1

Introduction

Cost–benefit analysis (CBA) offers a formal approach for decision makers to delineate the costs and benefits of different policies.[1] Although the conceptual underpinnings of CBA may be traced to the nineteenth-century French civil engineer and economist Jules Dupuit, see Dupuit (1844), extensive application had to wait until the twentieth century. Propelled by a rising demand for electricity and substantial damage from several serious floods, the US Congress passed two significant flood control acts (referred to here as the 1936 Act and the 1944 Act, respectively). The 1936 Act called for "works of improvement" on more than fifty major rivers throughout the United States and made flood control a federal government activity. Importantly, the Act introduced an approach to prioritizing projects:

> The Federal Government should improve or participate in the improvement of navigable waters or their tributaries, including watersheds . . . for flood control if the benefits to whomsoever they may accrue are in excess of estimated costs.[2]

In 1981 President Reagan issued Executive Order 12291 that stated that "regulatory action shall not be undertaken unless the potential benefits to society for the regulation outweigh the potential costs to society." This is the essence of cost–benefit analysis: a project is recommended if the benefits (whomever they accrue to) exceed the costs. A critical issue is how to measure costs and benefits and what is meant by benefits exceeding costs. Pareto (1896–1897) argued that a project should be undertaken only if at least one individual is made better off while no one is made worse off. This is a nice criterion but it is seldom fulfilled by actual projects since they tend to generate both winners and losers. The Kaldor–Hicks "compensation principle" (Hicks 1939, Kaldor 1939) established the idea of

[1] In Europe the approach is typically denoted CBA while in the US it is often denoted benefit–cost analysis (BCA). The European tradition is followed throughout this text.

[2] http://supreme.justia.com/cases/federal/us/313/508/case.html. The US Army Corps of Engineers was heavily involved in the development of evaluation methods. The National Academy of Sciences, see Council (2004), reviews the Corps' analytical procedures and planning methods.

hypothetical compensation as a practical rule for deciding on policies and projects. When winners, at least *hypothetically*, are able to compensate losers there is a hypothetical or potential Pareto improvement. This might seem to be an elegant way of circumventing the problem of aggregating individual and unobservable utility changes due to a policy. Nonetheless it draws on a strong ethical assumption, namely, that the pure possibility of compensation is enough to motivate a project.

In reality a project might mean that those who are initially well off gain while those far down on the social ladder lose. Such an uneven distributional outcome might motivate that individuals are "weighted" according to some distributional criterion, say, income or wealth. If the weighted outcome is favorable, the project is deemed welfare improving. To illustrate, in a two-person society, Individual A gains EUR 10 while the poorer Individual B loses EUR 8. Hypothetical compensation is possible (for example, by letting A pay EUR 9 to B, turning both into winners) but if society attributes different weights to the individuals (say, 0.4 to A and 0.6 to B) and there is no compensation paid, the project might turn out to be socially unprofitable (in the example, $0.4 \cdot 10 - 0.6 \cdot 8 = -0.8$). However, then one runs into Arrow's *Impossibility Theorem*, see Arrow (1951), which states roughly that there is no social welfare function that simultaneously satisfies some reasonable conditions. The theorem is valid in an ordinal world where interpersonal comparisons of utility (changes) are meaningless, that is where our proposed welfare or distributional weights do not make sense. Allowing some stronger measurability requirements, like cardinal measurability,[3] and interpersonal comparisons opens up the possibility of meaningful social welfare functions. Such assumptions allow for the possibility of applied welfare economics, that is cost–benefit analysis. Otherwise, Arrow's theorem implies that only projects that produce winners but no losers or projects where compensation is actually paid could be recommended, assuming, of course, that the Pareto principle is accepted and that dictatorship, where projects are selected according to the preferences of a single person/body, is ruled out. (For an excellent discussion of the requirements needed in order to establish different classes of social welfare functions, for example Utilitarian, Rawlsian, Bergson–Samuelson, and so on, the reader is referred to Boadway and Bruce (1984); see also the Appendix to Chapter 7 for a brief summary.)

