10

Benefit-Cost Analysis

Objectives

1. Define the three metrics: (1) net present value (NPV), (2) benefit-cost (BC) ratio, (3) and internal rate of return (IRR).
2. Describe the general challenges in calculating these three metrics.
3. Consider how to interpret these metrics.
4. Review evidence on BC analysis across educational research.

BC analysis allows us to determine if an educational investment is socially efficient. This determination is made when the monetized benefits—resources accrued as a result of the investment—exceed the costs, which are all the resources used to implement the investment. The method for calculating the costs of an intervention was described in detail in Chapters 4 through 6, and the methods for estimating the benefits were covered in Chapter 9 (and prefigured in Chapter 7). We assume here that costs and benefits have been correctly measured and that they correspond to one another. This chapter serves as the capstone for application of the ingredients method and shadow pricing techniques. Here, we bring the costs and benefits together to derive an economic metric that is informative for decisionmakers who are interested in the efficiency of educational investments.

We begin by describing the economic metrics—the NPV, BC ratio, and IRR. The metrics presented are mathematically straightforward. But this simplicity belies the challenges in interpreting them and placing them in context. We devote the following section to discussing some of these challenges. Finally, we provide an illustrative review of evidence generated from BC analyses. This review is not intended to summarize the range of literature or adjudicate between educational investments. It is intended to illustrate the main areas of BC analysis in education.

This chapter may appear to be the culmination of all the research inquiry into costs and benefits. In fact, there are still several more important steps to follow. These relate to checking the robustness of the results and interpreting the results for policymakers. These steps are documented in Chapters 11 and 12. The material covered here is a necessary precursor to following these steps.

10.1. COMBINING BENEFITS AND COSTS INTO ECONOMIC METRICS

In this section, we discuss three economic metrics: (1) NPV, (2) the BC ratio, and (3) the IRR. We illustrate each of these metrics using a stylized example of an adult literacy program. The stylized dollar flows are shown in Table 10.1.

The literacy program helps adults adapt to the labor market and obtain more highly paid jobs. Using the ingredients method, it is determined that the 1-year program costs $300 per participant. After completion, the participants earn $150 extra per year compared to what they would have earned without the program; this gain lasts for 4 years. Having accurately collected this information, creating the

Table 10.1 Stylized Example: Literacy Program for Benefit-Cost Analysis

	Undiscounted Costs	Undiscounted Benefits
Year 1	$300	0
Year 2	0	$150
Year 3	0	$150
Year 4	0	$150
Year 5	0	$150

economic metrics is mathematically simple, although each one requires careful interpretation.

10.1.1. Net Present Value

The primary economic metric for BC analysis of educational interventions is the NPV. The NPV is the discounted value of the benefits minus the discounted value of the costs. To discount, we apply the formulae from Chapters 5 and 9 respectively:

$$B_{PV} = \sum_{t=1}^{n} \frac{B_t}{(1+i)^{t-1}} \text{ and } C_{PV} = \sum_{t=1}^{n} \frac{C_t}{(1+i)^{t-1}}$$

Where B_t and C_t are the benefits and costs, t is the year in a series ranging from 1 to n, and i is the discount rate. Hence, the NPV of a project is straightforwardly calculated as follows:

$$NPV = B_{PV} - C_{PV}$$

Interventions with higher NPV amounts are preferred, and there is a strong presumption to reject any interventions with NPV amounts less than zero.

We can apply this to the simple example that was described in Table 10.1. Assuming a discount rate of 3%, the discounted sum of benefits in our example is given by the following:

$$B_{PV} = \frac{150}{(1+0.03)^1} + \frac{150}{(1+0.03)^2} + \frac{150}{(1+0.03)^3} + \frac{150}{(1+0.03)^4}$$

$$B_{PV} = 146 + 141 + 137 + 133$$

$$B_{PV} = \$557$$

As all the costs are incurred immediately, the discounted costs are as follows:

$$C_{PV} = \frac{300}{(1+0.03)^0} = \$300$$

Therefore, the NPV is calculated straightforwardly as this:

$$NPV = \$557 - \$300 = \$257$$

Given the NPV is clearly above zero, we can conclude that investment in this literacy program yields a positive stream of resources if we discount the future at a rate of 3%. Prima facie, the program is a good investment. Certainly, if there is an alternative program that costs about the same ($300) and has an NPV of only $100, for example, then this literacy program is clearly preferred.

However, we need to interpret this result carefully. The extra earnings count as social benefits, but they flow to the participant directly. The investor (such as the local government) will not recoup these extra earnings. If we had adopted a narrower perspective, the NPV and actual dollar amounts would correspond more closely (e.g., if we had only counted the taxpayer benefits and the taxpayer costs). Also, this NPV = $257 amount should be compared to NPV figures for similarly sized projects. A program that costs $100 with discounted benefits of $357 may be preferred as less risky. As well, the NPV should be compared to programs of similar duration. For example, we can imagine an intervention that costs the same ($300) and yields the same discounted benefits ($557), but these benefits come from one lump-sum undiscounted gain (of $867) in 10 years' time. In this case, the literacy program is almost certainly preferred. It is less risky, and it has an option value in that the NPV is received within 5 years such that we can reinvest over the remaining 5 years before the second intervention yields its returns.

