

8 Steps in a cost-benefit analysis



A cost-benefit analysis of an early childhood intervention involves three steps: (1) an estimation of the net impact of the intervention; (2) an estimation of the social costs and the social benefits of the intervention in monetary terms; and (3) a calculation of the cost-benefit of the intervention.

Data for steps 1 and 2 can be collected from repeat surveys and linkages to administrative records of program participants. It is important that only the *net* costs and benefits are counted. In the case of costs, this is the costs incurred by society in the program scenario which are over and above the costs of the counterfactual scenario (usually a no program situation). This applies to net benefits as well. For example, if a program leads to the completion of Year 11 at secondary school for participants compared with Year 10 for non-participants, then only the value of the extra year needs to be monetised.

Pecuniary values for non-pecuniary costs and benefits (crime, loss of health, unemployment) are usually derived from secondary literature which has made these estimates. Most intervention costs are pecuniary in nature and include spending by government on resources such as buildings, equipment and facilities, and for the wages of social, health and education workers. Usually the non-pecuniary costs of a program, such as the loss of time incurred by the child and his or her family as a result of participating in the intervention are not counted. This is usually because evaluators either do not believe they are large or believe they are benefits, not costs. For prudential reasons, most evaluations err on the side of understating net benefits, especially when they have a speculative component, and consequently these costs are ignored.

By contrast, many net benefits of early childhood interventions are non-pecuniary. The obvious examples are reductions in the crime rate, but other benefits include the increased satisfaction associated with the schooling years, eventual employment and family functioning. Most evaluations do not seek to monetise changes to life satisfaction but they do enumerate the effects of reduced crime on prospective victims. Again, the dominant reason for ignoring the net changes to satisfaction reside in the conservative nature of cost-benefit analyses on the one hand and the more speculative conversion of some benefits into money equivalents on the other.

Many of the conjectural estimates in a cost-benefit analysis arise from making projections of pecuniary benefits, not from monetising non-pecuniary benefits. With respect to early childhood interventions, this usually involves making projections of, for example, how improvement in school retention will convert into more wages for the participant in adult life. Because the anticipated benefits from these interventions can have very long horizons, policy makers who wish to discern whether an intervention is making a positive social contribution before the full extent of benefits are “known”, have to rely on linking intermediate program net impacts with findings from other studies. In the case of the increment to labour incomes, this can mean linking changes in primary school achievement (however measured), with changes to subsequent educational attainment, occupational attainment and associated wages. Using the law of statistical average, the evaluator can then project from the age of five through to 65 years. However, the longer the projection period the greater the band of errors around the estimate. For example, we may say that a program leads to an increase in the present value of wages of say \$5000 ± 10 per cent up to the age of 30, but \$30,000 ± 25 per cent up to the age of 65.¹⁰

10 Karoly et al. (2001) provide a very good summary of the practical matters to be considered before attempting a cost-benefit analysis.

Step 1: Estimating the net impact of the intervention

The first step in conducting a cost-benefit analysis requires answers to two questions. First, what outcomes have the participants achieved in comparison to what they would have achieved if they had not participated in the intervention (called the *net impact* of the intervention)? Second, if this intervention is extended to other children, will it have the same net impact?

Tests of significance are a measure of the confidence we have in the size of the estimated net impact. Given a large representative sample, tests of significance will tell you whether another group of children, randomly chosen from the population, are likely to incur the same effects from undertaking the intervention. However, a test of significance is different from the actual size of the net impact. An estimate of a net impact, while significant, may be very small in absolute terms. On the other hand, the net impact may be very large, but insignificant (referred to as effect sizes in the previous sections). In the latter case, this is usually because there was a large variation in the net impacts of individuals in the intervention group, perhaps because relevant co-variables have not been modelled.

To obtain information on the net impact of an intervention requires either the selection of a control group who are similar in every respect except for their inclusion in the intervention; or a large enough sample that allows us to generalise the effects and abstract from random external factors that may have an effect on an individual's performance.

These issues are discussed in detail below.

Selection of a control group

The difference in outcomes for the intervention and control groups that can be ascribed to the intervention itself is called the net impact.¹¹

The ideal way to avoid the possibility that particular types of parents (for example, the more motivated or persistent parents) enrol for the intervention is to randomly assign children to the intervention and control groups before the intervention commences. This requires the evaluation to be designed as an integral part of the intervention. By using the law of large numbers, random assignment ensures that there is no systematic tendency for either group to have more or less favourable characteristics, either observable or unobservable.