We discuss these issues initially since they are at the very heart of CBA, but *the aggregation problem is common to all (nondictatorial) decision criteria* and for all possible evaluation methods and approaches. Thus, cost-effectiveness/cost–utility analysis, economic impact analysis (EIA), multi-criteria analysis, or any other

[3] For example, height, output, and income are cardinally measurable, while a variable is ordinally measurable if it is possible only to rank values of the variable, for example, you may prefer muffins to waffles, and waffles to crispbread. The reader is referred to the Appendix of Chapter 7 for a short presentation of measurability and comparability requirements.

evaluation approach ultimately faces the aggregation problem. In some countries or contexts, evaluations focus on efficiency-related issues; distributional issues are considered to be either a separate issue, or as too sensitive to address. In other countries or contexts, distributional issues are considered to be an equally important part of an evaluation.

In any case, CBA subsequently conquered new worlds and found new applications beginning in the 1950s, as it was applied to various types of public projects in Europe and later on in developing countries. As popularity grew, so did the literature. For a historical review of the development of CBA in the US, the reader is referred to Zerbe Jr. (2007), for Australia to Dobes (2008), and for some UK studies to Hanley and Spash (1993). The theoretical principles of project evaluation is found in Drèze and Stern (1987), Johansson (1993), Just et al. (2004), Lesourne (1972, 1975), Myles (1995), and Jenkins et al. (2011). These manuals are quite formal and demand some knowledge of general equilibrium theory.[4] There are also many "cookbook"-style manuals providing detailed advice on how to proceed in an application. See, for example, European Commission (2008), HM Treasury (2011), Pearce et al. (2006), US EPA (2010), European Investment Bank (2013), and European Commission (2014). The current status of CBA in World Bank projects is discussed in World Bank (2010). In fact, even the US Army has its own manual, see US Army (2013), and Mansfield and Smith (2015) who describe how to undertake the evaluation of security policies within the framework of CBA. An introduction to the underlying welfare theory is found in Johansson (1991), and more technical presentations in Boadway and Bruce (1984) and Mas-Colell et al. (1995). The reader is also encouraged to visit the website of the Society for Benefit–Cost Analysis, currently housed at the Evans School of Public Affairs, University of Washington, for useful information regarding and references on CBA and related techniques; http://benefitcostanalysis.org/.

Here extensive coverage of other evaluation approaches is not provided. However, a brief account of cost-effectiveness/cost utility analysis (CEA/CUA), multi-criteria analysis (MCA), and EIA is supplied in Section 10.5.

There are many "hands-on" manuals and "cookbooks" that are extremely useful; nonetheless, in our experience, all evaluations offer unexpected theoretical (as well as empirical) surprises and challenges: hence the importance of the ability to derive relevant cost–benefit rules. This manual offers a toolkit useful not only for graduate students but also for those involved in policy evaluations at national and international organizations and consulting firms. It also serves as suitable reference for professional economists.

[4] Other books that deserve mentioning include Campbell and Brown (2003), Boardman et al. (2006), Brent (2006), Mishan and Quah (2007), de Rus (2010), and at a slightly more advanced level, Jones (2005), while Florio (2014), provides an excellent treatment of the Drèze–Stern shadow-pricing approach.

Much has changed since the early and classic manuals of Dasgupta et al. (1972), Little and Mirrlees (1974), Squire and van der Talk (1975), and others. The world was quite different then, with fixed exchange rates, much less trade, and limited international capital markets; environment and natural resource stocks played a more marginal role. Many new technical tools are now available that simplify the life of a cost–benefit practitioner. The current manual draws on such developments and addresses a number of important issues.

- Envelope properties[5] make it remarkably easy to obtain *general equilibrium* cost–benefit rules.
- Such rules for tax-distorted economies are easily related to *partial equilibrium* concepts like the marginal cost of public funds (MCPF) and the marginal excess burden of taxes (MEB).
- General *disequilibrium* rules relevant for economies with unemployment and other market imbalances can be obtained.
- Recent envelope results can be used to derive intertemporal general equilibrium rules from dynamic Ramsey growth models.
- Consistent treatment of resource stocks and the environment's assimilative capacity is possible.
- Simple but consistent measures of the social discount rate follow from dynamic models.
- Such discount rates have recently been estimated for many countries in Europe and elsewhere.
- Non-constant discounting, that is hyperbolic discounting, is starting to be used in some countries, implying that there is some recent empirical evidence to draw on.
- It is now possible to obtain (interpretable) cost–benefit rules for irreversible and stochastic scenarios, for example, using the Black–Scholes model.
- Techniques for stochastic sensitivity analysis are becoming more accessible.

