index is used to guide social expenditures toward the very poor and to try to overcome traditional biases that result in middle-income earners being the principal beneficiaries of government programs.

11.1.4 Equity approach 4: Focus on visibility and transparency

The best way to handle distributional issues in benefit-cost analysis is to let the decision-makers decide. Distribution is primarily a political question. If the benefit-cost analyst compares the best alternatives, using good technique, then decision-makers will be able to weigh improvements in income distribution against any loss in economic efficiency.

To a large extent, this is simply a matter of showing what the costs and benefits are from several points of view. The analyst should set out a distributional chart or matrix, showing gains or losses on one axis and the identification of relevant groups on the other. Such a display should be included in all benefit-cost analyses, unless the distributional effects are few (even in this case, the text of the benefit-cost report should contain a section on distributional effects).

A potential difficulty, however, is that the matrix can become complicated if a large number of groups are involved. Some loss of detail is often acceptable, though, to keep the display understandable to the decision-maker.

11.2 Multiple objectives

> Things should be made as simple as possible, but not any simpler.
>
> - Albert Einstein, cited in *Readers Digest*, Oct. 1977

The benefit-cost analyst tries to express the value of all alternatives in terms of a single criterion with a single weight - that is, in terms of economic net benefit. This makes for a simple investment decision. In reducing everything possible to economic net benefit, however, the analyst may have ignored other important factors, either because they cannot be quantified or because they cannot be valued in dollars. In some cases, these factors are *so* different from the economic criterion as to force a choice between apples and oranges.

The choice between investment options is simple when they are the same types and their context is also identical. In the real world of decision making, however, nothing is quite that simple - a budget most likely has to be allocated among different types of investments; the context of each investment varies enough for the criteria weights to vary as well; scores on various criteria cannot be expressed in dollars or any single unit of measure (numeraire); and so on. Making comparisons becomes very difficult when the investment alternatives are really apples and oranges. In such a case, three factors can vary:

- the criteria applied to each project;
- the weights of the criteria; and
- the numeraires (units of value used in measuring costs and benefits)

A government invests in a variety of programs and projects within a single budget. For example, a department might have to decide whether to build a new road or a new office building. These are not just alternative investments, but alternative investments in different categories. Often, the decision is made in two stages: first, deciding how much of the budget to invest in roads and in office buildings; and, second, deciding on what roads to spend the roads portion of the budget and on what buildings to spend the buildings portion of the budget. Often, the first decision (how much to invest in roads) is made subjectively or politically; and the second (whether to invest in a particular road) is made on the basis of benefit-cost analysis.

11.2.1 Attribute scores and trade-off weights

When a decision involves investment alternatives in different categories, the relevant criteria and their weights vary from one alternative to another. Sometimes, though, the variations are not important enough to completely rule out direct comparisons. For example, suppose you had to prescribe a set of criteria and criteria weights for farmers to use when purchasing animals. Some criteria would be specifically applicable to one type of animal: a good dog is *friendly*; a good pig *increases its weight quickly*; a good horse is *swift* and *sure-footed*. Other criteria would be generally applicable, such as *age*, *health* and *likely profit*. As well, the importance of a particular criterion might vary, depending on the context. For instance, if the farmer has small children, the importance of friendliness in the dog might be greater.

In spending his or her budget, the farmer does not want to end up with a friendly pig, a sure-footed dog and a fat horse. Neither does the farmer want to end up with five dogs and no horses or pigs. If the farmer applies a fixed set of criteria to all animals, then this might be the result.

In summary, if you are faced with multiple objectives, you should approach the decision in two steps: assign resources for achieving each objective; then select investments to maximize the achievement of each objective, given its resources. If this is not possible and all investments must be considered and ranked together, then the best (approximately correct) procedure is to fix the criteria and the **trade-off** weights; standardize the scores on each criterion; and then use the scores and weights to rank the options.

If you are judging alternative investments on the basis of multiple criteria, some of which cannot be measured in dollars, then the best way to score the alternatives on each criterion is to define the range of achievement, using whatever unit of measurement makes sense, and then to standardize the scores on that criterion by expressing them as a percentage of the possible achievement (see Table 11.2.1).

Table 11.2.1: An example of standardized scores of variable criteria with weights

	Standard Score*	Weight †	Weighted standard score
Alternative 1			
Criterion A	0.3	0.7	0.3 × 0.7 = 0.21
Criterion B	0.5	0.3	0.5 × 0.3 = 0.15
		1.0	0.36
Alternative 2			
Criterion A	0.4	0.7	0.4 × 0.7 = 0.28
Criterion B	0.7	0.2	0.7 × 0.2 = 0.14
Criterion C	0.9	0.1	0.9 × 0.1 = 0.09
		1.0	0.51

* Each standard score must be between 0 and 1.0. †Each weight must be between 0 and 1.0; the weights for each alternative must total 1.0.