However, random assignment does not totally eliminate systematic differences between the two groups, even if large samples are involved. Because it is not possible to force families to participate in an intervention, some self-selection out of the intervention by families who disapprove or do not see value in the program, or who because of contemporaneous difficulties in their life, are not in a position to offer their time to the program, is expected. Participation rates as a threat to evaluation design has been discussed in detail in Section 4.2. This effect will tend to overstate the net impact of the intervention.

In cases where random assignment is not possible, the evaluator should try to construct a control group from children who match the intervention group on relevant observable characteristics. Usually, this means characteristics such as parents' socio-economic background, ethnicity, pre-program IQ and so on. If pertinent unobservables, such as parental motivation, could be assessed, then these could be used for matching as well. Choosing which characteristics to match on is usually informed by existing literature. In cases where the intervention is large and pervasive, such as universal maternal and child health or preschool programs, it may be difficult to find a population of children who have not received any intervention services.

11 It is also possible to control for unobservable characteristics using the difference-in-difference method. However, this method requires measures of outcomes before the intervention. This can be reasonable for adult interventions. For example, in employment programs, measures are taken of the outcome variable, employment status, both before and after the program. It is less likely to be suitable for child-orientated interventions where the main outcomes are not observable before the intervention has begun.

Random assignment and matching methods do not in themselves ensure that, on average, the non-program characteristics of the two groups are the same for any selected sample. If the samples are small, then the chance that any one pair of program and control groups will not be similar is high. Evaluators will often control for other factors in analyses of intervention effects (such as multiple regression), even when random assignment or *ex ante* matching has been undertaken, to iron out any remaining differences between the two groups.

Generalisability

The second question is whether the same net impacts will occur if the intervention is extended to other groups of children. This requires the evaluator to discuss the sub-groups that can potentially benefit from the intervention, and to assess whether the evaluation results can be generalised to these other groups. For example, an evaluation may only consider children from low-income refugee families or mothers, and the same net impact may or may not accrue if the intervention was applied to low income indigenous families or fathers.

Step 2: Estimating the social costs and benefits of the intervention in monetary terms

The second step in conducting a cost-benefit analysis involves an estimation of the social costs and social benefits of the intervention in monetary terms.

Estimating benefits

There are two philosophical considerations that should be made explicit with regard to estimating benefits of an intervention. The first relates to how extensively the measured benefits of a intervention are defined, or how many people who are indirectly affected by the intervention are counted. The second consideration is whether the outcome is real, or merely an intermediate result which is valued for its potential to affect “real” outcomes.

It is possible to argue that most people in a society can be affected by an intervention, but clearly a line must be drawn, otherwise the evaluator will spend excessive time making calculations of secondary and uncertain detail.

Determining at what point a result, especially an ephemeral one, is a benefit in itself or an indicator of a potential benefit is a rather more difficult question to answer. For example, is higher educational attainment valued in itself, or only as an indicator that the person will go on to find more rewarding and stable employment? Further, if benefits are not sustained over time, are these early benefits “real” or just promises that amounted to nothing? There may be a case for arguing that an intervention that only affects educational outcomes for a few years post-intervention, has zero benefits since the higher attainment in these years may have had no long term effect on the happiness of children, or their families, who participated in the intervention.

Without a clear notion of what defines a benefit for the participant, evaluations are reduced to measuring the effects on broader society. If only higher educational attainment is valued because it leads to a higher paying job, then it is not appropriate to count higher net years of education as a benefit as well as the additional wages. If however, more schooling is considered to have intrinsic value to the child, then some estimate should be made of this additional value at the time that the education was received.

In the cost-benefit evaluations considered later in Section 11 the final chosen outcomes are usually the impact on future wages and crime rates. If this is what the program designers intended, then it is appropriate for the evaluator to concentrate on recording these benefits. However, intermediate milestones, such as educational attainment, are worth measuring if it is intended to project final outcomes before enough time has elapsed to allow the full extent of the final outcomes to be apparent.