In order to keep the presentation as simple and transparent as possible, the focus is on models involving a few commodities. This avoids messy (summation/vector) notation and leads to no real loss of generality as, say, x could be interpreted as a single good or as a vector. Deviations from the perfect market economy are introduced one at a time. The reader can aggregate these as needed; there are indeed many possible combinations or evaluation puzzles offered by real-world project appraisal.

[5] The envelope theorem, in economics dating back to at least Auspitz and Lieben (1889), says that only the direct effects of a change in an exogenous variable need to be considered (if evaluated at a (constrained or unconstrained) optimum). Hotelling's lemma, Shephard's lemma, and Roy's identity are (static) examples of the envelope theorem. For an intertemporal variation, see Section 4.3.

The text is structured as follows. Chapter 2 introduces a few very simple cost–benefit rules. Although they are simple they are nonetheless *general equilibrium* rules. The first couple refer to a perfect market economy, but public goods, externalities, and non-use values are subsequently introduced. Non-use values are those derived in the absence of consumption of the resource or its services. Examples include altruistic concerns for others and pure existence values when people care for endangered species. The chapter also addresses the question of how to handle externalities or damage occurring abroad, for example, if winds take a factory's harmful emissions into another country. Different kinds of market distortions are considered in Chapter 3. Sometimes agents have market power, for example, there might be a monopoly or a monopsony in a market. The chapter derives cost–benefit rules that account for market power. Another distortion relates to taxes such as value-added taxes, unit taxes, income taxes, and so on. The manual suggests one principal way to handle distortionary taxes in project appraisal. We also demonstrate that this approach yields evaluation rules that are identical to those obtained by calculating the marginal cost of public funds (MCPF) or the marginal excess burden of taxes (MEB). It is demonstrated that a second-best approach where welfare is maximized subject to a government budget constraint results in the same basic evaluation rules as the other approaches. As tradeable permits are an alternative to (emission) taxes, the treatment of permits in social evaluations of projects is addressed. Finally, sometimes prices are sticky, that is markets do not clear. The final section discusses how to handle market imbalances such as excess supply or unemployment and excess demand or shortages in a CBA.

The treatments in Chapters 2 and 3 are restricted to single-period evaluations. Chapter 4 introduces intertemporal models. Some basic concepts like discounting and comparing projects with different life spans are discussed. Then a simple model of optimal control theory is introduced and used to derive a dynamic cost–benefit rule. A Ramsey growth model is used to derive the discount rate expression most often used in CBA in the EU and in many other countries. The discussion then turns to modification of the expression if there is a capital income tax (or other capital market distortions causing a wedge between gross and net return on investments). A related question relates to the timing of an investment. How to handle this issue in project appraisal is also discussed. Finally, the chapter introduces what might be termed "Student's little helper," a simple two-period discrete-time model that is somewhat easier to handle than the more complex optimal control theory (or dynamic programming) models found in advanced macroeconomic textbooks. The model is used to derive a simple dynamic cost–benefit rule in the presence of a distortionary tax on capital. Chapter 5 introduces renewable and nonrenewable natural resources and discusses how to treat such resources in project appraisal. Forests are given a treatment separate from other renewable resources since there is a crucial time dimension due to long growth cycles. Often the extraction of natural resources cause environmental damage. The final two sections derive a dynamic cost–benefit rule given such a negative externality; both first-best and second-best rules are considered.

Thus far the text is devoted to evaluations of infinitesimally small projects. Chapter 6 provides a formal definition of small and large projects, then turns to different ways of evaluating discrete or large policy changes. One approach is Taylor series approximations. Evaluating line integrals is another. Next consumer and producer surplus changes are estimated. One challenge is that not all functions are such that the surplus measures are path independent. If the integrals are path dependent, the order in which prices and other parameters are changed affects the size of the total surplus change, making the evaluation meaningless. The chapter also introduces the non-marginal definitions of compensating and equivalent variation when a policy proposal significantly affects many prices (and other parameters). Other issues discussed are the problems in transferring an estimate of a benefit from one context to another (say, trying to value salmon fishing in one river from data for another river) and the interesting possibility that a seemingly infinitesimally small price increase might cause large effects. A final section introduces ultra-large or megaprojects that require special tools such as computable general equilibrium (CGE) models.

