

(SOCIAL) COST-BENEFIT ANALYSIS IN A NUTSHELL

RUFUS POLLOCK

EMMANUEL COLLEGE, UNIVERSITY OF CAMBRIDGE

1. INTRODUCTION

Cost-benefit analysis is a process for evaluating the merits of a particular project or course of action in a systematic and rigorous way. *Social* cost-benefit analysis refers to cases where the project has a broad impact across society – and, as such, is usually carried out by the government.

While the cost and benefits may relate to goods and services that have a simple and transparent measure in a convenient unit (e.g. their price in money), this is frequently not so, especially in the social case. It should therefore be emphasized that the costs and benefits considered by (social) ‘cost-benefit’ analysis are not limited to easily quantifiable changes in material goods, but should be construed in their widest sense, measuring changes in individual ‘utility’ and total ‘social welfare’ (though economists frequently express those measures in money-metric terms).

In its essence cost-benefit analysis is extremely, indeed trivially, simple: evaluate costs C and benefits B for the project under consideration and proceed with it if, and only if, benefits match or exceed the costs. That is:

$$\underbrace{B}_{\text{Benefits}} \geq \underbrace{C}_{\text{Costs}} ?$$

So what makes things more complex? There are a variety of factors:

- Benefits and costs may accrue to different sets of people. If this is so we need some way to aggregate and compare different benefits and costs across people.
- Benefits and costs may occur at different points in time. In this case we need to compare the value of outcomes at different points in time.

Email: rufus [at] rufus [dot] pollock [dot] org. This paper is licensed under Creative Commons attribution (by) license v3.

- Benefits and costs may relate to different types of goods and it may be difficult to compare their relative values. This usually occurs when one of the goods does not have an obvious and agreed upon price. For example, we may be spending standard capital goods today in order to obtain environmental benefits tomorrow.
- Benefits and costs may be uncertain.
- Benefits and costs may be difficult to calculate and, as a result, there may be widely differing views about their sizes. One might think this could be subsumed under uncertainty, however the two points are rather different: two people agreeing that an outcome follows some probability distribution is different from them arguing about its mean and variance.

Usually, in real-world cases the dominant issue is usually the last one: the basic job of calculating estimates for the project's costs and benefits. This especially true in the 'social' case where the projects under consideration may involve costs and benefits that very difficult to quantify – what is the benefit of the national security derived from military spending, how large are the benefits from education, etc etc. Necessarily this quantification only makes sense on a case-by-case basis. Here we are concerned with general principles and we therefore focus only on the preceding four items and look at how they can be incorporated into the analysis in a general way.

2. THE BASIC MODEL

We are considering a project with known (though perhaps uncertain) benefit B and cost C . Our task is to decide whether it is worthwhile. As already discussed, if B and C were denominated in exactly the same terms (i.e. the same good, at the same time) for a single person and with no uncertainty things would be straightforward: check whether $B \geq C$. However, this is unlikely to be the case so we will need to do more work. All of this work, in its essence involves converting benefits and costs into values that can be easily compared. Equivalently we need to have benefits and costs denominated in terms of some standard good or measure of value. We shall term this good or measure of value the *numeraire*.

In theory, this numeraire could be anything: apples, years of life, acres of rainforest etc. However, given that many (though by no means all) goods are already denominated in

terms of money, it is often natural to use a numeraire that is money-metric. We also need to specify money in whose hands (for example, £1 in the hands of someone on breadline is likely of different value to £1 in the hands of a billionaire).

For the purposes of *social* cost benefit analysis a very natural numeraire is ‘(uncommitted) government funds’, that is money the government has but is not yet allocated to any given project. This is a natural numeraire since it is likely to be government funds which are used in paying for the project being considered.

We will also assume, to make our lives straightforward, that, unless specified otherwise, the cost of a given project is exactly one unit of government funds today. This allows us to focus only on the benefits which makes things simpler while sacrificing no real generality.

We begin, in the section that follows, by focusing solely on the distributional issues and ignoring temporal and uncertainty. We then introduce temporal considerations and discounting and conclude by discussing uncertainty.

3. DISTRIBUTIONAL CONCERNS

Let us suppose there are N people or groups. The benefit to group i from the project under consideration is b_i of *income/consumption*.¹ Income/consumption is construed broadly here to include not only normal traded goods such as food or digital music players but also things like security, a stable biosphere etc. Individuals existing income/consumption is denoted by x_i . Those operating the project (the government) have a utilitarian welfare function:²

$$W = \sum_i u(x_i)$$

The change in welfare (assuming away any changes in behaviour) arising from the individual benefits is:

$$\Delta W = \sum_i u(x_i + b_i) - u(x_i)$$

If the gains are small relative to existing income we may approximate the change in individual welfare using the derivative: $u(x_i + b_i) - u(x_i) = u'(x_i)b_i$. Defining, $w_i = u'(x_i)$ we then have a set of ‘distributional weights’, that is weightings for individuals such that

¹We won’t distinguish between income and consumption here since it will not matter for our purposes.

