1 Benefit-Cost Analysis: Introduction and Overview

Introduction

Social benefit-cost analysis is a process of identifying, measuring and comparing the social benefits and costs of an investment project or program. A program is a series of projects undertaken over a period of time with a particular objective in view. The project or projects in question may be public projects – undertaken by the public sector – or private projects. Both types of projects need to be appraised to determine whether they represent an efficient use of resources. Projects that represent an efficient use of resources from a private viewpoint may involve costs and benefits to a wider range of individuals than their private owners. For example, a private project may pay taxes, provide employment for the otherwise unemployed, and generate pollution. These effects are termed social benefits and costs to distinguish them from the purely private costs and returns of the project. Social benefit-cost analysis is used to appraise private projects from a social viewpoint as well as to appraise public projects.

It should be noted that the technique of social benefit-cost analysis can also be used to analyse the effects of changes in public policies such as the tax/subsidy or regulatory regimes. However a very broad range of issues can arise in this kind of analysis and, for ease of exposition, we adopt the narrower perspective of project analysis in this study.

Public projects are often thought of in terms of the provision of physical capital in the form of infrastructure such as bridges, highways and dams. However there are other less obvious types of physical projects that augment environmental capital stocks and involve activities such as land reclamation, pollution control, fishery management and provision of parks. Other types of projects are those that involve investment in forms of human capital, such as health, education, and skills, and social capital through drug-use and crime prevention, and the reduction of unemployment. There are few, if any, activities of government that are not amenable to appraisal and evaluation by means of social benefit-cost analysis.

Investment involves diverting scarce resources – land, labour and capital – from the production of goods for current consumption to the production of capital goods which will contribute to increasing the flow of consumption goods available in the future. An investment project is a particular allocation of scarce resources in the present which will result in a flow of output in the future: for example, land, labour and capital could be allocated to the construction of a dam which will result in increased electricity output in the future (in reality there are likely to be additional outputs such as irrigation water, recreational opportunities and flood control but we will assume these away for the purposes of the example). The cost of the project is measured as an opportunity cost – the value of the goods and services which would have

been produced by the land, labour and capital inputs had they not been used to construct the dam. The benefit of the project is measured as the value of the extra electricity produced by the dam. Chapter 2 discusses the concept of investment and investment appraisal in more detail.

The role of the benefit-cost analyst is to provide information to the decision-maker – the official who will appraise or evaluate the project. We use the word "appraise" in a prospective sense, referring to the process of actually deciding whether resources are to be allocated to the project or not. We use the word "evaluate" in a retrospective sense, referring to the process of reviewing the performance of a project or program. Since social benefit-cost analysis is mainly concerned with projects undertaken by the public sector the decision-maker will usually be a senior public servant acting under the direction of a minister. It is important to understand that benefit-cost analysis is intended to inform the existing decision-making process, not to supplant it. The role of the analyst is to supply relevant information about the level and distribution of benefits and costs to the decision-maker, and potentially to contribute to informed public opinion and debate. The decision-maker will take the results of the analysis, together with other information, into account in coming to a decision. The analyst's role is to provide an objective appraisal or evaluation, and not to adopt an advocacy position either for or against the project.

An investment project makes a difference and the role of benefit-cost analysis is to measure that difference. Two as yet hypothetical states of the world are to be compared – the world **with** the project and the world **without** the project. The decision-maker can be thought of as standing at a node in a decision tree as illustrated in Figure 1.1. There are two alternatives: undertake the project or don't undertake the project (in reality there are many options, including a number of variants of the project in question, but for the purposes of the example we will assume that there are only two).

The world without the project is not the same as the world before the project; for example, in the absence of a road-building project traffic flows may continue to grow and delays to lengthen, so that the total cost of travel time without the project exceeds the cost before the project. The time saving attributable to the project is the difference between travel time with and without the project, which is larger than the difference between travel time before and after the project.