The NPV metric has the advantage of being the most straightforward to report and interpret. The school invests $300, and there are benefits in extra earnings of $557, yielding a surplus (akin to a profit) of $257. However, it is often difficult to compare NPVs because the scale of the program makes such a difference to the final number. A $20 million program—even with very modest benefits—will almost certainly yield a higher NPV than a $2 million program. It is possible to express the NPV per participant. But then it may be unclear what total amount of resources is required, and it may give the impression that the NPV is constant as the number of students expands. The simplicity of the NPV is therefore traded off against the difficulty of comparing benefits and costs across programs in ways decisionmakers can use.

10.1.2. Benefit-Cost Ratio

The BC ratio is a simple adaptation from the NPV metric. Instead of taking the difference between present value benefits and costs, we divide benefits by costs:

$$BCR = \frac{B_{PV}}{C_{PV}}$$

A BC ratio greater than 1 is one where the benefits exceed the costs. Interventions with higher BC ratios are preferred, and there is a strong presumption that interventions with BC ratios less than 1 (i.e., where costs exceed benefits) should be rejected.

For our stylized literacy program in Table 10.1, the BC ratio is as follows:

$$BCR = \frac{557}{300} = 1.86$$

This ratio is clearly greater than one and therefore indicates that the program is a good investment with a discount rate of 3%. As a shorthand explanation, the BC ratio is often interpreted as "for every dollar invested in this literacy program, there will be a return of \$1.86." As with the NPV, however, we caution that this interpretation does not imply that the program yields the same NPV for each marginal dollar invested.

Here, we can see why we emphasize the term *negative benefits* rather than induced costs when performing BC analysis. Imagine a college access program with a cost of \$2 million, benefits of \$4 million in extra earnings, but also with \$1 million spent on newly induced college enrollment. If the extra college enrollment is counted as a negative benefit, the BC ratio is 3/2 = 1.5. However, if the extra college enrollment is counted as an induced cost, the BC ratio is 4/3 = 1.33. As there may be many negative benefits that necessitate discounting, the former way of calculating the ratio is preferred.

The advantage of the BC ratio is that it can be easily applied in comparisons of investments as if it is a return on investment. For example, a program with a BC ratio of 4 offers a higher return than a program with a BC ratio of 2. Looking across a set of different interventions, the decisionmaker might rank BC ratios and choose those that are the highest. This approach is valid in some contexts—for example, when the programs are of similar scale, riskiness, and duration—but not always. Moreover, the BC ratio may offer a simplistic interpretation: The shorthand explanation might lead decisionmakers to think that they can invest in the program in any dollar amount, whereas measured benefits in this case are tied to implementing the program at a particular scale, in this case the \$300 level.

10.1.3. Internal Rate of Return

A third economic metric is the IRR. This is the rate of interest that equates the present value of benefits to the present value of costs.

Formally, the IRR is defined as the discount rate (i) that causes the NPV or net benefits to equal zero:

$$NPV = \sum_{t=1}^{n} \frac{B_t}{(1+i)^{t-1}} - \sum_{t=1}^{n} B_{PV} - C_{PV} = 0$$

Or, equivalently, this:

$$NB = \sum_{t=1}^{n} \frac{B_t - C_t}{(1+i)} = B_{PV} - C_{PV} = 0$$

The IRR for an educational investment can be compared to the interest rate on investments of comparable size and duration. So, if the funding agency can invest in programs that yield returns of 10%, any educational intervention with an IRR above 10% would represent a good investment; any intervention with an IRR below 10% would be presumed to be rejected.

In our stylized literacy program example as per Table 10.1, the IRR turns out to be approximately 0.349 (or 34.9%). The discounted costs of the program are $300, so we need to find the discount rate that will make the discounted sum of benefits equal to $300. We can identify this IRR by calculating the discounted sum of benefits when i = 0.349, which can be obtained iteratively or by using automated software algorithms (e.g., Excel or many electronic spreadsheets):

$$\begin{aligned} B &= \frac{150}{(1+0.349)^1} + \frac{150}{(1+0.349)^2} + \frac{150}{(1+0.349)^3} + \frac{150}{(1+0.349)^4} \\ &= 111 + 82 + 61 + 45 = \$300 \end{aligned}$$

This IRR at 34.9% is considered attractive. There is no mathematical threshold for interpreting the IRR (other than that it should be positive). However, given that most government agencies work with interest rates below 10%, a figure of 34.9% appears to be very high.

There are advantages to using the IRR as an economic metric for evaluating educational programs. Inherently, calculating the IRR does not set out an assumed discount rate. As the value for the discount rate can make a big difference to the NPV, this freedom may be useful. Of course, it is still necessary to compare the IRR to another threshold interest rate for decisionmaking purposes for comparing the profitability of the investment. So this freedom is not fully liberating.

Intuitively, the IRR may be appealing because it can be readily placed in context. Most investors would regard somewhere between

5% and 10% as a reasonable interest rate on their investments. (Private individuals tend to discount at much higher rates—for example, 20%; see Warner and Pleeter [2001]—but a reasonable social interest rate should be below that of private investors.) So any IRR above 10% would suggest that the program represents a good investment. However, the IRR is quite sensitive to changes in the stream of benefits and costs and so must be interpreted cautiously. For example, the IRR for the literacy program given its current flows is 34.9%, which is a very high rate. If we learn that in fact the undiscounted costs were 10% higher ($330) and the undiscounted benefits in each year were 10% lower ($135), this does not mean that the IRR falls by 10% or even by 20%. Instead, the new IRR becomes 23.1%, a reduction of one third from its initial value. If the stream is longer, this sensitivity is magnified: If the original undiscounted benefits are spread over 12 years instead of 4, the IRR falls to 13%.