Suppose alternative projects to provide office accommodation are to be judged according to four criteria: NPV (dollars); location (average commuting time for staff); availability (earliest occupancy date); and healthiness (air exchange rate for the building).

Let's look a one criterion, location, measured as the average commuting time for staff. Suppose that the lowest possible average commuting time is 10 minutes and that the highest average commuting time that would be acceptable under any circumstances is 45 minutes. Also suppose that the staff preferences are a linear function - that is, the benefit from decreasing their average commuting time by one minute is the same whether the decrease is from four minutes to three minutes or from 17 minutes to 16, etc. The maximum *possible location benefit* would result from changing an average commuting time of 45 minutes to an average commuting time of 10 minutes, a gain of 35 minutes (score of 100). You can standardize an average commuting time of, say, 20 minutes by expressing it as a percentage of the maximum possible location benefit: A commuting time of 20 minutes represents a gain of 25 minutes and would receive a standardized score of (25/35) × 100 = 0.714.

To complete the example, suppose we decide that the location is worth a weight of $6,000 (or 6,000 points, if a dollar weight is inappropriate). When the average commuting time changes from 45 minutes to 10 minutes, the maximum possible gain, the benefit is 6,000 units. Therefore the location benefit of our example project is 0.714 × 6,000 = 4,284. This is a weighted standard score. Once we have the weighted standard scores on each criterion, we can add them to obtain a total standard score for the project, and we can compare projects on this basis.

11.2.2 Limits to the weight of non-economic factors

If the primary purpose of a project or program is economic benefit, then non-economic factors should not receive more than a small percentage, say 15 per cent, of the total weight.

Best practice – equity analysis

- Distributional issues are important to the Government of Canada and should be considered in-depth in each benefit-cost analysis. Even a simple analysis showing who benefits and who pays can often be helpful to decision-makers.

- There are no uncontentious ways to combine efficiency and equity objectives in the same set of figures, although attempts have been made to use various types of weights to this end.

- Distributional issues should be covered in every benefit-cost analysis but kept separate from the economic-efficiency analysis. If a recommendation to approve a particular alternative hinges on equity objectives, then the net cost of choosing the equity-based recommendation must be made visible to the decision-makers.

Best practice - Ranking by multiple objectives

- Sometimes, important factors cannot be expressed in dollars no matter how ingenious and skilful the analyst is. In this case the decision-maker needs other techniques to compare alternatives against multiple criteria.

- If not all the criteria apply to all the alternatives, or if the criteria have different importance (weight) in different cases, then the analyst has a particularly difficult task. A two-stage approach (allocating budget to categories and then, within categories, to cases) is generally the best way to handle the multi-criteria problem. In cases where this is not possible, the weighted-score approach, using trade-off weights and standardized scores against the various criteria, is best. If the primary purpose of a project or program is economic, then non-economic factors should not receive more than a small percentage, say 15 per cent, of the total weight.

12. Key best practices

A good benefit-cost analysis meets the following criteria:

- the objectives and priorities are clear;
- the best alternative ways of achieving the objectives are identified for analysis;
- the alternatives are defined in a way that enables fair comparison;
- the 'point of view' of the analysis is stated;
- assumptions and calculations are visible to the reader at every stage of analysis;
- benefits and costs are estimated in detail for every time period, without short cuts;
- the technical analysis is well done (in regard to discount rates, inflation adjustments, choice of decision rule, etc);
- uncertainty and risk are carefully considered;
- distribution effects are clearly set out (who pays, who benefits?); and
- the recommendation is well reasoned and gives fair consideration to all alternatives.

Appendix A: Glossary

Accounting price. Reflects the **economic** value of **inputs** and **outputs,** rather than a **financial** or market value. Synonymous with **shadow price** and **social price.**

***Adjustment factor*.** The percentage by which the **financial** price of an **input** or **output** must be raised or lowered to reflect its true **economic** value. Synonymous with *conversion factor*.

***Appraisal*.** A before the fact (ex ante) evaluation.

Asset. Anything of value, but especially physical assets, such as machinery or farmland, or monetary assets (which can be used to finance the purchase of physical assets).

Base case. The optimised without-project **scenario**. Not the same as *do nothing* or *status quo*.

Benefit-cost analysis. A procedure that evaluates the desirability of a program or project by weighing the **benefits** against the **costs**.

Benefit-cost ratio. The ratio of **benefits** to **costs**. It should be calculated using the present values of each, discounted at an appropriate accounting rate of interest. The ratio should be at least 1.0 for the project to be acceptable. Inconsistent benefit-cost ratios may arise because they are dependent on arbitrary accounting conventions.