Which is the better path to choose? The **with-and-without** approach is at the heart of the benefit-cost process and also underlies the important concept of **opportunity cost**. **Without** the project – for example, the dam referred to above – the scarce land, labour and capital would have had alternative uses. For example, they could have been combined to increase the output of food for current consumption. The value of that food, assuming that food production is the best (highest valued) alternative use of the scarce resources, is the opportunity cost of the dam. This concept of opportunity cost is what we mean by "cost" in social benefit-cost analysis. **With** the dam project we give up the opportunity to produce additional food in the present, but when the dam is complete it will result in an increase in the amount of electricity which can be produced in the future. The benefit of the project is the value of this increase in the future supply of electricity over and above what it would have been in the absence of the project. The role of the benefit-cost analyst is to inform the

decision-maker: if the **with** path is chosen additional electricity valued by consumers at $X will be available; if the **without** path is chosen extra food valued at $Y will be available. If X>Y the benefits exceed the costs, or, equivalently, the benefit/cost ratio exceeds unity. This creates a presumption in favour of the project although the decision-maker also has to take distributional effects into account – who would receive the benefits and who would bear the costs?

How do we measure the benefit of the additional electricity produced by the project? The gross value of the project output is measured by the amount consumers are willing to pay for it. In the case of a small increase in output **willingness-to-pay** (WTP) is measured by market price. However where the project output is substantial, relative to the original quantity of the good produced and consumed, willingness-to-pay for additional units of the good will be lower than market price because of the downward slope of the demand curve. In these circumstances marginal willingness-to-pay (WTP for an additional unit of output) declines as a result of the project and consumer benefits are measured as an area under the demand curve known as **consumer surplus**.

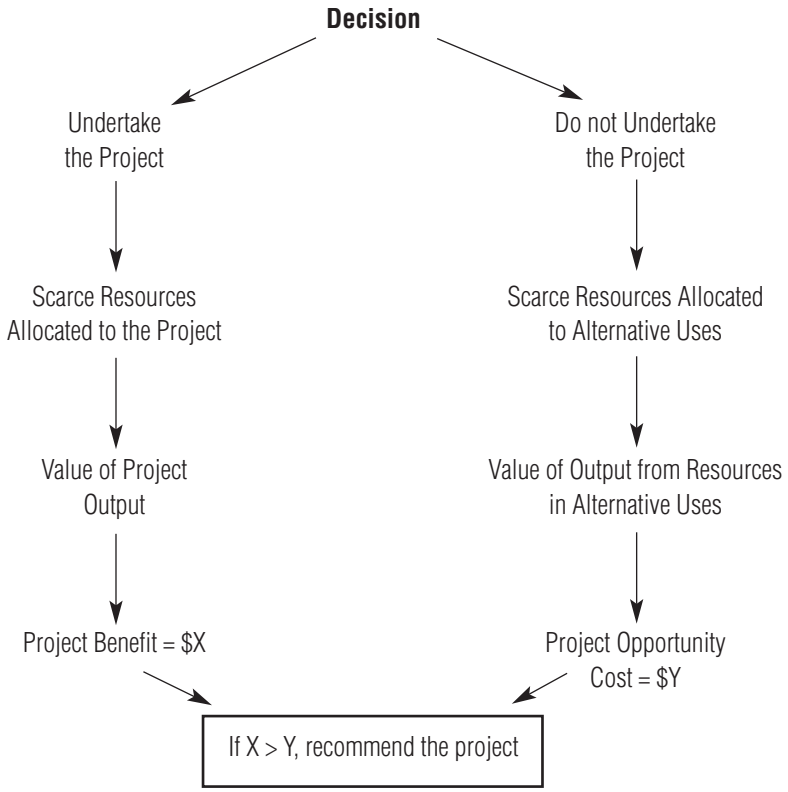

Figure 1.1 The "With and Without" Approach to Benefit-Cost Analysis

The concepts of consumer surplus, and the corresponding measure of **producer surplus**, which measures benefits or costs to suppliers, are discussed in detail in Chapter 7, in particular the concept of consumer surplus is illustrated by Figure 7.1. We now explain why we defer discussion of these important economic concepts until later in the book.