Also, the IRR does not provide any indication of the project scale. Thus, we could estimate identical IRRs for two separate projects, indicating that they are equally desirable, even when the NPV of one project is larger. An IRR of 34.9% is attractive to a policymaker for a project of $1 million, but it is especially attractive for an investment of $100 million. Finally, it is sometimes difficult to calculate a unique value for the IRR. This does not occur in instances like our numerical example in which all the costs occur at the beginning of the project and benefits come later. This stream closely parallels most projects, but it need not. If costs and benefits are dispersed unevenly throughout the project cycle, it is sometimes possible to calculate more than one IRR.

Finally, we distinguish the IRR from the idea of the social return on investment (SROI). This term is often used loosely to refer to the returns an enterprise obtains on its philanthropic investments (for a definition, see Millar & Hall, 2013). In theory, the enterprise should conduct BC analysis to derive the IRR. However, the enterprise may be interested in its own independent impact along certain unique dimensions. For example, a private company might contribute $1 million to a mentoring program for at-risk youth; this contribution triggers matched funding of $1 million by the local government. If an evaluation establishes that delinquent behaviors fell by 10%, this reduction in delinquency can be valued by its shadow price. The private company might calculate the value of all reduced delinquency and compare that to its $1 million contribution rather than the total resource cost of $2 million: The private company regards the matched funding as leveraged funding, which its initial contribution has created. (In other examples, the benefits might be defined more narrowly than they would for a full social BC analysis.) This SROI is therefore not strictly based upon a full social perspective. As such, it is not an IRR.

10.1.4. Break-Even Analysis

A final metric can be derived using break-even analysis. In this case, the analysis is constrained such that the present value benefits equal the present value costs within a given period of time (sometimes referred to as the payback period). So, costs are calculated over the entire project length, but the benefits are only counted over the time period until they sum to the value of the costs. The analyst then reports the time period as the break-even point for the intervention. Formally, we are identifying the value for *n* in the following equation:

$$NPV = \sum_{t=1}^{n} \frac{B_t}{(1+i)^{t-1}} - \sum_{t-1}^{n} \frac{C_t}{(1+i)^{t-1}} = B_{PV} - C_{PV} = 0$$

For the stylized example, the break-even point is partway through the third year (assuming a discount rate of 3%). After 2 years, the discounted benefits are $279, which is just below the $300 cost. After 3 years, the benefits exceed the costs.

Clearly, this is a shorthand metric for an economic evaluation: It simply indicates how long the investor must wait before recouping the investment. The advantage is that it is very simple to explain in terms of, for example, "after three years, the benefits will cover the costs." This might be helpful to a policymaker with a limited time horizon. It also might be useful in the context of uncertain policy contexts and outcomes; in some contexts, projecting forward 10 or 20 years may seem to require a tremendous leap of faith.

However, this metric is simplistic. It does not provide information on the complete value of the investment: Lagged but high-return education programs will look less valuable than programs with an immediate payback. The metric invites comparisons—based on the shortness of the break-even time period—which may not be legitimate. Early education programs, in particular, will have later break-even points than youth or college programs but may have higher NPVs because of higher lifetime efficacy. Finally, this break-even metric may be misleading if benefits are spread across groups. For most educational investments, the benefits accrue to society or the broader community and not to any specific individual or government agency. Thus, a preschool program run by the state will accumulate benefits to the participants, local taxpayers, and the broader community. The total social returns may be equal to the costs after 5 years, but the separate groups do not receive all the returns. Adopting a narrow perspective, some investments may yield returns that are never recouped in actual money terms.

10.2. PERFORMING BENEFIT-COST ANALYSIS

The main challenges in performing BC analysis are measuring costs, calculating benefits, and making sure it is legitimate to combine them—the mathematics of the metrics is straightforward. Indeed, given that they are directly linked, it makes sense for the analyst to report all three metrics where possible. This obviates the need to consider which one is the most relevant. It also provides more information. Regardless of the number of metrics reported, the analyst should explain the results from BC analysis in a way that allows the reader to interpret the results, to place them in a policy context, and—essentially—to make better decisions.

As a summary for our stylized example in Table 10.1, we can say that the present value benefits are $557 compared to costs of $300; this is a BC ratio of 1.86 and yields an IRR of 34.9%. The break-even point for this investment is in the third year, post-intervention. Each metric conveys useful information, and we note that an IRR of 34.9% might "sound" better than $1.86 returned for each $1 invested. Nevertheless, in this case, the results prompt the same conclusion: The program is efficient from this social perspective. See Example 10.1 for the results for the HighScope Perry Preschool Program.

Example 10.1 Benefit-Cost Results for the HighScope Perry Preschool Program

Here, we bring together the evidence on costs and benefits of the HighScope Perry Preschool Program (see the table after this paragraph). The costs were calculated using the ingredients method (see details in Example 5.1 in Chapter 5). The benefits were derived from surveys of the subjects at ages 27 and 40, many years later to compare outcomes by random assignment to the program relative to a control group. Benefits were valued by shadow prices using a range of methods (see details in Chapter 9). The costs and benefits are expressed relative to the status of a child who did not participate in the program.