***Budget-Year Dollars*.** Face value dollars of varying purchasing power (depending on when a transaction is undertaken). Synonymous with **nominal dollars** and *current dollars*.

Capital. Resources that will yield benefits gradually over time. Related to investment (in contrast to consumption). May be divided into physical and financial; fixed and working; etc. Sometimes defined more broadly to include human capital (for example, in regard to an education that yields benefits over time).

Cash flow. The funds generated or used by the project. Reflects the **costs** and benefits over time from a stated point of view.

Cash-flow statement. A financial statement that records the **cash flow** of a project or financial entity. Synonymous with *sources-and-uses-of-funds statement*.

***Certainty equivalent*.** See **expected value**.

***Constant dollars*.** Dollars of constant purchasing power. The units of purchasing power are fixed by stating the base year, for example, 100 in 1995 constant dollars. Constant purchasing-power units. A better term is **real dollars**.

Constant price. A price that has been deflated to real terms by an appropriate **price index**.

Consumer surplus. The value consumers receive over and above what they actually have to pay. Varies from one person to another and is measured by willingness to pay.

***Contingent valuation*.** A method of inferring the value of **benefits** and **costs** in the absence of a market. What people would be willing to pay to gain a benefit (or willing to accept in recompense for a loss) if a market existed for the good.

Cost. An expense related to purchase of **inputs**, including capital equipment, buildings, materials, labour and public utilities. Costs such as environmental damage or injuries to health are sometimes referred to as negative **externalities**.

Cost-effectiveness analysis. A type of analysis commonly used to compare alternative projects or project designs when the value of **outputs** (benefits) cannot be measured adequately in dollars. If it can be assumed that the benefits are the same for all alternatives being considered, then the task is to minimize the cost of obtaining them through cost-effectiveness analysis. Synonymous with *least-cost analysis*.

Crossover point. The value that equalizes the **net present values** of benefits and **costs**.

***Decision rule*.** A criterion for accepting or rejecting a project or for ranking investments in order of desirability.

***Delphi method*.** A technique for obtaining subjective judgmental values through iterative estimations by a group of experts.

Demand. Need or desire for a good or service. The need varies with person, price and circumstance, so demand is usually expressed in terms of the quantities demanded at various prices. The demand curve usually slopes downward, indicating that people will demand more at lower prices than at higher prices. Opposite of **supply**.

Depreciation. Not a term used in **benefit-cost analysis**. In other **financial** frameworks, depreciation is the allocation of the **cost** of an **asset** over time. This is necessary for a working estimate of production costs, but because rates of depreciation are usually determined primarily by legal and accounting requirements, the amount of depreciation often has little relationship to the actual rate of use or cost of replacement.

***Deterministic model*.** A **benefit-cost** model that uses single fixed values for each **input** (rather than a range of values and **probabilities**).

Distributional gain or loss. A change in the distribution of wealth or income.

Discounted cash flow. The **costs** and **benefits** (**cash flows**) discounted to present values to give a common basis for comparison.

Discounting. The process of adjusting future values to an equivalent present value at a stated point in time by a discount rate.

Discount rate. The interest rate at which future values are discounted to the present and vice versa. Either the **opportunity cost of capital** (applied to investment dollars) or the time preference for consumption (applied to consumption dollars).

Distortion. A difference between **market prices** and true values (**economic prices**).

Distributional effect. A change in the income or wealth of the people from whose point of view the **benefit-cost analysis** is done.

Economic. Having to do with the national economy, especially as in *economic value*. The value of a good or service to the country as a whole, as opposed to its private or commercial value.

Economic price. Price that reflects the relative value that should be assigned to **inputs** and **outputs** if the economy is to produce the maximum value of physical output efficiently. There is no consideration of income distribution or other non-efficiency goals in such a price. Synonymous with *efficiency price* and *true price*.

Economic rate of return. An **internal rate of return** based on **economic prices**.

Expected net present value (ENPV). The sum of all of the possible **net present values** multiplied by their **probabilities**.

Expected value*.** The sum of all possible outcomes, each multiplied by its ***probability. For example, if there are two possible outcomes, $100 and $200, and their respective probabilities are 0.3 and 0.7, then the expected value is ($100 × 0.3) + ($200 × 0.7) = $170. Synonymous with *certainty equivalent*.

Externality. A **benefit** or **cost** falling on third parties who normally cannot pay or be compensated for it through a market mechanism. An external benefit is a positive externality; an external cost is a negative externality. Externalities are not reflected in the **financial** accounts. For example, a project may harm the environment, train workers, or make it easier for other firms to get started in a related line of business, but these effects do not show up in the project's financial statements. For **economic** analysis, however, it is necessary to take such externalities into account and place a value on them.

