Introduction to Cost-Benefit Analysis: Looking for Reasonable Shortcuts

Book · January 2010

CITATIONS
49

READS
14,944

1 author:

Ginés de Rus
Universidad de Las Palmas de Gran Canaria
106 PUBLICATIONS **1,610** CITATIONS

SEE PROFILE

Chapter 1

Introduction

> Any project can be viewed as a perturbation of the economy from what it would have been had some other project been undertaken instead. To determine whether the project should be undertaken, we first need to look at the levels of consumption of all commodities by all the individuals at all dates under the two different situations. If all individuals are better off with the project than without it, then clearly it should be adopted (if we adopt an individualistic social welfare function). If all the individuals are worse off, then clearly it should not be adopted. If some individuals are better off and others are worse off, whether we should adopt it or not depends critically on how we weight the gains and losses of different individuals. Although this is obviously the 'correct' procedure to follow in evaluating projects, it is not a practical one; the problem of benefit-cost analysis is simply whether we can find reasonable shortcuts.
>
> Joseph E. Stiglitz (1982)

1.1. The rationale of cost–benefit analysis

Cost–benefit analysis is not about money. It is not about inputs or outputs either. It is about welfare. Money is central to financial analysis but only instrumental in the economic appraisal of projects and policies. Money is the common unit in which economists express the social costs and benefits of projects. Volume of drinking water, accidents avoided, time savings and energy and labour consumed are measured in different units and we need a common unit of measure to express all these heterogeneous items in a homogeneous flow. This is the key role of money in cost–benefit analysis.

The creation of jobs is frequently presented as a benefit of a project, but labour is an

input not an output. A motorway is not constructed to create jobs but to move people and goods. Workers building and maintaining a motorway represent a social cost equal to the net value lost in the next best use of this input. It is true that if a worker is unemployed, society does not lose as much as in the case of a similar worker already employed, but this only shows that cost values are context dependent.

The output of a project is easier to measure than its welfare effects. Public agencies report their activities with indicators such as passengers, water, electricity or the number of students taught within a training programme, but cost–benefit analysis sees output as a means to increase welfare. The success of a new facility cannot be explained by the number of users. It is possible to subsidize prices to induce people to use the new facility without increasing social welfare. Therefore, cost–benefit analysis is interested in the social value achieved from the outputs of the project to compare with the value of other goods sacrificed elsewhere for the sake of the project.

Cost–benefit analysis is about the well-being of individuals affected by the project and not about the number of trips or visits. The change in welfare is what economists want to measure, and this is quite a challenging task because welfare cannot be measured. To solve this problem, economists have found an alternative: to use money as an expression of welfare. I do not know how great the utility[1] of a particular individual is when driving his car from *A* to *B* at a particular date and time, but if I am able to determine the amount of

money to charge for this trip that makes him indifferent between driving or not, then interesting things can be said. Cost–benefit analysis is not about money but money helps.

Cost–benefit analysis conceived as a toolkit for the selection of projects and policies, in the general interest of the society, presupposes the existence of a social planner, a benevolent government that compares benefits and costs before the implementation of projects and policies. Many economists and non-economists would consider such a view as naïve, to say the least. An alternative view[2] explains a government's action by the political power of different interest groups. Subsidies to agriculture, for example, could be better explained by the pressure of farmers than by an independent assessment of the social benefits and costs of agricultural policy.

Do we need to believe in the goodwill of the government to practise cost–benefit analysis? The answer is no. If we believe that a government's acts are better explained by the influence of interest groups, cost–benefit analysis can show who benefits and who loses as a result of particular projects, and the magnitudes of the gains and losses. This assessment can be very helpful in explaining which policies are adopted or even in influencing a government's decision. 'Cost–benefit analysis may be in the battle against misleading information spread by self-interested political pressure groups. Still, these analysts can influence political outcomes by making enough voters aware of the true effects of different policies' (Becker, 2001).