(Continued)

(Continued)

Benefit-Cost Analysis: HighScope Perry Preschool Program

	Present Value at Program Start From the Perspective of:		
	Participants	**General Public**	**Total (Society)**
Program Costs [C]	–	$20,947	$20,947
Measured benefits up to age 40 (child care, K–12 education, adult education, college, earnings, crime, welfare)	$14,944	$105,060	$120,004
Projected benefits after age 40 (earnings, crime, welfare)	$18,233	$44,859	$63,092
Total Benefits [B]	$33,177	$149,919	$183,096
Net Present Value [B – C]	$33,177	$128,972	$162,149
Benefit-Cost Ratio [B/C]	n/a	7.2	8.7
Internal Rate of Return [IRR]*			8.1%

Source: Adapted from Barnett (1996); Nores, Belfield, Barnett, and Schweinhart (2006).

Notes: *IRR from reanalysis by Heckman, Moon, Pinto, Savelyev, and Yavitz (2010). 2015 dollars rounded. Discount rate 3%. Annual cost adjusted for years of participation.

The results for the economic metrics are given in the table. For participants, there are no costs; the net present value (NPV) is therefore the sum of the benefits at $33,177. This is the gain in economic well-being for participants accrued over the lifetime but expressed as a lump sum at age 4. For the general public, the program cost is $20,947. This amount is what the program is expected to cost in 2015 dollars, but it is not adjusted for possible changes in ingredients based on changes in either relative prices or technologies. The benefits for the general public—again, expressed as a lump sum at the same time the costs are incurred—are $149,919. The NPV is therefore $128,972, and the benefit-cost (BC) ratio is 7.2. From the social perspective, the costs are $20,947, and the benefits are the sum of the benefits to participants and the general public at $183,096. The NPV is therefore $162,149, and the BC ratio is 8.7. Finally, from a reanalysis by Heckman and colleagues (2010) using alternative models of benefits spread over 36 years, the IRR for society is 8.1%.

Overall, the BC analysis of the HighScope Perry Preschool Program indicates that the program is efficient from all perspectives and most likely yields returns that exceed alternative uses of investment funds.

We emphasize that reporting the results of BC analysis is not a substitute for decisionmaking—that is, a formal "reasoned determination" based on the evidence and other considerations. As we discuss in Chapter 11, policymakers should not simply rank interventions based on their net benefits and choose the one with the highest net benefits. Here, we note a number of important issues that might undermine our ability to compare programs according to their NPVs, BC ratios, or IRRs.

To begin, it is essential to perform BC analysis from the appropriate perspective and for affected populations. The default perspective is that of an entire society—counting all resources—but alternative perspectives for subpopulations or constituencies are often informative. Just as the ingredients spreadsheet can be divided according to funding agency, the stream of benefits and hence NPV results can also be divided according to funding agency. Taxpayers, in particular, may be interested in the amount of public funding versus the amount of public benefits from each educational investment (e.g., Trostel, 2010). Shaffer (2011) recommended performing multiple account BC analyses—that is, analyses that take a particular perspective within an overall BC analysis and identify the winners and losers (beneficiaries and payers or students and taxpayers). For health evaluations, Neumann, Sanders, Russell, Siegel, and Ganiat (2016) stated that results should always be reported from both a societal and payer perspective. Certainly, results will vary considerably depending on the perspective adopted, so it is important the perspective adopted is clearly stated and justified by the analyst.

Also, it may be informative to assess the distribution of benefits across different groups in society. Groups can be defined quite broadly. Among program participants, for example, we may wish to separately calculate benefits by income level, gender, or race, in order to assess whether one group obtains a larger share of benefits. As with costs, it is common that benefits received by one group are not the same as benefits received by another group. The BC ratios may therefore be very different across subgroups. But heterogeneous results should not be overly interpreted: Policymakers should not infer that when, for instance, NPVs are higher for boys than girls, the analyst is recommending greater investments in boys. There are important issues related to equity and equality that should determine investments across different groups. This is another reason why it is important to separate out the BC result from policy decisions.

There are many factors associated with research study design that might influence BC analysis across different interventions. For example, interventions may start at different ages (preschool through to college). They may count a different array of benefits (including some crime or health status impacts for example) and look over a different time horizon. Studies may also apply a different set of shadow prices or apply different benefit transfer procedures. Finally, interventions may cost very different amounts such that some interventions may not be financially viable in all contexts. For example, the Abecedarian preschool program is over 4 times as expensive as the HighScope Perry Preschool Program to implement (Barnett & Masse, 2007, Table 2). As the interventions vary across more dimensions, comparisons become even less tenable (see the discussion in Harris, 2009).

One critical element of most research designs is the ability to test for statistical significance. In most empirical research, the underlying framework is one of hypothesis testing: By convention, the determination to reject the null hypothesis is based on the significance level of the *p* value, which is usually 1%, 5%, or 10% (but for a critique of statistical significance, see Wasserstein & Lazar, 2016). In contrast, BC analysis is not motivated by hypothesis testing but instead by guiding decisionmakers. Therefore, statistical significance is less important, and it need not be influential in selecting benefits for inclusion in BC analysis. As declared by Farrow and Zerbe (2013, p. 370), "Statistical significance levels for program and policy effect size are not relevant to BCA [benefit-cost analysis]. Regardless of the associated level of significance, all estimated effects should be included in the BCA model with the appropriate standard error." In other words, the researcher should include the average of each estimated effect in the baseline BC analysis and then report results for the variance as part of the sensitivity testing. Insofar as BC analysis involves hypothesis testing, the hypothesis is that the NPV is greater than zero. This is clearly different from testing if each impact is statistically significant. We recognize that this practice—considering statistical significance irrelevant—may not be agreed upon or followed in all research studies. The analyst must therefore be clear when combining benefits and costs what assumptions are made about statistical significance.