Financial. Using **market prices** and taking a commercial point of view.

Financial rate of return. The **financial** profitability of a project. Usually refers to an annual return on net fixed **assets** or on investment but may refer to the **internal rate of return**, which is determined through **discounted cash-flow** analysis.

***First-year return*.** The net **cash flow** during one year of a project, including the one-year cost of **capital** invested. A measure that may indicate the best time to begin a project.

Fixed costs. The **costs** such as management salaries, interest and loan repayments that must be met, at least in the short term, regardless of production volume.

Hurdle rate. The **rate of return** the project must achieve to be acceptable.

Incremental. Additional or **marginal**.

***Index number*.** Any index calculated to compare an amount in one period with that in another, for example, growth of production, population. See **price index**.

Inflation. A general increase in **market price** levels (a fall in the general purchasing power of the currency unit).

Input. That which is consumed by the project (as opposed to the project's **output**). Usually refers to the physical inputs used by the project, including materials, **capital**, labour and public utilities. Inputs like environmental quality, foreign exchange and workers' health are usually termed **externalities**.

Internal rate of return (IRR). The yield or profitability of a project based on **discounted cash-flow** analysis. The IRR is the **discount rate** that, when applied to the stream of **benefits** and **costs** reflected in the **cash flow** of a project, produces a **net present value** of zero.

Investment horizon. The period over which **benefits** and **costs** will be compared.

Marginal. Last, in the sense of the last additional unit. For example, the marginal benefit is the value of one more (or one less) unit of **output**. Synonymous with **incremental**.

Marginal productivity of capital. The productivity of the last unit of investment that would be undertaken if all investment alternatives were ranked in descending order according to their **economic** profitability and the available funds were distributed until exhausted. More loosely, the profitability of the marginal project (the project that should receive the last dollar of investment).

Market price. (a) The price of a good in the domestic market (see **financial);** as opposed to the **economic price**, **efficiency price**, **shadow price** or **social price**; (b) the cost of a good, including indirect taxes and subsidies.

Market risk. The **risk** to which all enterprises are exposed through cycles in the economy. Unlike some other risks, market risk cannot be diversified away.

Model. A representation or **simulation** of a system or process showing how parameters, benefits and **costs** interact to produce a bottom-line result by which the project can be judged.

Multiplier. The ratio of a change in the total community income to the initial change in expenditure that brought it about.

Mutually exclusive. Alternatives that cannot be undertaken simultaneously: if one alternative is carried out, the other cannot be. The alternatives may be mutually exclusive because they represent alternative times of beginning the same project, because funds are limited, or because if one is carried out the other will not be required (for example, a choice between a thermal and a hydro power station).

National parameter. A **shadow price** (or **accounting price**) that is the same for all projects in the country. In most cases, the shadow price for foreign exchange and the premium on savings over consumption are national parameters.

Net present value (NPV). The net value of an investment when all **costs** and benefits expressed in standard units of value (**numeraires**) are summed up. Synonymous with net present worth.

Nominal dollars, nominal prices. Prices prevailing in a particular year. Synonymous with **budget-year dollars**.

Non-tradable. Referring to a good that cannot be exported (for example, building foundations, haircuts).

Non-traded. Referring to an inherently **non-tradable** good or to **tradable** good that for **economic** or policy reasons is neither imported nor exported.

Numeraire. The standard unit of value that makes it possible to add and subtract **costs** and benefits that are otherwise expressed in unlike units. For example, apples and oranges, as everyone knows, should not be added up. But if they are expressed in terms of a common numeraire, such as pieces of fruit, kilograms or dollars, it is then possible to say that we have 20 pieces, three kilograms, or $4 worth of fruit. Common numeraires in **benefit-cost analysis** are dollars of investment, dollars of consumption, and dollars of foreign exchange.

Operational and maintenance costs. The recurring **costs** for operating and maintaining the value of physical **assets**.

Opportunity cost. The value of something foregone. For example, the direct opportunity cost of a person-day of labour is what the person would otherwise have produced had the person not been taken away from his or her best alternative occupation to be employed in the project.

Opportunity cost of capital. The best alternative return foregone elsewhere by committing **assets** to the project.

Option. The opportunity to invest in a particular program, project or course of action.

Output. That which is produced. Usually refers to the physical product of the project. Other effects of the project, such as housing for workers, employment, training of labour, and foreign-exchange savings, are usually called **externalities**.

Payback period. The time required for the cumulative present value of benefits to become equal to the cumulative present value of **costs**.

Present value. A future value discounted to the present by the appropriate discount rate.