To present the conceptual foundations and methods of cost–benefit analysis we will proceed 'as if' the government would aim for the best projects in the general interest of the society. Although we know of many cases that show that such an assumption is unrealistic, the simplification is harmless. As we proceed to identify benefits and costs, winners and losers and try to measure and value the main effects of the project under evaluation, the analysis is not going to change whatever our particular beliefs on the government's behaviour are.[3]

We have started assuming the existence of a benevolent government. This is not the only assumption and simplification in this book; in fact, there is no way to deal with the analysis of the economy but through the use of simplifying assumptions, replacing the actual world with a model that reflects the essence of the more complex reality that we want to understand.

To move forward, we need to clarify what is understood by acting in the general interest of the society. Let us consider that our benevolent government is evaluating the construction of a dam and a hydroelectric power station. The government doubts whether it should accept the project. By undertaking the project, the region would obtain electricity at a lower cost than without the project, recreation benefits, both in the stock of water (e.g. fishing and boating) and on the flow of reservoir release (e.g. rafting), and some jobs would be created at the time of its construction and during the lifetime of the project. Furthermore, there might be a multiplier effect, as the project would create new economic activity induced by the expenditure associated with the construction and operation of the project.

Economists point out that from the benefits described above we have to deduct some costs. Firstly, the construction and maintenance, equal to the net benefits of alternative needs that have not been attended to because the public money has been assigned to the dam and power station, have to be deducted. They also argue that labour is an input, not an output, so it is a cost of the project, though its magnitude will depend on what is lost when the worker is employed within the project. The multiplier effect, if it exists, turns out to be irrelevant if it is also associated with the alternatives.

Secondly, all the other costs associated with the relocation of the inhabitants of the village in the area where the dam would be built and with the people negatively affected by the alteration of the flow and course of the river should be deducted. The magnitude of these costs could be substantial.

The government considers all the relevant benefits and costs regardless, in principle, of who the beneficiaries and the losers are (assume for simplicity that all the effects are inside the country) and the government decides to undertake the project if, given the available information, the society improves. Its decision is not based on the arguments of the private companies that will build the dam and power station, nor on the campaign of the opponents. The decision takes into account the whole society, with social welfare as the unique reference. The challenge for our benevolent government is how to value all the benefits and costs and how to compare them given that beneficiaries and losers are individuals with different income, education, health, etc, and are affected at different

moments during the lifespan of the project.

This water project, as any other public infrastructure such as parks, high speed rail, highways, ports or the introduction of policies such as environmental regulations, can be interpreted as perturbations in the economy affecting the welfare of different individuals at different moments in time compared with the situation without the project or policy, which does not necessarily mean the *status quo* but what would have happened in the absence of the project or policy.

The assessment of the effects of the project requires a benchmark. It is necessary to compare the world *with* and *without* the project: to recreate an alternative world, or the so-called counterfactual. Cost–benefit analysis practitioners have to solve two main problems. Firstly, they have to build the counterfactual and this means to replicate the world without the project, a dynamic world that evolves without the perturbation introduced by the project. This is not an easy task because the time period for this exercise may be quite long, 40 years or more, and the values of key variables will possibly change in each one of these years, only some of them in predictable ways. Secondly, the practitioner has to imagine the world with the project, forecasting the main changes with respect to the counterfactual that he has previously created.

The expected changes when the project is implemented are then the result of the comparison with the counterfactual: the worse the counterfactual, the better the project. Hence, it is important to present all the assumptions and the data used to complete this

exercise. Transparency and *ex post* evaluation can help to avoid both innocent errors and strategic misrepresentation.

Suppose the counterfactual and the world with the project have been properly designed and the expected changes have been estimated: time savings, enhanced water quality or a reduction in the number of fatal accidents. Now, the analyst has to convert these values into monetary units ($)[4] assuming that this is technically possible and morally acceptable.

We want to measure changes in the welfare of the individuals who compose the society; however individuals' utility cannot be measured in the same way as the amount of electricity produced or the number of people displaced to build the dam. To decide on the goodness of the project we need to measure something that is unobservable. Furthermore what it is observable – the production of electricity, number of individuals involved, extension of flooded surface, etc – is not very useful if we do not translate the physical units into a common measure related to changes in individual utility, which allows the comparison between what is gained and what is lost.