Having considered representative individual's economies, Chapter 7 turns to the aggregation of monetary benefits and costs across individuals. The basic problem is that simply summing across individuals is not a legitimate procedure unless the project is infinitesimally small and the welfare distribution is optimal. The chapter addresses a few different ways of handling this problem. Hypothetical compensation is one such possibility. A policy is recommended if those who gain, at least hypothetically, are able to compensate those who lose. Unfortunately, the approach does not work for policies that are so large that they significantly affect relative prices (at least this is the common view). Another possibility is to attribute explicit weights to different groups of individuals. The two final sections discuss some of the different options that are available in empirical studies and the approaches taken by a few leading empirical manuals.

Uncertainty is introduced in Chapter 8. Some conventional cost–benefit rules for a risky world are introduced. After a brief digression into the value of flexibility, which modifies the conventional net present value (NPV) criterion common in both CBA and much academic teaching, the chapter turns to a simple illustration of the Black–Scholes model, familiar from financial economics. In particular, the illustration aims at showing how to use the model in empirical CBA. A more general stochastic cost–benefit rule is also supplied. Finally, the chapter discusses the treatment of fatal risks (mortality) in project evaluations, both marginal changes of such risks, relevant, for example, small road and rail investments, and non-marginal ones, for example, nuclear disasters, some medical treatments, and similar events. A brief discussion of the implications of the risk of "doomsday" for discounting follows, and a discussion of how to evaluate natural disasters closes the chapter.

The next two chapters are devoted to a wide range of issues. Valuing benefits or costs for commodities that are not priced is an important and difficult part of many empirical evaluations. In principle, there are two available approaches, and both are covered in Chapter 9. The first approach is called a stated-preference method as surveys are used to estimate the willingness to pay (WTP) for a project (or sometimes the compensation needed in lieu of a beneficial project). In turn, there are two main variations of this approach, contingent valuation and discrete choice experiments, but Carson and Louviere (2011) have questioned this "dichotomy" and suggested a new nomenclature since the boundaries between the different approaches are becoming increasingly difficult to identify as researchers mix "ingredients" from both approaches. According to the contingent valuation method, respondents are, typically, asked about their WTP for a policy change, an open-ended approach, or whether they are willing to pay a particular amount of money for a project, a closed-ended approach. The second major variation is discrete choice experiments, where the respondent has to rank different combinations of attributes and cost is one attribute. As is (typically) the case for a closed-ended valuation question, a discrete choice experiment requires the use of econometric methods in order to arrive at an estimate of the mean or median willingness to pay. The second major approach is termed revealed preference. It exploits markets that are related to the commodity under evaluation. For example, in order to visit a natural park or undergo a particular medical examination, travel costs are incurred. Often these can be used to estimate how the related commodity is valued. In the workplace, there is sometimes a trade-off between risks and monetary rewards that can be used to assess how fatal or nonfatal job risks are valued. The question, though, is how to interpret market demand functions: can they be interpreted as reflecting the preferences of a representative consumer? This issue, among others, is also addressed in the chapter.

Chapter 10 is devoted to a few important topics in evaluation. The first section reports the magnitude of the discount rate in CBA for different countries. The next topic is whether to include wider economic benefits in a cost–benefit analysis and what such wider benefits might consist of. The chapter then turns to a brief discussion of the difference between the approach used in this text and the classic shadow-price approach associated with Drèze and Stern (1987) in the *Handbook of Public Economics*. The conventional cost–benefit approach has been challenged by recent developments in behavioral economics and happiness economics, which are briefly discussed. The chapter's final section is devoted to a brief presentation of cost-effectiveness analysis, cost–utility analysis, MCA, and EIA.

Since there typically is considerable uncertainty with respect to "true" parameter values it is important to add a sensitivity analysis to a point estimate evaluation. Such an analysis might be deterministic, providing upper and lower bounds for critical parameters, or stochastic, drawing on simulation tools such as Monte Carlo techniques. Both deterministic and stochastic approaches are discussed in Chapter 11. A brief section on due diligence and evaluations is added. This is an important issue since it is far from unusual that CBA as well as other evaluation techniques are intentionally or unintentionally misused so that the reported results are biased and misleading; the benefits of a preferred project or the costs of an opposed project are exaggerated.