Price index. The market value of a fixed basket of goods and services at one date divided by the market value of the same basket at some base date. Subtracting 1.0 from the index gives the decimal equivalent of the percentage increase in prices between the two periods. Useful in measuring rates of **inflation**.

Probability. The quantified likelihood of something occurring.

Probability distribution. A graphic representation of the likelihood of something occurring.

Producer surplus. The value a producer receives over and above the minimum payment needed to continue to supply the good.

Rate of return. The profitability of a project. A shorthand term usually applied in **economic** analysis to the internal **economic rate of return** and in **financial** analysis to the annual return on net fixed **assets** or to the internal **financial rate of return** (it is important to specify which).

Real dollars, real prices. Standard units of purchasing power, defined by stating a base year.

***Residual value*.** The **market value** of an **asset** at the **investment horizon**.

Risk. The degree to which outcomes are uncertain. The extent of possible variation in the outcome.

Risk analysis. A **benefit-cost analysis** that recognizes the simultaneous variation of the values of several **inputs**, according to specified ranges and **probabilities**, and analyzes the resulting variability in the bottom line.

***Risk variable*.** A variable in **risk analysis**, chosen because of its likely importance to the outcome of the analysis.

Salvage value. The **residual value** of an **asset** at the **investment horizon**.

Scale. The size of a project.

Scenario. An outline or portrait of a possible future; usually portrays unfolding events, rather than being static in time.

Sensitivity analysis. An examination of the effect that a change in a single variable (parameter, **cost** or benefit) has on the outcome of a project.

Shadow price. The true or **economic value** of a good (as opposed to the **market price,** which might be distorted). Synonymous with **accounting price** and **social price**.

Simulation. A mathematical model that sets out a system of interacting parameters, **costs** and benefits to predict the likely outcome of an investment.

Social net present value. The **net present value** of a project, calculated using true or **economic values** (**social prices** or **shadow prices**).

Social price. A price that reflects the true value to the country of **inputs** and **outputs** of the project. Synonymous with **accounting price** and **shadow price**.

***Standard deviation*.** A statistical measure of how spread the values are in a distribution.

Supply. Willingness to provide. Willingness varies with supplier, price, and circumstance, so supply is usually expressed in terms of the quantities that would be supplied at various prices. The supply curve usually slopes upward, indicating that suppliers will supply more at higher prices than at lower prices. Where economies of scale exist, however, the supply price may drop as scale increases over the range where such economies prevail. Opposite of **demand**.

Switching value. The value of an **input** that reverses the ranking of two alternative projects. For example, alternative A will produce shoes with sophisticated modern equipment and very few workers, whereas alternative B will consist of a network of small workshops employing many poor artisans and very little capital equipment. If income going to the poor is given a **weight** of up to 1.5, alternative A has a higher **rate of return**. If income going to the poor is given a weight of greater than 1.5, however, alternative B has the higher rate of return. Thus, 1.5 is the switching value.

Tradable. Referring to a good that could be traded internationally in the absence of restrictive trade policies.

Trade-off. The give and take involved in compromise or deal making; the negatives that come along with the positives and vice versa.

Transfer payments. Payments that redistribute wealth but do not use up resources or create them.

Weight. A factor that, when multiplied by the value to be **weighted**, adjusts that value to reflect certain considerations.

Weighted. Adjusted to reflect relative importance.

Willingness to pay. What consumers are willing to pay for a good or service. Consumers willing to pay substantially more than the actual market price enjoy a consumer surplus (the amount they would pay minus the amount they actually have to pay).

Appendix B: Questions to ask about a benefit-cost analysis

A quick guide

1. Is the problem or opportunity clearly stated? Is there a compelling rationale for the federal government acting in this situation? Are the objectives clear and coherent?

2. Is the analysis set out separately from the point of view of each important actor?

3. Are the alternatives defined in a fair and comparable way? Are the important alternatives analyzed?

4. Is this an open and transparent analysis? Is each stage of the analysis set out so that you can follow the reasoning and the numbers?

5. Are the likely incremental effects of the project or program alternatives well analyzed?

6. Are the costs and benefits of these effects measured well and set out in detail over the full life of the project? Are likely changes in relative prices taken into account or does the analyst take short cuts?

7. Are inflation adjustments and discounting done separately? Are the price index and discount rate the appropriate ones?

8. Does the analysis take into account uncertainty in the data and risk in the investment?

9. Does the analysis describe who pays and who benefits?

10. Does the analysis make a reasoned recommendation and give a fair showing to the alternatives it does not recommend?

Appendix C: Selected readings

General readings

Australia, Department of Finance. 1991. *Handbook of Cost-Benefit Analysis.* Australia, Department of Finance, Canberra, Australia.