Traditional expositions of benefit-cost analysis usually start with the notion of consumer and producer surplus. However these concepts are relevant to the analysis only if output or input prices (or, in the case of non-marketed output, imputed prices) change as a result of undertaking the project. In many cases, including the case study which is developed in the Appendices to Chapters 4–6, no price changes can be identified. However all social benefit-cost analyses face the difficult task of social accounting – working out how the overall net benefits (or net costs) of the proposed project will be shared among the interested parties, including foreign and domestic, public and private, and consumers and producers. Entitlement to shares in net benefits is governed by a complex array of fiscal, regulatory and financial arrangements. Failure to understand these relationships can lead to fundamental errors of omission and double-counting in social benefit-cost analysis, and we have given these matters priority in the order of presentation.

The example of the electricity project was presented as if the benefit-cost analysis directly compares the value of extra electricity with the value of the forgone food. In fact the comparison is made indirectly. Suppose that the cost of the land, labour and capital to be used to build the dam is $Y. We assume that these factors of production could have produced output (not necessarily food) valued at $Y in some alternative and unspecified uses. We will consider the basis of this assumption in detail in Chapter 5, but for the moment it is sufficient to say that in a competitive and undistorted market the value of additional inputs will be bid up to the level of the value of the additional output they can produce. The net benefit of the dam is given by $(X–Y) and this represents the extent to which building a dam is a better (X–Y>0) or worse (X–Y<0) use of the land, labour and capital than the alternative use.

When we say that $(X–Y)>0 indicates a better use of the inputs than the best alternative use we are applying a measure of economic welfare change known as the **Kaldor–Hicks criterion**. The K–H criterion says that, even if some members of society are made worse off as a result of undertaking a project, the project confers a net benefit if the gainers from the project could compensate the losers. In other words, a project does not have to constitute what is termed a **Pareto improvement** (a situation in which at least some people are better off and no one is worse off as a result of undertaking the project) to add to economic welfare, but merely a **potential Pareto improvement**. The logic behind this view is that if society believed that the distributional consequences of undertaking the project were undesirable, the costs and benefits could be redistributed by means of transfer payments of some kind. The problem with this view is that transfers are normally accomplished by means of taxes or charges which distort economic behaviour and impose costs on the economy. The decision-maker may conclude that these costs are too high to warrant an attempt to redistribute benefits and costs. We return to the issue of the distributional effects of projects in Chapter 11.

Since building a dam involves costs in the present and benefits in the future the net benefit stream will be negative for a period of time and then positive, as illustrated in Figure 1.2. To produce a summary measure of the net benefits of the project all values have to be

converted to values at a common point in time, usually the present. The net present value is the measure of the extent to which the dam is a better (NPV>0) or worse (NPV<0) use of scarce resources than the best alternative. Converting net benefit streams, measured as net cash flows, to present values is the subject of Chapters 2 and 3.

When we compute present values for use in a social benefit-cost analysis we need to make a decision about the appropriate rate of discount. The discount rate tells us the rate at which we are willing to give up consumption in the present in exchange for additional consumption in the future. A riskless market rate of interest, such as the government bond rate, provides a measure of the marginal rate of time preference of those individuals participating in the market. However it can be argued that future generations, who will potentially be affected by the project, are not represented in today's markets.

In other words, in using a market rate of interest as the discount rate, the current generation is making decisions about the distribution of consumption flows over time without necessarily consulting the interests of future generations. This raises the question of whether a social discount rate, as opposed to a market rate, should be used to calculate the net present values used in public decision-making. This issue is considered further in Chapters 10 and 11.