These issues are best discussed with examples. In the next section, we provide examples of BC analyses that have been performed to illustrate the main results and reporting conventions. For these examples, we focus on the baseline results. In Chapter 11, we provide a detailed discussion of uncertainty, sensitivity testing, and distributional analysis. In that chapter, we discuss more thoroughly how the results might vary depending on the assumptions.

10.3. EXAMPLES OF BENEFIT-COST ANALYSIS

10.3.1. Investments in Preschool

By far the most attention in BC analysis of educational interventions has focused on preschool and early childhood education. Early studies applied experimental methods, which can more readily accommodate BC methods, and the strong results from these studies stimulated broader interest in the returns to early education. Summaries of the evidence on the benefits and costs of early education program are given in Barnett and Masse (2007); Bartik, Gormley, and Adelstein (2012); Duncan and Magnuson (2013); Institute of Medicine and National Research Council (2014); and Karoly (2012).

One prominent program is the Chicago Child-Parent Center program (Reynolds, Temple, White, Ou, & Robertson, 2011). This program provides support for children from age 3 to third grade, including reading and math instruction, a parenting program, outreach to help enrollment, and health/nutrition services (Reynolds, Temple, Robertson, & Mann, 2002, p. 272). The preschool program was delivered through preschool centers that were located to serve children in low-income neighborhoods who had limited access to alternative early education services.

In Table 10.2, we report the BC results for the preschool component of the Chicago Child-Parent Center program (Reynolds et al., 2011). The comparison group received business-as-usual services available in the local area or no comparable services. Following the published evidence, we report the results per child rather than in total; on average, each center served 100 to 150 children so from a policy perspective these per-child numbers might be multiplied by 100 to 150 to indicate the amount of total funding needed per operating unit.

The program was extremely valuable for the families: Benefits by age 26 were valued at $10,080 with an estimated $26,520 in benefits from earnings and other behaviors. As there was no cost to the families, the NPV is $36,600. Of key interest are the BC results for the general public and for society (the sum of individuals and the general public). Expressed at the initial time of the program (when the child is age 3), the discounted cost of the program is $10,060. Looking across the measured benefits up to age 26, the general public accrues benefits of $37,020; projecting forward, additional benefits of $33,700 are expected. The NPV is therefore $60,660, and the BC ratio is 7.1—that is, the program yields general public benefits that are 7 times the costs of the program. From the social perspective—that is, including the benefits to

Table 10.2 Benefit-Cost Analysis of the Chicago Child-Parent Center Program

	Present Value at Age 3 From the Perspective of:		
	Participants	General Public	Total (Society)
Program Costs [C]	–	$10,060	$10,060
Estimated Benefits at age 26 [B1]			
Child care, child abuse and neglect, K–12 education, juvenile crime (net of college)	$10,080	$37,020	$47,100
Projected Benefits [B2]			
Earnings, adult crime, health	$26,520	$33,700	$60,210
Total Benefits [B = B1 + B2]	$36,600	$70,720	$107,310
Net Present Value [B – C]	$36,600	$60,660	$97,250
Benefit-Cost Ratio [B/C]	n/a	7.1	10.8

Source: Adapted from Reynolds et al. (2011, Table 4).

Notes: 2015 dollars rounded to nearest 10. Discount rate 3%. Preschool program.

participants and the general public net of transfers between the two groups—the NPV is $97,250 per child and the BC ratio is 10.8. Finally, from detailed tabulations in Reynolds et al. (2011, Table 4), the program's break-even point is before the end of high school (age 18). For society as a whole, this investment appears to provide a very high yield.

Example 10.2 Rate of Return Studies

In many examples of educational benefit-cost (BC) analysis, the internal rate of return (IRR) is not calculated, and analysts rely on the net present value (NPV) and the BC ratio. However, the IRR is used almost exclusively in estimating the benefits and costs of obtaining additional years of schooling.

In many low-income countries, a large portion of the young population does not attend school, even at the primary level. Governments are forced to make difficult decisions about which levels of education—primary, secondary, or higher—should be the recipients of scarce investment funds. To allocate these resources across levels of education, one could attempt to compare the costs and the benefits of each of the three alternatives. The investment that yields the highest net benefits—or BC ratio or IRR—would produce relatively greater benefits for a given cost. In fact, hundreds of studies have done exactly that, albeit with somewhat restricted definitions of what constitutes benefits and costs. For an extensive review of this literature, see Psacharopoulos and Patrinos (2004).

Figure 10.1 illustrates a basic schematic that is followed in estimating the benefits and costs of education (for further details on the method, see Barrow & Malamud, 2015; Carnoy, 1995). The researcher first obtains data on individual earnings, usually from a census or household survey. Using these data, the researcher constructs an "age-earnings profile" for each level of education, which traces out the average lifetime earnings of individuals who have attained a given level of education. This approach for estimating benefits is discussed in Chapter 9. The figure depicts hypothetical age-earnings profiles for two levels of education: (1) secondary and (2) postsecondary. The secondary profile begins at age 18, following graduation from secondary school; the postsecondary education profile begins at 22, after graduation from university. Both end at the retirement age of 65, when individuals cease working.