Hence, though the ideal way of measuring the impact of our project is through utility functions (we would measure the change in utility of each individual) the problem is that these utility functions and the associated utility changes are unobservable. Converting the unobservable utility changes through an 'exchange rate' between utility and income to observable monetary units gives us a way of calculating the impact of the project.

One alternative solution might be to submit the project to a referendum and to accept the outcome: that is, accepting the view of the majority. Let us have a look at this in more detail. Table 1.1 collects the information (expressed in monetary units) of the benefits and costs of those affected by the construction of the dam and the hydroelectric power station. Our society consists of five individuals.[5]

Table 1.1.
Benefits and costs of a hydroelectric power station
(values in $)

Individual	Benefits	Costs	Net Benefits
A	7	0	7
B	2	8	–6
C	3	4	–1
D	9	1	8
E	1	6	–5

The individual *B*, for example, benefits from cheaper energy (+$2) but he also fishes downstream and the dam prevents him from practising his favourite sport in the initial conditions (–$8). The result is a net loss of $6 for individual *B*. We could interpret the values in the column 'net benefits' as the monetary compensation that will be needed with the project to leave each individual indifferent without the project: for example, the individual *B* would be willing to accept $6 and the individual *A* would be willing to pay $7.[6]

The column 'net benefits' allows us to anticipate that the project would be rejected in a referendum. Individuals *A* and *D* would vote in favour, but individuals *B*, *C* and *E* would vote against it. Would it be a good decision to reject the construction of the water project? To answer this question we have to check whether the construction of the dam is a social improvement and for this purpose we need to define a decision criterion.

A possible criterion is the strong Pareto improvement. To move from one situation to another is a social improvement (in the sense of Pareto) if at least one person is better off without making anyone else worse off. There are winners and nobody loses. We have seen that the referendum would result in the rejection of the project. Would it be possible, in these circumstances, a Pareto improvement despite the outcome of the ballot?

Although it seems clear that the project under discussion would not be approved in a referendum, the society may gain from the project if, as it happens to be in this case, the benefits ($22) outweigh the costs ($19). Suppose the project is carried out and part of the benefits is used to compensate individuals *B*, *C* and *E*, so that their net benefit is zero, leaving them indifferent. Table 1.1 shows that, after compensation, there is a net benefit of $3 to share out as deemed appropriate. If the project is rejected this net gain would be lost.

On the other hand, when comparing the benefits and costs of the project, the magnitude of gains and losses counts. Individual *C* is against the project because it costs him $1, while *D* stands for the project because he gains benefits of $8. If we ignore the

intensity of preferences, like in a referendum, we lose the potential gains arising from the project.

As we have seen, the Pareto improvement criterion requires no losers (i.e. there is full compensation to those initially harmed by the project). This rarely happens in the real world since, in many cases, the situation is similar to that described, but without full compensation to the losers.[7] If a project produces a positive balance of benefits to the society as a whole and there are losers who, for some reason, can not be fully compensated, it is normal practice to undertake such a project (the winners could have compensated the losers and still remain winners).

This criterion, in which the compensation is only hypothetical, is known as the potential compensation criterion, or Kaldor–Hicks criterion.[8] If the losers are compensated, it would result in a Pareto improvement. Unless the project has unacceptable distributional consequences, the economic evaluation of projects and policies rests basically on the criterion of potential compensation just described.

To be more precise, we need – at least for a small project – to weigh individual (or group) *i*'s monetary gain with the marginal utility of income of individual *i*, and with the social welfare weight attributed to individual *i* reflecting the social welfare function. Hence, the marginal social utility of income (see Chapters 2 and 11) attributed to *i* depends on what social welfare function[9] we assume and the income distribution.

We can multiply the social marginal utility of income by the monetary valuation of the project (willingness to pay or willingness to accept) of individual *i* and sum over all individuals. So if initial welfare distribution is optimal (and where the marginal utility of income might vary across individuals since they might have very different utility functions: one being a dedicated wild life person, another being a passionate consumer of diamonds) the marginal social utility of income is the same for everyone. Then a project so small that it leaves the welfare distribution unchanged as a linear approximation can be evaluated simply by summing willingness to pay or willingness to accept across individuals/groups.