Much of what has been said to this point also applies to projects being considered by a private firm: funds that are allocated for one purpose cannot also be used for another purpose, and hence have an opportunity cost. Firms routinely undertake investment analyses using techniques similar to those of social benefit-cost analysis. Indeed the appraisal of a proposed

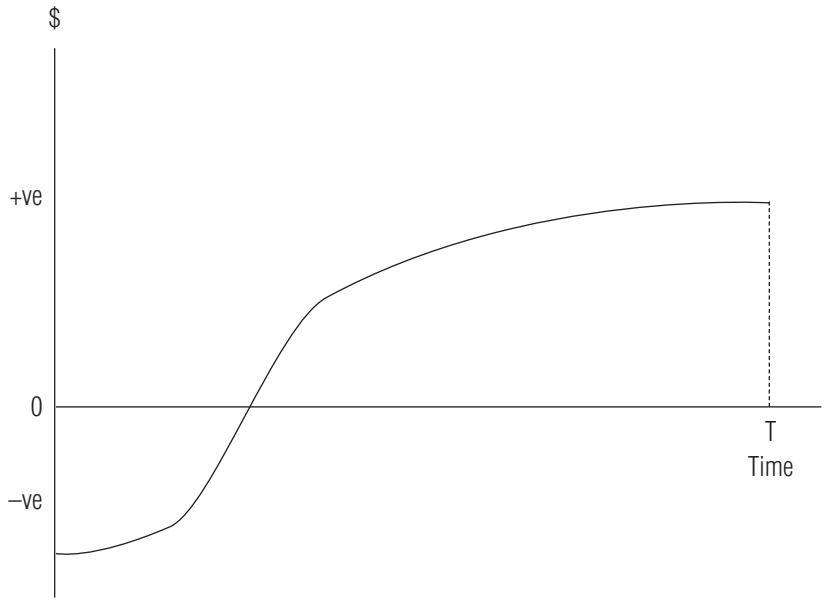

Figure 1.2 Typical Time-Stream of Project Net Benefits

project from a private viewpoint is often an integral part of a social benefit-cost analysis, and for this reason the whole of Chapter 4 is devoted to this topic. A private investment appraisal takes account only of the benefits and costs of the project to the private firm – its effect on revenues and costs and hence on profit. The project may have wider implications – environmental and employment effects, for example – but if these do not affect the firm's profits – its "bottom line" – they are omitted from the analysis. In contrast a social benefit-cost analysis takes a wider or "social" perspective – it measures and compares the costs and benefits experienced by all members of "society". In the context of social benefit-cost analysis "society" is to be interpreted in a relatively narrow way: it is simply that group of individuals deemed by the decision-maker to be relevant, and it is usually termed the **referent group**. Before undertaking a social benefit-cost analysis the analyst needs to ascertain from the decision-maker the composition of the referent group. Often the referent group consists of all the residents of a country, but it may be more narrowly defined in terms of sub-groups such as residents of a State or region, or social groupings such as the poor, unemployed, elderly, or people of Aboriginal descent.

It is clear that benefit-cost analysis can be conducted from different viewpoints: for example, it can take account of only the benefits and costs to the owners of the equity (the shareholders) in a private firm; an analysis from this perspective shall be referred to in this book as a **private benefit-cost analysis**. Alternatively, it can be broadened to include all benefits and costs to members of the referent group.

In what we term a **project benefit-cost analysis** estimates of all project benefits and costs are calculated at market prices; the project analysis tells us whether, in the absence of loans and taxes, the project has a positive NPV at market prices. The project NPV calculated in this way is neither the private NPV (the value of the project to private equity holders) nor the social NPV (the value of the project to the referent group). The equity holders do not stand to receive all the benefits of the project or incur all of the costs: for example, taxes may be due on project income, and loans may be obtained to finance part of the project, with consequent outflows in the form of interest payments. Whether the return to equity and debt holders is relevant in a social benefit-cost analysis depends on whether these groups are part of the referent group, but tax revenues paid by the project to the domestic government are certainly social benefits. Furthermore, by pricing inputs and outputs at market prices the project benefit-cost analysis ignores various types of referent group effects such as employment benefits, measured as the project wage bill less the opportunity cost of supplying the labour. It also excludes the benefits or costs of non-marketed commodities such as pollution. We discuss the former type of benefits and costs in detail in Chapter 5, and the latter in Chapter 12.