The benefits of higher education are calculated as the difference at each age between what individuals earn as higher-education graduates and what they might have earned as secondary graduates. This is represented by the Area B. The costs of higher education are divided into two components. The first is the cost of income forgone while receiving a university degree—an opportunity cost of studying instead of working (represented by the area C1). The second includes all the direct costs of studying, like books, tuition, and so forth (represented by the area C2).

In this simple framework, the IRR for university education is calculated by finding the discount rate that equalizes the discounted sum of benefits

Figure 10.1 A Schematic for Calculating Rates of Return

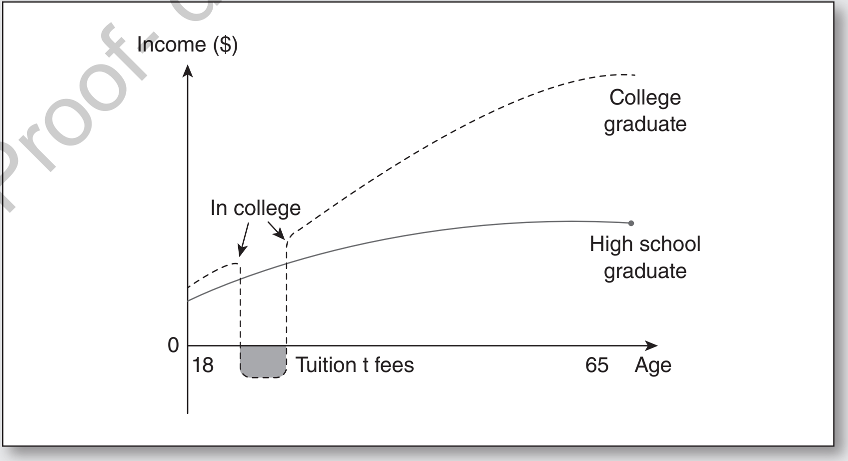

(Continued)

(Continued)

and the discounted sum of costs. Often, a shortcut is applied, using education premia from cross-sectional earnings equations. The general consensus from this approach, mostly based on studies spanning back over 30 years, is that the average IRR to university education is somewhat lower than that of primary education (see Psacharopoulos, 1994, Table 1; Psacharopoulos & Patrinos, 2004). In a simplistic way, this suggests that primary education is a better candidate for scarce investment funds than higher education.

Despite the relative ease of conducting these studies, there are also many critiques of the method (Glewwe, 1999; McEwan, 1999). First, the method assumes that the sole benefit of education is higher earnings, despite suggestive evidence of many other benefits. Second, it focuses exclusively on the rate of return to a greater *quantity* of education, when decisionmakers may be more interested in the return to the *quality* of education. Third, many authors assume that the only cost of education is forgone income (C1) when it is well known that the direct costs of education may be substantial (C2). Fourth, the age-earnings profiles are usually constructed with data from a single cross section, rather than longitudinal data that track a group of workers over their careers. Implicitly, this assumes that the earnings of a 65-year-old today, adjusted for price level, are a good approximation of what a 25-year-old will earn in 40 years. Finally, it assumes that the earnings of a current high school graduate are a good approximation of what current university graduates *would have* earned without a degree.

10.3.2. Investments in Youth

Educational interventions to enhance youth development and economic well-being can also be evaluated using BC analysis. There is a case that early investments have the highest payoff (NPV); this is represented graphically as the Heckman curve (www.heckmanequation.org) and is summarized by the idea of "prevention being more efficient than remediation." But there are good reasons why investments in youth might be a priority. The time between costs and benefits is shorter: Most educational benefits are from higher earnings, and youth are closer to the labor market (obviously) than preschoolers. Also, youth behavior has more behavioral consequences, not just in the labor market but also with respect to crime, health status (e.g., through substance abuse), and other social acts (e.g., teenage pregnancy). Thus, there are potentially more benefits during youth. Finally, investments in youth might be targeted more accurately. An educational intervention such as class size reduction in kindergarten might have considerable benefits (Chetty et al., 2011), but it is typically applied to all students within a given classroom and is therefore expensive per at-risk child. By high school, there is more evidence on which subset of students need remedial or developmental supports; resources are therefore required for fewer students. Ultimately, of course, whether investments in youth or preschool pay off is an empirical question.

Youth programs do appear to have many benefits that exceed costs. Indeed, programs that are targeted at at-risk youth, often high school dropouts, can be very high yield especially if they reduce involvement in the criminal justice system (for YouthBuild, see Cohen & Piquero, 2015; for Job Corps, see McConnell & Glazerman, 2001). A detailed BC analysis of National Guard Youth ChalleNGe Program (NGYCP) was performed by Perez-Arce, Constant, Loughran, and Karoly (2012). NGYCP is an intensive program for high school dropouts aged 16 to 18. In a residential program initially for 22 weeks, National Guard participants receive discipline training, academic instruction, and fitness programs The following year entails the next phase, comprising more education, training, and employment counseling services. Outcomes were modeled as earnings gains mediated through changes in educational attainment, plus changes in welfare receipt and criminal activity (service to the community was also counted as a benefit). The analysis was novel in that it included an estimate of the marginal excess tax burden (METB) and the opportunity cost of time for the participants (for a cost-effectiveness [CE] analysis of NGYCP, see Hollands et al., 2014). Expressed as present values at age 17 using a 3% discount rate, the cost per admittee was $16,825, and the present value benefit was $44,674 (2015 dollars). This yields an NPV of $27,848, a BC ratio of 2.66, and an IRR of 6.4%. Recognizing that the measured benefits were not exactly the same, these values are comparable to some early childhood investments (see Karoly, 2012, Table 2; Temple & Reynolds, 2015).