These decisions over the scope of the population and the designation of ultimate benefits should be decided before the evaluation is designed. Not only does it determine what sort of benefits to measure, but also the length of time the children should be followed.

The ideal length of time allowed for data collection is determined by the balance between the costs of data collection and the attrition of participants from the evaluation study. In practice, evaluators typically have to work with short- and medium-term data, using projections to complete the cost-benefit analysis. This means drawing on existing studies on the relationship between anticipated benefits such as health and educational attainment, and longer-term outcomes such as earnings. These reduce confidence in the final result, but it is an acceptable method and possibly the only way to conduct a cost-benefit analysis in years immediately following an intervention.

Scope of included benefits

It is not necessary to know at the commencement of the intervention all types of potential benefit, as survey items can be decided at a later stage. If however, the intervention expects a change in the families' status (such as a change to mothers' education), then this does have to be considered *ex ante*. It is easy for older interventions that have extensive data on adult experiences to estimate final pecuniary benefits from an intervention but considerably harder for interventions that are more recent (for example, Bolivia PIDI) and only have intermediate outcomes, such as cognitive development and need for remedial schooling. Interpreting the results from these newer interventions requires a more flexible and sophisticated approach.

Scope of potential beneficiaries

Earlier studies only considered monetising the effects on the child, victims of crime and the taxpayer (for example, savings in remedial education, higher earnings and lowered crime rates). More recent studies have also considered the broader effects on the mother (for example, better health and less substance use). In most cases, a qualitative discussion is made of the effects on people who are likely to be affected. Groups for whom quantitative measures are required generally need to be known *ex ante*.

Estimating costs

Scope of included costs

Deciding which cost to include is usually the least controversial part of a cost-benefit analysis, and is usually limited to the direct government costs of running the intervention. Recall that all costs are benefits foregone, or rather, represent the money value of benefits society would have otherwise gained had the government spent the money on another program. This may be another type of child related intervention, a welfare program or simply reduced taxes. Theoretically, changes incurred by staff in the counterfactual program should also be considered, but this is generally too much detail for most evaluations.

With respect to early childhood interventions, most of the resources used are the labour services of professional or paraprofessional staff who deliver health, psychological and educational services to young children and their families.

Some evaluations attempt to distinguish between fixed and variable costs, but it is unclear how valuable this information is to the overall study, relative to the cost of collating it, since most fixed costs are only fixed over limited ranges of production. In-kind resources, such as free rent or facilities, should be included as well as budgeted items. Often, no account is taken of the costs to families, as it is often assumed that the time they provide has no opportunity cost. This may or may not be true. The skill in measuring costs is usually to avoid double counting. Karoly, Kilburn, Bigelow, Caulkins, Cannon and Chiesa (2001) provide further details of hazards to watch for when measuring costs.

Fixed costs associated with the establishment of the intervention are only relevant if they will be incurred every time the intervention is extended. Fixed costs associated with the initial design of the intervention that are one-off are not relevant.

Most costs are part of the funding costs and are already monetised. However, in the case where idle resources are used (such as unemployed professionals), the costs of employing the resource will not be true costs and should be excluded. In cases where there are indirect costs, such as a cost of the families' time and costs to the schools that are not directly funded, then some estimate of these should also be made, but only if it is believed they are significant. If they only represent minor costs, then they should be just referred to in the text.

It is rare for the costs of an intervention to extend beyond the intervention period. If they occur, they will in most cases be revealed as negative benefits and will be monitored through the benefits section.

Step 3: Calculating the cost-benefit

Cost-benefit analyses produce a single metric, which is a summary measure of the difference between the costs and benefits of an intervention. In most cases, however, a sensitivity analysis is done, and this can produce a range of figures for one intervention. A sensitivity analysis involves inputting different assumptions about, for example, the value of crime, or the rate of discount, to see how much the overall figure changes.

Net present value (NPV) is the preferred method to calculate the net impact of an intervention (see Section 9 for a full discussion of cost-benefit calculations), although many studies convert this to the ratio of net present benefits to net present costs. Present value (PV) calculations can be calculated at any time after the intervention has ended, although (as noted above) the earlier the evaluation the less clear the results and the more conjectures are needed about how early indicators map into later outcomes. It is advisable to calculate net present values for several rates of discount (rates of interest) to make the sensitivity of the result to variation in the rates apparent.