The common justification of the Kaldor-Hicks compensation criterion in practice is based on the argument that redistribution can be performed more efficiently through the fiscal system and that, overall, given the large amount of different projects being carried out, the positive and negative distributional effects tend to offset each other, and everybody wins in the long run, or are unimportant or the costs of identifying winners and losers and paying compensation are higher than the benefits.

1.2. Steps of cost–benefit analysis and overview of the book

The economic appraisal of investment projects and public policies must be flexible enough to capture the specific characteristics of each case study; however there exist some steps that have to be followed independently of the particular aspects of the project under evaluation. They are the following:

(i) *Objective of the project and examination of the relevant alternatives*

Before evaluating, the objective of the project, the problem to be solved, has to be clearly defined and the relevant alternatives identified. To analyse an isolated project without considering its role within the programme or policy where it belongs can lead to wrong conclusions. Moreover, before working with data and applying the methodology of economic evaluation, it is essential to analyse the relevant alternatives that allow the achievement of the same objective. An improper analysis of available alternatives can lead to important errors despite the methods and techniques being rigorous.

There are two *a priori* approaches for the practitioners in the appraisal of a project: firstly, when the analyst has to evaluate a particular project, for example, a price reduction in a public service; secondly, when the project is the improvement of a public service. If the goal of the regulator were to benefit consumers without damaging service quality, a possible measure could be to reduce prices, keeping the financial equilibrium with public subsidies. However there are also other policies to achieve this goal. An alternative could consist of introducing a system of incentives that compensates for efforts in the reduction of costs that allow price cuts. Another policy could be a private concession of the public service.

The consideration of different projects to achieve the same goal is a previous stage to the identification and quantification of benefits and costs in the evaluation given that the

omission of more efficient alternatives is to lose the opportunity to gain better results. It is not enough to have positive social benefits, it is required that those benefits are greater than the benefits in the best available alternative. The same happens with investment projects. The question 'is the investment the best way of solving the problem?' must be answered. Other possible reversible and less costly options must be analysed, such as different management of the facility.

In the stage of the search for relevant alternatives it is very useful for the economist to interact with and receive feedback from other specialists more familiar with the technology or field related to the project. The objective of this step is to avoid errors because of a lack of precise information about more efficient methods to achieve the same goal. The greater refinement in the evaluation methodology would be useless if better alternatives had not been taken into consideration.

Finally, it is not convenient to define projects with too broad a scope because a positive evaluation of the aggregate can hide separable projects with negative expected returns. Therefore their inclusion, without differentiation, in a programme or a more global project can lead to wrong conclusions. To establish the limits of a project is not always easy but a careful discussion of the project with experts can allow us to distinguish intrinsic complexity from the inclusion of independent projects that are perfectly separable.

On the other hand it is nonsense to evaluate a project narrowly defined in the sense

that its existence is not possible without complementary actions. Suppose that an investment project is composed of two main parts (e.g. a port and an access road) and the social net present value (*NPV*) is negative. The strategy of promoters could be to evaluate only the first part (the construction of the port) and once it has been built to present another project consisting of the complementary infrastructure to connect the port with the road network. In this case, the road project will probably be socially worthy because the investment cost of the already existing port is now irrelevant in the evaluation of the construction of the road, principle that it is not applicable to the lost benefits if the port can not be operated.

In this chapter, we present the basic concepts of cost–benefit analysis and make explicit the simplification used to deal with the economic evaluation of projects that have medium and long term effects.

(ii) *Identification of costs and benefits*

Once the project is defined and bounded, it is necessary to identify the benefits and costs derived from its implementation. In some cases, this step is immediate and must not bring up greater difficulties, for example when projects have only direct effects (Chapter 2).

The identification of the costs and benefits of a project with significant indirect effects on other markets is more complicated because the impact must be located in markets different from the ones where direct effects are produced. The more reasonable

approximation, when the analysis is not conducted within a general equilibrium framework, consists of identifying the main secondary markets affected by the project as it happen to be the case with the evaluation of a new railway line that reduces the demand of an existing airport (Chapter 3).