10.3.3. Benefit-Cost Analysis in Developing Countries

Increasingly, educational interventions in developing or low-income countries are being evaluated using BC analysis (Dhaliwal, Duflo, Glennister, & Tulloch, 2012; McEwan, 2012). Given the variation in education systems, it is difficult to generalize from this literature. Also, these interventions are often a bundled set of services that includes an education component and other investments and program services. One prominent program in low-income countries, the Graduation Program, is an intensive, multicomponent program involving health education and technical skills training along with asset transfer, consumption support, and home visits (see Banerjee et al., 2015). The influence of just the educational services on outcomes is hard to isolate. Moreover, the benefits are specified—and valued in shadow prices—in diverse ways, over different time horizons, and in very different economic contexts.

One general conclusion is that investments in education for girls have an extremely high payoff (World Development Report, 2011). Typically, girls' schooling has been very limited, yet their role in the household and with respect to fertility and child-rearing is so important that the social benefits of education are high. There are also significant benefits mediated through the effect of education on health. This is illustrated in BC analysis of expanding primary school enrollments for girls in Tanzania by Brent (2009). The expansion—a projected 1% increase in enrollments—is valued in terms of reduced rates of HIV/AIDS infections. The shadow price of HIV/AIDS is derived from the value of human capital in the labor market. Even applying a conservative specification of benefits, expansions of primary schooling yield social benefits with a BC ratio of 1.3 to 2.9 times the present value costs of education provision.

10.3.4. Programs to Increase Wages for Welfare Recipients

There is sizable evidence on the benefits and costs of training programs. Much of this evidence applies experimental methods to identify the returns to program participants (see Redcross, Deitch, & Farell, 2010).

One of the largest examples was the experimental evaluation of the Job Training Partnership Act (JTPA) between 1987 and 1992. Over 20,000 potential job training participants were randomly assigned to either receive training or serve as a control group. BC analyses of JTPA and other training programs looked at the costs of the program and compared these to the benefits in terms of earnings and employment subsidies net of lost welfare payments for participants (see the framework in Chapter 9 adapted from Orr, 1999). For JTPA, a comparison across treatment and control groups showed the impact of training was earnings gains of 19% to 21% overall, although these gains were concentrated in the adult sample and varied across the quantiles of trainees' earnings (see Abadie et al., 2002, Tables II and III; Orr, 1999; and Orr et al., 1996). Other benefits were generated because individuals required fewer government-provided services (e.g., public assistance and criminal justice involvement) as a result of their job training. These benefits were received by taxpayers and other members of society. Interestingly, not only did the benefits of JTPA vary across groups but the program costs did also: From a social perspective, the benefits for adult men were $2,700 and the costs $1,640; for male youth, the benefits were –$1,940 and the costs were $3,910 (Bloom et al., 1997, Table 8, adjusted to 2015 dollars).

We might expect welfare participants would respond to education programs to improve their labor market prospects. Greenberg, Deitch, and Hamilton (2009) summarized results of BC analyses from 10 welfare-to-work programs that relied on educational investments. Five were "education-first" programs: Participants were provided with educational supports to help them obtain credentials (e.g., GEDs) and so become more successful in the labor market. The other five were "mixed-initial-activity" programs: Participants were either assigned to educational supports or job search assistance depending on need. The education-first programs typically failed the BC test (Greenberg et al., 2009, Table 5.1). From the individual trainees' perspectives, three of the five programs had negative benefits (with zero costs). From the government perspective, four programs had negative present values (with BC ratios of less than 1). Finally, from the social perspective, all five education-first programs had negative present values and BC ratios less than 1. As a contrast, mixed-initial-activity programs showed more positive BC results (Greenberg et al., 2009, Table 6.1). Only one program site yielded losses from the each of the participant, government, or social perspectives. From the social perspective, the BC ratio ranged from 0.7 to 3.79.

Overall, welfare-to-work programs that rely primarily on educational supports appear to have a negative return with costs greater than their benefits. When targeted based on need or income at baseline, the returns are positive but not always. Thus, it is not always the case that educational investments have positive NPVs. Of course, we have not accounted for any justification of investments for these populations on equity grounds.

10.4. CONCLUSIONS

In this chapter, we have presented the basic structure of BC analysis and how to present the results using three economic metrics. This presentation reflects the fact that BC analysis is a numerical method. Although there are different ways to estimate costs and shadow price benefits, costs and benefits are combined in a specific way. With respect to costs, it is expected that all resources are valued even if a school or district does not directly pay any money for their use. For benefits, the analyst applies quantitative measures of educational outcomes and models of these for as long as they persist. The consequence of these prescriptions is that the metrics for comparing policies are simple and easy to understand; thus, the appropriate conclusions can be drawn. The disadvantage is that some aspects cannot be incorporated into the BC analysis: Any outcomes that cannot be shadow priced, for example, are excluded; issues of fairness are not addressed.