2

The basic cost–benefit (C–B) model

After a brief introduction of some necessary tools, such as Lagrange functions, indirect utility functions, and profit functions, the chapter turns to the basic model relating to a single-household economy, used to derive a few general equilibrium cost–benefit rules. As not all goods and services are priced in markets, the chapter discusses how to handle non-priced private commodities as well as public goods and bads (externalities) in project evaluations. A particular class of values, non-use values, is also introduced. Non-use values include concern for other people living now or in the future (altruism) and concern for endangered species (existence values). Sometimes a project causes effects in foreign countries; such international spillovers and their inclusion in evaluations are also discussed.

2.1 A quick refresher course in micro

In this section we provide a brief presentation of some simple maximization problems that constitute the foundation for the toolkit used in deriving cost–benefit rules in this and the next chapter. However, because we focus here on the principles, the problems are slightly simpler in terms of number of commodities than they are later on. Therefore, we refer the reader to the Appendix at the end of this chapter for a full treatment. Consider a simple constrained utility maximization problem:

$$\max_{[x^{dn},x^{de}]} U(x^{dn}, x^{de})$$
$$\text{s.t.} \quad y = p^n \cdot x^{dn} + p^e \cdot x^{de}, \tag{2.1}$$

where the utility function is assumed to be well behaved so that first-order and second-order conditions for an interior optimum are satisfied. There are just two commodities, x^{dn} and x^{de}, that the household consumes. As only relative prices matter in this kind of model the price of commodity n (the numéraire) is set equal to unity (and this price, $p^n = 1$, is suppressed in what follows); recall that if all prices and income double, real income remains unchanged.[1] The price of the other commodity is p^e, and lump-sum income is y. To solve the problem, formulate the Lagrangian function:

$$L(.) = U(x^{dn}, x^{de}) + \lambda(y - x^{dn} - p^e \cdot x^{de}), \tag{2.2}$$

where λ is a Lagrange multiplier. First-order conditions for an interior solution are obtained by differentiating the Lagrangian with respect to x^{dn}, x^{de}, and λ and setting the derivatives equal to zero; $\partial U(.)/\partial x^{dj} = U_{x^{dj}} = \lambda p^j$ for $j = e, n$. Differentiating with respect to λ yields the budget constraint. Solving these conditions obtains the *Marshallian demand functions* $x^{dn}(p^e, y)$, $x^{de}(p^e, y)$ and the function $\lambda(p^e, y)$. Using these in the direct utility function obtains the maximum-value function, that is, the *indirect utility function*:

$$V(p^e, y) = U[x^{dn}(p^e, y), x^{de}(p^e, y)]. \tag{2.3}$$

What do the partial derivatives of this function look like? The easiest way to arrive at the answer is to take the partial derivative of the Lagrangian with respect to p^e and y, respectively, to obtain

$$\begin{aligned} \frac{\partial L(.)}{\partial p^e} &= -\lambda \cdot x^{de} \\ \frac{\partial L(.)}{\partial y} &= \lambda. \end{aligned} \tag{2.4}$$

The (constrained) *Envelope Theorem* says that the *total* effect of a small parameter change can be obtained by simply taking the *partial* derivative of the Lagrangian with respect to the parameter and evaluating it at the *optimum*. Note that x^{dn} and x^{de} are changing in response to the price or income increase, but because the consumer is optimizing the entire time and facing a budget constraint, these indirect effects sum to zero; compare Equation (2.6) below. Refer to a textbook on microeconomics, for example, Jehle and Reny (2011, pp. 604–607) or Varian (1992, pp. 502–503), for details. Thus we have

$$\begin{aligned} \left.\frac{\partial L(.)}{\partial p^e}\right|_{\mathbf{x}^d(p^e,y),\lambda(p^e,y)} &= \frac{\partial V(.)}{\partial p^e} = -\lambda(p^e, y) \cdot x^{de}(p^e, y) \\ \left.\frac{\partial L(.)}{\partial y}\right|_{\mathbf{x}^d(p^e,y),\lambda(p^e,y)} &= \frac{\partial V(.)}{\partial y} = \lambda(p^e, y), \end{aligned} \tag{2.5}$$

where $\mathbf{x}^d(p^e, y) = [x^{dn}(.), x^{de}(.)]$, and the vertical bar to the right of the partial derivative denotes that we are to make this evaluation at the optimum. The impact

[1] Alternatively, lump-sum income can serve as a numéraire while both commodity prices are flexible.