In financial analysis, the identification is much simpler: benefits are revenues and costs are the payment of inputs valued at market prices. However, in economic analysis, benefits are those that are enjoyed by the individual independently of their conversion into revenues, and costs are net social benefits lost in the best available alternative.

Moreover, it is necessary to decide 'who stands' in cost–benefit analysis. Generally, country frontiers delimitate who must be included. Citizenship is the reference when the project has no global or controversial effects beyond the national boundaries. Sometimes it depends on who finances the project. In a co-financed project with supranational funds it would not be acceptable to exclude citizens of countries that contribute with their taxes to financing the project. On the contrary, it is not uncommon for a region, only to consider local benefits and costs, ignoring positive or negative effects that take place in the rest of the country.

(iii) *Measurement of costs and benefits*

Projects' benefits can be measured through individuals' willingness to pay (or willingness to accept). Sometimes a monetary measure of the utility change that is derived from the

project can be obtained observing the behavior of consumers in the market; i.e., from market data. This is the case of the measurement of direct benefits in the primary market affected by the project (Chapter 2), the indirect effects in the secondary markets related to the primary market (Chapter 3), and in the valuation of non-marketed goods when the analyst can find an 'ally' market where some useful information is revealed about the willingness to pay of the individuals (Chapters 5 and 6)

On other occasions, economists have to estimate project benefits by asking people directly about their willingness to pay (stated preferences); this consists of asking individuals about monetary quantities that reflect the change in their welfare thanks to the project. This approximation is used for non-marketed goods like environmental impacts or safety changes (Chapter 6).

In general, projects and public policies that are subject to evaluation imply the use or saving of resources. The costs of a standard investment project can be classified as: construction, maintenance, labour, equipment and energy – costs that are measured from the quantity of the inputs valued by their respective prices. From an economic point of view, the cost of input use is the net social benefit lost in the next best alternative. Market prices will sometimes be a good approximation of the opportunity cost, but in other cases it will be necessary to introduce some correction in market prices to approximate the social opportunity cost of the inputs, and this is what are called shadow prices (Chapter 4).

(iv) *Benefits and costs aggregation*

Benefits and costs occur in different periods of time and affect different individuals. The aggregation requires homogeneity, but benefits and costs that occur in successive years or affect individuals with different social conditions are not homogeneous. If they are directly summed, the implicit weight associated with each benefit or cost is the unity: a unit of benefit is identical disregarding the year or the individual.

Many infrastructure projects have lifetime periods over thirty years. Moreover, in the case of public policies that modify educational or health programmes, introduce or eliminate taxes, etc., the *ex ante* lifespan is practically infinite. To discount future benefits and costs is a process of homogenizing to allow comparison. The discounting is performed using a discount rate greater than zero. This implies that the value of the benefits and costs decreases with time. The basic idea consists of the fact that individuals generally give more value to present than to future consumption and, therefore, future units of consumption are counted with a lower present value (Chapter 7 and 8).

Project costs and benefits affect individuals' utility. To go from net individual benefits to aggregate social benefits implies redistributive effects. If the society gives, for example, more weight to the income of poor people, then the benefits and costs of a project cannot be added without social weighting. The net social benefit of the project should ideally be obtained as the weighted sum of the individual net benefits (Chapter 2 and 11).

(v) *Interpretation of results and decision criteria*

The aspiration of the practitioner of cost–benefit analysis is to obtain a figure that summarizes the flows of benefits and costs. This figure is the *NPV* of the project, and helps with the decision accept–reject or to choose between a set of projects.

To obtain a unique figure is not always easy. There are positive or negative impacts that resist the reduction to a monetary figure as happens to be the case with some environmental impacts. There are situations in which it can be appropriate to make a qualitative description of some effects, and then to attach this information to the NPV obtained with the effects that can be unambiguously measured.

Decision criteria based on the discounted value of benefits and costs are straightforward. If the NPV is positive and the redistributive effects of the project are positive or unimportant, the project increases social welfare. If we have to compare projects it may be necessary to homogenize before choosing the project with the higher NPV (Chapter 7 and 8).