Regardless, the number of BC studies in education is growing, in part because of the increase in random assignment methods for evaluating effectiveness (and the use of earnings as measures of benefits).[1] There is also interest from government departments, state legislatures, and nonprofit agencies in identifying high-return social programs (for an inventory of BC studies for Washington State, see Lee et al., 2012; for a broader discussion, see Revesz and Livermore, 2008). Yet, there are still many fields in education policy to which BC analysis has not been widely applied—for example, value-added modeling, continuing education programs, education for select groups such as ex-offenders or veterans, and parental education programs. Presently, BC analysis is a long way from being a default approach to economic evaluation.

Thus far, the evidence suggests that—in many but not all cases—the NPV from educational investments is positive. Large labor market effects, as well as broad and persistent effects on behavior, translate into large benefits when programs are effective. Indeed, some educational programs may be efficient even if they do not improve the educational standing of the participants; for example, Papay and Johnson (2012) described a program that helps teachers improve their practice but that primarily yields cost savings. Ultimately, though, education is an investment for future benefits and so we cannot be certain that each investment will yield a positive NPV. Also, even if the expected NPV is positive, it may be that the estimate cannot be precisely bounded above zero. In the next chapter, we consider how to address this uncertainty.

[1] For an investigation of the IRR for the Japanese ronin examination process, see Ono (2007). For a baseline calculation of the IRR for underperforming teenagers in Israel, see Lavy and Schlosser (2005). For a BC of financial incentives in community colleges, see Barrow, Schanzenbach, and Claessens (2015). For a BC analysis that reduces child maltreatment, see Maher, Corwin, Hodnett, and Faulk (2012).

Discussion Questions

1. How would you interpret each of the following metrics: NPV, BC ratio, and IRR? Under what circumstances would you use each metric, and why?
2. What types of educational interventions lend themselves best to BC analyses? Why?
3. Under what conditions might BC analysis be preferable to CE analysis for evaluating educational interventions?
4. Why might a decisionmaker reject an educational evaluation with a positive NPV?

Exercises

1. The Chicago Child-Parent Center program serves approximately 3,000 3- to 5-year-olds from low-income families in over 20 centers across Chicago. Students receive 3 hours of instruction per day, 5 days per week during the school year and for 6 weeks over the summer. The program featured structured math and reading activities led by certified teachers and aides in small classes, with a substantial parent component and supplemental health and nutrition services. Reynolds et al. (2002) estimated the per-student costs, present value at age 3 with a 3% discount rate, at $8,510, adjusted to 2007 dollars. Adapted from Reynolds et al. (2011), selected benefits of the program are estimated as follows:

	Present Value Benefits/Savings (2007 dollars)
Child care for families	$4,387
Child abuse/neglect public services	$7,330
Grade retention school services	$880
Special education school services	$5,317
Earnings in adulthood	$28,844

(Continued)

(Continued)

	Present Value Benefits/Savings (2007 dollars)
Criminal activity expenditures	$42,462
Health expenditures	$3,294
College tuition	−$294
Total Benefits	$92,220

a. Calculate the social NPV and BC ratio for the program.
b. Divide the benefits into private benefits and fiscal benefits. What is the private NPV and BC ratio and the fiscal NPV and BC ratio?
c. Under what circumstances would distributional weights change your recommendation on implementing the program?

2. The HighScope Perry Preschool Program was a high-quality preschool education to 3- and 4-year-old African American children living in poverty and at high risk of low academic performance. The program involved an active learning curriculum and weekly 1.5-hour home visit to each mother and child, designed to involve the mother in the educational process and help implement the preschool curriculum at home. To test the benefits of the program, 123 students were randomly assigned either to the control group or preschool treatment group. These 123 students were followed up at ages 19, 27, and 40. The present value cost of the program at age 4 was $15,000 (in 2000 dollars). Adapted from Nores et al. (2006), selected results were as follows:

	HighScope Perry Preschool Group	Control Group
Earnings at age 27 (2000 dollars)	$13,328	$11,186
Earnings at age 40 (2000 dollars)	$24,466	$19,699
Felony total count ages 19 to 27	0.12	0.26
Felony total count ages 28 to 40	0.05	0.07
Misdemeanor total count ages 19 to 27	0.12	0.22

	HighScope Perry Preschool Group	Control Group
Misdemeanor total count ages 28 to 40	0.17	0.25
Months on welfare entire ages 19 to 27	4.74	4.71
Months on welfare entire ages 28 to 40	2.15	2.02
HS graduate by age 19	67%	46%

Perform a social BC analysis of this program. You may find the following helpful: shadow prices for social willingness to pay for crimes of $47,000 per felony, $7,200 per misdemeanor (McCollister, French, & Fang, 2010), and welfare payments of $400 per month.

3. A school district is looking to boost high school graduation and college completion rates. It has a choice of two summer programs for students between eighth and ninth grade. Without the programs, the district has a high school graduation rate of 40% and a college completion rate of 20%. Program A is delivered to 1,000 students and costs $2,400 per student; for these students, the high school and college completion rate is 66%. Program B is delivered to 500 students and costs $400 for each year per student but only when the student is enrolled in either school or college; for these students, the high school graduation rate is 90%, and the college completion rate is 60%. As an economist, you calculate the lifetime earnings of high school dropouts at $300,000, of high school graduates at $400,000 and of college graduates at $600,000 (present values at age 20 using a 3% discount rate). Assume there are only these three possible educational states. Perform BC analysis of the two programs relative to the status quo.