---1992. *Introduction to Cost-Benefit Analysis for Program Managers*. Australia, Department of Finance, Canberra, Australia.

---1993. *Value for Your IT Dollar: Guidelines for Cost-Benefit Analysis of Information Technology Proposals.* Australia, Department of Finance, Canberra, Australia.

Canada, Treasury Board of Canada Secretariat and Consulting and Audit Canada. 1995. *Benefit-Cost Analysis Guide for Regulatory Programs*, Treasury Board of Canada Secretariat, Ottawa, ON, Canada.

Gramlich, E.M. 1981. *Benefit-Cost Analysis of Government Programs*. Prentice-Hall, Englewood Cliffs, NJ, USA.

Great Britain, Overseas Development Administration. 1988. *Appraisal of Projects in Developing Countries: A Guide for Economists.* HMSO, London, UK.

Harberger, A.C. 1997. *Economic Project Evaluation: Some lessons for the nineties*, Canadian Journal of Program Evaluation, Special Edition.

Jenkins, G., and Harberger, A.C. 1995. *Cost-Benefit Analysis of Investment Decisions.* Harvard Institute for International Development, Harvard University, Cambridge, MA, USA.

Sassone, P.G. 1978. *Cost-Benefit Analysis - A Handbook*. Academic Press, New York, USA.

Stanbury, R., and Vertinski, I. 1989. *Guide to Regulatory Impact Analysis*. Office of Privatization and Regulatory Affairs, Ottawa, Canada.

Sugden, R., and Williams, A. 1985. *The Principles of Practical Cost-Benefit Analysis*. Oxford University Press, Oxford, UK.

Van Pelt, M., and Timmer, R. 1992. *Cost-Benefit Analysis for Non-Economists*, Netherlands Economic Institute, Amsterdam, Netherlands.

Watson, K. 1997. *Cost-Benefit Analysis in the nineties*, Canadian Journal of Program Evaluation, Special issue.

World Bank. 1996. *Economic Analysis of Investment Operations.* World Bank, Washington, DC, USA.

World Bank and Economic Development Institute. 1991. *The Economics of Project Analysis: A Practitioner's Guide.* World Bank, Washington, DC, USA.

Defining fair options

Public Works and Government Services Canada. 1995. 2.0 "Real Property Investment Strategies;" and 3.0 "Policies and Practices." *Real Property Investment Guide*. Public Works and Government Services Canada, Ottawa, Canada.

Transport Canada 1994. 4.0 "Option Identification;" 5.0 "A Common Frame for Comparison." *Guide to Benefit-Cost Analysis.* Transport Canada, Ottawa, Canada.

Being clear on the point of view

Boardman, Vining, A., and Waters,W.G., 1993, "Cost and Benefits through Bureaucratic Lenses*:* Example of a Highway Project", *Journal of Policy Analysis and Management*, 13(3).

Inflation adjustments and discounting

Burgess, D.F. 1981. "The Social Discount Rate for Canada: Theory and Evidence." *Canadian Public Policy,* **7**, Summer, pp. 383-344.

Jenkins, G.P. 1981. "The Public Sector Discount Rate for Canada: Further Implications." *Canadian Public Policy*, **7**, Summer, pp. 399-407.

Lowenstein, G., and Elster, J., eds. 1992. *Choice over Time*. Russell Sage Foundation.

Watson, K.F. 1992. "The Social Discount Rate for Canada." *Canadian Journal of Program Evaluation*, **7**(1).

Decision rules

Bierman, H., and Smidt, S. 1993. *The Capital Budgeting Decision: Economic Analysis of Investment Projects*. Maxwell Macmillan Canada, Toronto, Canada.

Watson, K.F., 1990. "Using the Internal Rate of Return in Public Sector Project Evaluations: A Comment." *Canadian Journal of Program Evaluation,* **5**(1).

Multipliers

Schwartz, H. 1982. *A Guide to Regional Multiplier Estimation*. Department of Regional Economic Expansion, Ottawa, Canada.

Public Works and Government Services Canada. 1995. *Real Property Investment Guide*. Public Works and Government Services Canada, Ottawa, Canada.

Time savings

Culley, E.K., and Donkor, F. 1993. *Valuation of Passenger Travel Time Savings.* Transport Canada, Ottawa, Canada.

Value of safety, health and life

Cropper, M.L., et al. 1990. "Discounting and the Evaluation of Life Saving Programs," *Journal of Risk and Uncertainty*. **3**, pp. 369-79.

Keeney, R. 1994. "Decisions About Life-Threatening Risks." *New England Journal of Medicine*, **331**, pp. 193-196.