The important concept of the referent group is illustrated in Figure 1.3, which deals with an example which will be developed in Chapters 4–6 of this book. Suppose that a wholly foreign-owned company proposes to set up a factory in a developing country. The government wishes to appraise the proposal from the point of view of residents of the host country – the referent group. The firm has two questions to consider. First, is the overall project efficient from a market viewpoint? This is determined by the project benefit-cost analysis which compares the benefits and costs associated with undertaking the project, where benefits and costs are calculated at market prices; the present value of the net benefits is represented by

Area A+B in Figure 1.3 (the interpretation of the breakdown of the project net present value into the components A and B will be explained shortly). Second, is the project profitable from the perspective of the firm's owners, or, equity holders? This is determined by the private benefit-cost analysis. If the project is to be wholly internally financed the answer to this second question is obtained by deducting tax payments from the project NPV. However we will assume that there is to be some debt participation in the project in the form of a loan from a financial institution in the host country. The amount of the loan must be deducted from the project cost and the loan repayments and interest charges deducted from the project's after-tax benefits to give the benefits and costs of the project to the equity holders; the private benefit-cost analysis.

In this example we shall assume that the firm's equity holders are not considered part of the referent group. This being so, in Figure 1.3 Area A represents the net present value of the project net benefits to the members of the referent group: the lenders of the firm's loan (the bank) and the recipients of the firm's tax payments (the government). The net benefit of the project to the non-referent group members, the firm's equity holders, expressed as a net present value, is represented by Area B. Only if the net benefit to equity holders is positive is the project worthwhile from the firm's viewpoint. Areas A and B together amount to the project NPV.

As noted above the project may have a wider impact than that summarized by the project benefit-cost analysis. The project may generate benefits or costs to various groups within the host country. For example, some people who would otherwise have been unemployed may obtain jobs: the pay that they receive from the firm may be higher than the value

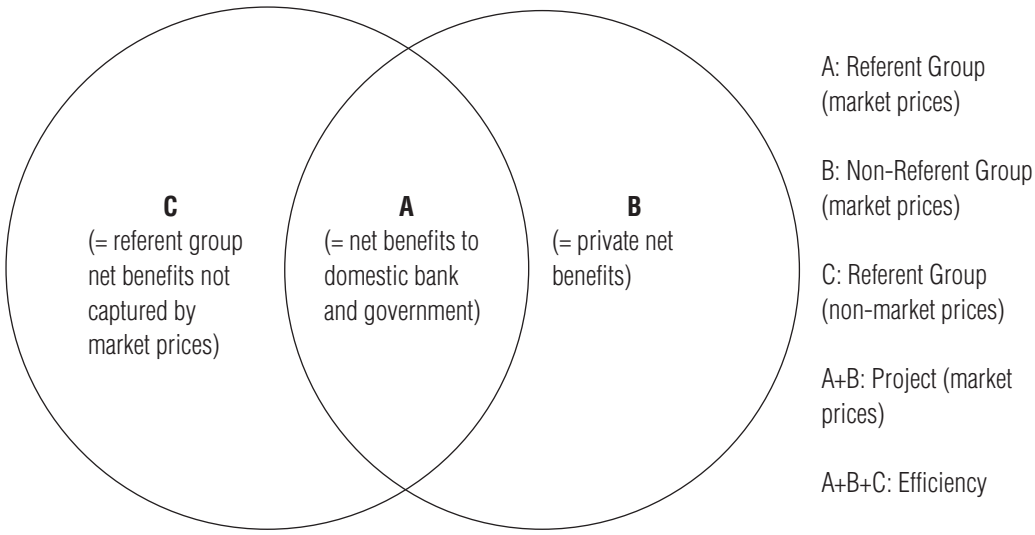

Figure 1.3 Relationship between the Project, Private, Efficiency and Referent Group Net Benefits

of their time in some non-market activity, thereby resulting in a net benefit to them. The firm may purchase various goods and services, such as water and electricity, from government agencies, paying prices in excess of the production costs of these inputs, again generating net benefits for this section of the referent group. The project may generate pollution which imposes health and other costs on residents of the host country. In Figure 1.3 Area C represents the set of net benefits (present value of benefits net of costs) accruing to the referent group as a result of divergences of market prices from referent group valuations of benefits or costs, or as a result of non-marketed benefits and costs. We shall refer to these as **non-marketed net benefits/costs** accruing to members of the referent group. The total referent group net benefit is given by Area A+C.