Results must be subject to a risk analysis with the objective of determining the sensitivity of the NPV to changes in key variables. Ideally, it is preferable to compute a probability distribution of NPV instead of obtaining a unique NPV figure. Risk analysis allows the decision maker to have some information about the likelihood of the feasible results. Risk analysis does not eliminate the risk of the project but makes more evident the

actual risk of the project to the decision maker (Chapter 9).

(vi) *Comparison of the project with a base case*

In cost–benefit analysis, it is important to avoid a comparison of the project with an irrelevant *base case*. For example, comparing with the situation before the project, the NPV can be high because it can be hiding the fact that without the project that situation does not remain constant. There are maintenance policies or a minimum renewal of equipment, etc., that could be implemented without the project. Thus, in the evaluation of the construction of a high speed rail line, we have to compare with a counterfactual where the demand changes and the supply of conventional rail and competing modes also changes.[10]

The situation without the project is known as the base case. We have to distinguish between *do nothing* and *do minimum*, which consists of the minimum intervention that is foreseen without the project. The distinction can be made clearer with an example. Consider that the objective is to substitute the water pipes, that supply water to the city, due to excessive water leaks. In this project the reasonable base case is a *do minimum* because, without the project, there will be maintenance operations and selective actions to avoid greater damage. Now consider that the project consists of a maintenance plan instead of investing in a new network. In this case the base case would be a *do nothing*.

(vii) *Economic return and financial feasibility*

Cost–benefit analysis compares the social benefits and cost in contrast with financial analysis, which uses revenues instead of social benefits and private costs instead of social costs. However it is very important for the analyst to have a report that not only includes the economic return of the project but also the financial result or commercial feasibility of the project.

It is perfectly possible for a project or public policy to generate social benefits that exceed social costs and, at the same time, to present a negative financial result. Let us consider, for example, the case of a reforesting policy that reduces land erosion and delivers new space for recreation. Moreover, the responsible public agency obtains some revenues from charging for parking close to the recreation area. It is likely that this project presents a positive social NPV and a negative financial result. The analyst must present both results to the decision maker for two main reasons.

Firstly, the real world is characterized by the presence of budget constraints; therefore it is really useful for the public agency to have information on the social net benefit of the project as well as on the proportion of costs that are covered by revenues. Secondly, many projects produce a wide range of NPVs as a function, for example, of the pricing policy applied. It is usual for projects that admit the possibility of charging to users to present different possible combinations of social NPV and financial NPV. For example, a road project can be evaluated as a free access road or as a toll road. If the second option is

chosen, several possible price structures exist: it is possible to discriminate by time, vehicle type or use intensity. It is likely that social benefits diminish with the toll; however collected revenues can contribute to fixed costs. To report the different options available and their social and financial NPV increases the usefulness of cost–benefit analysis.

---

[1] We use 'utility' and 'individual welfare' as synonymous.

[2] George Stigler, Gary Becker and Sam Peltzman are among the top economists promoting the 'interest group competition model'.

[3] Nevertheless, once the conventional cost–benefit analysis is completed, we need to address explicitly the institutional design and the possible conflicting objectives, and we attend to that in Chapter 10.

[4] We use $ as representing an undefined monetary unit without any relation to its actual market value.

[5] We assume the individual is the best judge of his own interest; hence we ignore problems derived from distorted preferences (see Adler and Posner, 2001).

[6] We assume for simplicity here that willingness to pay and willingness to accept coincide. For a technical discussion on the why these monetary measures of utility changes differ, see Chapter 11.

[7] As it is the case with legal expropriation.

[8] There is a simplification here with the Kaldor-Hicks hypothetical compensation. Boadway pointed out that a positive sum of compensating variations, for example, is not equivalent to gainers being able to compensate losers. The problem known as the Boadway-paradox is that the act of compensation might affect relative prices and hence welfare/utility of an individual (see Boadway, 1974; Jones, 2002).

[9] For a Utilitarian society the marginal social utility or welfare weight is unity for all *i*, for a Rawlsian society it is equal to zero for all but the worst-off individual or group.

[10] The comparison with an irrelevant alternative can be an error of the analyst, or a strategy for getting the project through.