Cropper, M.L. 1992. "Rates of Time Preference for Saving Lives." *American Economic Review,* May.

Lawson, J.J. 1989. *The Valuation of Transportation Safety*. Transport Canada, Ottawa, Canada.

Nathwani, J.S., et al, 1990. "Safety, Social Well-Being and Its Measurement." Institute for Risk Research, University of Waterloo, Waterloo, Canada. Paper No. 21.

Wildavsky, A. 1980. "Richer Is Safer." *Public Interest*. **60**, pp. 23-39.

Environment

Pearce D.W., and Markandya, A. 1989. *Environmental Policy Benefits: Monetary Valuation*. Organisation for Economic Co-operation and Development, Paris, France.

N. Hanley, N., and Spash, C.L. 1993. *Cost-Benefit Analysis and the Environment.* Edward Elgar Publishing Ltd., Brookfield, VT, USA.

Job creation

Kuo, C.Y. 1997. "Evaluating the Social Cost of Job Creation." *Canadian Journal of Program Evaluation.*

Foreign exchange

Industry Canada and the Centre for the Study of International Economic Relations, University of Western Ontario. 1986. *The Shadow Price of Foreign Exchange in the Canadian Economy.* Industry Canada, Ottawa, Canada

Approaches to uncertainty

Baird, F.B. 1989. *Managerial Decisions under Uncertainty*. John Wiley & Sons, New York, NY, USA.

Breyer, S. 1993. *Breaking the Vicious Circle - Toward Effective Risk Regulation*. Harvard University Press, Cambridge, MA, USA.

Center for Risk Analysis, Harvard School of Public Health. 1994. *A Historical Perspective on Risk Assessment in the Federal Government*. Harvard University Press, Cambridge, MA, USA.

Dawes, R.M. 1988. *Rational Choice in an Uncertain World*. Harcourt Brace Jovanovich, New York, NY, USA.

Dixit, A.K., and Pindyck, R.S. 1994. *Investment under Uncertainty*. Princeton University Press, Princeton, NJ, USA.

Glickman, T.S., and Gough, M., eds. 1990. *Readings in Risk*. Resources for the Future, Washington, DC, USA.

Morgan, M.G., and Henrion, M. 1993. "*Risk Analysis* and Management." *Scientific American*, July, p. 3241.

Morgan, M.G. 1990. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge University Press, Cambridge, UK.

Morgan, M.G. et al. 1992. "Communicating Risk to the Public." *Environmental Science and Technology*. **26**(11), pp. 2048-2056.

Society for Risk Analysis. 1997. *Risk Analysis*. Plenum Publishing, NewYork, NY, USA. (Quarterly.)

Reutlinger, S. 1970. *Techniques for Project Appraisal under Uncertainty*. World Bank, Washington, DC, USA. World Bank Staff Occasion Paper, No. 10.

National Academy of Sciences and National Research Council. 1994. *Science and Judgement in Risk Assessment*. National Research Council, Washington, DC, USA.

Zeckhauser, R.J., and Viscusi, W.K. 1990. "Risk within Reason." *Scientific American*, **248**.

Risk analysis

Anderson, J.R. 1989. "Re-considerations on Risk Deductions in Public Project Appraisal." *Australian Journal of Agricultural Economics*, **33(**2), pp. 136-140.

Arrow, K.J., and Lind, R.C. 1970. "Uncertainty and the Evaluation of Public Investment Decisions." *American Economic Review*, **60**, pp. 364-378.

Carter, E.E. 1972. "What Are the Risks in Risk Analysis?" *Harvard Business Review.* July-August.

Fairley, W., Jacoby, H.D. 1975. "Investment Analysis Using the Probability Distribution of the Internal Rate of Return." *Management Science*, **21**(12), pp. 1428-1437.

Great Britain, Overseas Development Administration, 1990. *Risk Analysis - A Guide to Its Use in Project Appraisal*. Overseas Development Administration, London, UK.

Hertz D.B., and Thomas, H. 1983. *Risk Analysis and Its Applications*, John Wiley & Sons. New York, NY, USA.

---1984. *Practical Risk Analysis: An Approach through Case Histories*

Hull, J.C. 1980. *The Evaluation of Risk in Business Investment.* Pergamon, Oxford, UK.

National Academy of Public Administration. 1995. Chapter 3, Part 1 "Analysing Risks." and Part 2 "Using Risk Analysis in Regulatory Decisions." *Setting Priorities, Getting Results: A New Direction for the Environmental Protection Agency*. World Bank, Washington, DC, USA.

Pouliquen, L.Y. 1975. *Risk Analysis in Project Appraisal*. World Bank, Washington, DC, USA. World Bank Staff Occasional Paper, No. 11.