What then does the whole area, A+B+C, represent? This can be thought of as representing the **efficiency net benefits** of the project – the present value of benefits net of their opportunity cost, and irrespective of whether they accrue to members of the referent group or not. Area B represents the net benefits to the non-referent group equity holders, which will determine the firm's decision whether to undertake the project or not. Area A+C represents the net benefits to the referent group, which will determine the government's decision as to whether or not to allow the project to proceed. Referent group net benefits are a subset of the efficiency net benefits. The composition of the referent group follows from the definition of the scope of the benefits and costs to be counted. As noted earlier, it is essentially a policy decision as to who the relevant stakeholders or referent group members are. The composition of the referent group net benefit is the main issue which the benefit-cost analyst is called upon to address, although in negotiating with the firm the decision-maker may also be interested to know how attractive the project is from a private viewpoint. It should also be noted that the definition of the referent group can be controversial especially in situations where there are transboundary externalities such as pollution affecting citizens of other states or countries.

Apart from measuring the aggregate referent group net benefit, the analyst will also need to know how this is distributed among the different sub-groups as the decision-makers will, most probably, want to take into consideration the distribution of net gains and losses among the referent group members: this is referred to as the **referent group analysis**.

In summary, the hypothetical project discussed above (or any other project) can be appraised from four different points of view:

(i) the **project** benefit-cost analysis: this is represented by Area A+B and is obtained by valuing all project inputs and outputs at private market prices;
(ii) the **private** benefit-cost analysis: this is obtained by netting out tax and interest and debt flows from the project appraisal, and, if the firm's equity holders are not part of the referent group as in our example illustrated in Figure 1.3, it will be given by area B: which, in this example, is the non-referent group project net benefit;
(iii) the **efficiency** benefit-cost analysis: this is represented by Area A+B+C and is obtained in a similar way to the project appraisal, except that the prices used to value inputs or outputs are shadow- or accounting-prices, which are discussed in Chapter 5, or are derived from the application of non-market valuation techniques as discussed in Chapter 12;
(iv) the **referent group** (or social) benefit-cost analysis: this is represented by Area A+C and can be obtained in two ways as noted below – directly, by enumerating the costs and

benefits experienced by all members of the referent group; or indirectly, by subtracting non-referent group net benefits from the net benefits calculated by the efficiency analysis. In our example, the non-referent group net benefits are summarized by the private appraisal (Area B), although in other cases the private project owners may be part of the referent group.

In the course of undertaking a complete social benefit-cost analysis the project analyst will therefore need to follow a sequence of steps:

- First, calculate the **project** cash flow at market prices (Area A+B in Fig. 1.3)
- Second, calculate the **private** cash flow at market prices (Area B in Fig. 1.3)
- Third, recalculate the **project** cash flow at **efficiency** prices (Area A+B+C)
- Fourth, disaggregate the efficiency cash flow among the **referent group** (and non-referent group) members.

It is clear that there are two ways of going about the task of estimating Area A+C – the net benefits to the referent group: directly, by listing all the benefits and costs to all members of the referent group – in this example, labour, government organizations, and the general public – and measuring and aggregating them; or indirectly by measuring the efficiency net benefits of the project and subtracting from them the net benefits which do not accrue to the referent group. Under the first approach Area A+C is measured directly; under the second approach Area A+B+C is measured and the net benefits to those not in the referent group (represented in the example by Area B) are subtracted to give Area A+C.