Wiliams, T. 1995. "A Classified Bibliography of Recent Research Relating to Project Risk Management." *European Journal of Operations Research*. **85**, pp. 18-35.

Sensitivity analysis

Baird, B. 1989. Chapter 11, "Sensitivity Analysis," *Managerial Decisions Under Uncertainty.* John Wiley & Sons, New York, NY, USA. pp. 352-375.

Eschenbach, T.G. 1992. "Quick Sensitivity Analysis for Small Projects and Feasibility Studies." *AACE Transactions.* **L-6**, pp. 1-8

Rappoport, A. 1967. "Sensitivity Analysis for Decision Making," *Accounting Review*, July, pp. 441-456.

Watson, K. 1995. "Sensitivity Analysis in Outcome Evaluations: A Research and Practice Note." *Canadian Journal of Program Evaluation*. **10** (2).

Probabilities

Hampton, J.M., et al. 1973. "Subjective Probability and Its Measurement." *Journal of the Royal Statistical Society*. **136**, pp. 21-42.

Henley, E.J., and Kamamoto, H. 1992. *Probabilistic Risk Assessment*. IEEE Press, Piscataway, NJ, USA.

Kabus, I. 1976. "You Can Bank on Uncertainty." *Harvard Business Review*, May-June.

Kachelmeier, S.J., and Shebata, M. "Examining Risk Preferences under High Monetary Incentives." *American Economic Review*. **82**, pp. 1120-1141.

Keeney, R.L., and von Winterfedlt, D. 1991 "Eliciting Probabilities from Experts in Complex Technical Problems." *IEEE Transactions in Engineering Management*, **38 (**3), pp. 191-201.

Law, A.M., and Kelton, D. 1991. *Simulation Modelling and Analysis*. McGraw-Hill, New York, NY, USA.

Spetzler, C.S. 1975. "Probability Coding in Decision Analysis." *Management Science*. **22**, pp. 3

Transport Canada and Consulting and Audit Canada. 1995. *A Guide for Assessing the Benefits of Reducing the Risk of Marine Oil Spills*. Transport Canada, Ottawa, Canada.

Walpole, R.E. and Myers, R.H. 1993. *Probability and Statistics for Engineers and Scientists*. Macmillan, Toronto, Canada.

Income distribution

Australia, Department of Finance. 1991. Chapter 7, "Analysing Distribution Effects." *Handbook of Cost-Benefit Analysis*, Australia, Department of Finance, Canberra, Australia.

Londero, E. 1987. *Benefits and Beneficiaries: An Introduction to Estimating Distributional Effects in Cost-Benefit Analysis*. Inter-American Development Bank, Washington, DC, USA.

Nwaneri, V.C. 1970. "Equity in Cost-Benefit Analysis: A Case Study of the Third London Airport." *Journal of Transport Economics and Policy*.

Weisbrod, B.A. 1968. "Deriving an Implicit Set of Governmental Weights for Income Classes." *In* Chase, S.B. ed. *Problems in Public Expenditure Analysis*. The Brookings Institute.

Multiple objectives

Baird, B.F. 1989. Chapter 13, "The Art of Decision Making with Multiple Criteria," *Managerial Decisions under Uncertainty*. John Wiley & Sons, Toronto, Canada.

Clemen, R.T. 1990. Chapters 15 and 16, "Conflicting Objectives," *Making Hard Decisions.* PWS-Kent Publishing, London, UK.

Finney, C.E. 1990. "A Consultant's Guide for the Economic Ranking of Public Sector Projects." *Project Appraisal*, **5** (1).

Keeney, R., and Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. John Wiley & Sons, New York, NY, USA.

Kirkwood, C.W. 1992. "Multi-Attribute Evaluation." *Management Science*, **38**, pp. 819-826.

Public Works and Government Services Canada. 1995. Chapter 8, "Analysis of Non-Financial Factors," *Real Property Investment Guide*. Public Works and Government Services Canada, Ottawa, Canada.

Criticisms and limitations of benefit-cost analysis

Carter, E.E. 1972. "What Are the Risks in Risk Analysis?" *Harvard Business Review*, July-August, pp. 72-82.

Downs, G.W., and Larkey, P.D. 1986. *The Search for Government Efficiency: From Hubris to Helplessness*. Temple University Press, Philadelphia, PA, USA.

Gutin, J. 1991. "At Our Peril: The False Promise of Risk Assessment." *Greenpeace*. March-April.

Latham, G.P., and White, G. "The Futility of Utility Analysis." *Personnel Psychology*. **47**, pp. 31-46.

McCloskey, M. 1994. “Problems with Inappropriately Broad Use of Risk Terminology.” *Comparative Risk Bulletin*. January.