At first sight it might seem strange to consider using the indirect approach. However as we will see in Chapters 4 and 5 it is relatively easy to measure the net benefits represented by areas A+B+C and B respectively. The net efficiency benefits of the project are obtained by valuing all project inputs and outputs at marginal values to the world economy: these marginal values may be represented by accounting- or shadow-prices which are artificial rather than observed market prices, and which are relatively easy to calculate, as discussed in Chapter 5, or by prices obtained from the application of non-market valuation techniques as discussed in Chapter 12. The net private benefits are obtained by using market prices which are directly observable, and deducting tax and debt flows: this calculation simply mimics the process which the firm undertakes internally to decide whether or not to proceed with the project. Measuring Area A+C directly is more difficult because each subset of the referent group which is affected by the project has to be identified and their costs and benefits measured. In summary, the indirect approach produces an aggregate measure, whereas under the direct approach the social net benefits are measured in disaggregated form and assigned to various groups. While the disaggregation provides important information which relates to the income distributional concerns of the decision-maker it is more difficult to obtain than the summary figure.

In this book we advocate the use of both approaches: measure Area A+C as A+B+C less B, and then measure its component parts directly and sum them to get Area A+C. If the same answer is not obtained in both cases an error has been made – some benefits or costs to members of the referent group have been omitted or incorrectly measured. A check of this nature on the internal consistency of the analysis is invaluable.

An analogy which may assist in determining what is to be measured and where it belongs in the analysis is to think of the project as a bucket. Costs go into the bucket and benefits

come out: however the range of benefits or costs which go in or come out depends on the perspective that is taken. In the **efficiency analysis** we count all the costs and benefits measured at the appropriate shadow-prices, and the latter minus the former – the net benefits of the project – is equivalent to Area A+B+C in Figure 1.3. The **project analysis** is similar to the efficiency analysis except that all the costs and benefits are measured at market prices, where these exist, to obtain an estimate of Area A+B, and non-marketed benefits and costs are ignored. In the **private analysis**, which, in the example, measures non-referent group net benefits accruing to the foreign firm, we count all sums contributed or received by the firm's equity holders to calculate the non-referent group net benefits (Area B). In the example this consists of the project cost less the loan obtained from the domestic bank, and the project revenues less the interest and principal repayments, and the tax payments to the host country. In the **referent group** analysis we count all contributions and receipts by referent group members to estimate Area A+C: in our example this consists of Area A – the capital contribution of the domestic financial institution, together with the loan repayments and interest payments, and the taxes received by government, and Area C – the employment benefits received by domestic labour. (If there are other, non-marketed net benefits or costs accruing to *non-referent group* members, we would need to add this as an additional category, say D, which would be included in the efficiency benefit-cost analysis, but, like area B, deducted from the aggregate efficiency net benefit to arrive at the total referent group net benefit. This is considered in more detail in Chapter 6.)

At this point it should be stressed that many projects will not correspond exactly to the above example, and what is to be included in Areas A, B and C will vary from case to case. Furthermore, additional categories of project effects may be required. For example, suppose, as in the case study developed in the appendices to Chapters 4–6, that part of the cost of the project was met by a loan from a foreign bank which is not part of the referent group. To incorporate this possibility we would add the foreign bank's net benefits to Area B in Figure 1.3. Area B still forms part of the efficiency benefit-cost analysis but this area has to be subtracted from the total efficiency NPV to calculate the referent group NPV. Another possibility is that instead of paying taxes the foreign firm receives a subsidy. Areas A and B would then need to take account of the subsidy: as a credit item in Area B and a debit item in Area A. A comprehensive framework which takes account of all possible categories of benefits and costs is presented in Chapter 6.

We have discussed four ways of looking at a benefit-cost analysis, and the interrelationships among these four points of view. The client for whom we are conducting the analysis will rarely attach equal importance to each. For example, an international organization such as the IMF or World Bank may be primarily interested in the best use of scarce resources, as summarized by the efficiency analysis. A regional organization, such as a regional development bank, will also be interested in efficiency but will want to know how the host country – the referent group – is affected. The host country will be primarily concerned with referent group effects but, where a private firm is involved, may also be concerned with the viability of the project from a private viewpoint for negotiation purposes. The private firm will similarly be interested in the referent group analysis as it tries to obtain favourable terms from the host country.