²We could avoid any reference to individual utility here by positing a social welfare function defined directly in terms of the outcomes being affected be that money, education, security etc.

the total (welfare) benefit is just the sum of the weights times the individual benefits:

$$\Delta W = \sum_i w_i b_i$$

One last step remains: we need to convert utility back into our numeraire (government funds) via multiplication by some constant θ – the overall benefit B will then be $\theta\Delta W$. To specify this constant we choose a benchmark project and then define its benefit B to exactly one unit of the numeraire – since costs are also 1 this implies this project is just worthwhile. The standard approach is for the benchmark to be a project which generates benefits equivalent to one unit equally divided equally among all groups, i.e. $b_i = 1/N$.³ This implies that:

$$\theta \sum_i w_i \frac{1}{N} \equiv 1 \Rightarrow \theta = \frac{N}{\sum_i w_i}$$

Note that if income were already equally distributed so $x_i = x$ and utility (which is only defined up to a constant) were normalized so that that the marginal utility of income at the reference income x were exactly one we would have $w_i = w = 1 \Rightarrow \theta = 1$. To summarize:

$$\begin{aligned} N &= \text{Number of beneficiaries} \\ b_i &= \text{Benefits to group } i \\ x_i &= \text{Income of group } i \\ w_i &= \text{Weights} = \text{Marginal Utility of group } i \\ \Delta W &= \text{Welfare benefit} = \sum_i w_i b_i \\ \theta &= \text{Conversion factor from SW to Numeraire} \\ B &= \text{Benefit} = \theta \sum_i w_i b_i \end{aligned}$$

3.1. Calculating the Conversion Factor θ . With a few assumptions on the form of the utility function and knowledge of the distribution of income we can calculate an actual

³Clearly the choice of the benchmark project is important. Why then choose this project? The answer is that £equally distributed is what standard government projects like provision of free education or free healthcare actually amount to. As such such this ‘equal distribution’ project is definitely included in the government portfolio and furthermore no worse project should be worthwhile because we could reallocate from that project to the ‘equal distribution’ and improve well-being.

figure for θ . Assume that utility takes CES form:

$$u(x) = \frac{x^{1-\alpha} - 1}{1 - \alpha}$$

Thus, the weights (equal to marginal utility) are $w_i = x_i^{-\alpha}$ and hence $\theta^{-1} = E(x^{-\alpha})$.

At this point, it is useful to move to continuous rather than discrete variables so:

$$\frac{1}{\theta} = \int w_i(x) dF(x) = \int x^{-\alpha} dF(x) = E(x^{-\alpha})$$

Now, the distribution of income x is (approximately) log-normal $LN(\nu, \sigma)$ in which case using the formula for the moment generating function of the normal we have:

$$E(x^{-\alpha}) = e^{-\alpha\mu + \frac{\alpha^2\sigma^2}{2}}$$

3.2. Example 1: Benefits in Proportion to Income. Suppose the project generates value V which is then distributed in proportion to income. Let λ be the ratio of individual benefit to income so $b_i = \lambda x_i$. Now $\sum_i b_i = V \Rightarrow \lambda = V / \sum_i x_i$. Thus using our formula from above the total benefit in terms of the numeraire is:

$$\begin{aligned} B &= \theta \sum_i w_i b_i \\ &= \theta \lambda \sum_i w_i x_i \\ &= \theta \frac{NV}{\sum_i x_i} \frac{1}{N} \sum_i w_i x_i \end{aligned}$$

Using a CES utility function so $w_i x_i = x_i^{1-\alpha}$ and using expectations we have:

$$B = \frac{VE(x^{1-\alpha})}{E(x^{-\alpha})E(x)}$$

Using a log-normal distribution for income and the expression for the MGF as before this further reduces to $B = e^{-\alpha\sigma^2}$. For log-utility, $\alpha = 1$, and a reasonable estimate of σ is 0.47 (Newbery 2008) which implies $B = 0.8V$. Thus each pound/euro/dollar distributed generates a benefit in terms of the numeraire of 0.8 and, if the project is to be worthwhile, it must have a direct yield of at least 25% ($= 1/0.8$).

3.3. Discounting. We now come to the time factor: benefits of effort or expenditure today may not be realised until tomorrow. In the spirit of keeping things simple let

us fix everything about the problem except the temporal aspect. In particular, ignore distributional issues, uncertainty and variations in the types of goods involved.

Assume that by giving up one unit of expenditure today we gain V units at time T . Let expenditure today be x and at time T (in the absence of the project) x_T . There is a utility function u which converts expenditure into *contemporaneous* utility (i.e. utility in that period). Let the numeraire be present period utility normalized so that the marginal utility today equals 1, i.e. $u'(x) = 1$.

Thus there are two major factors to take into account. First, how to convert utility from period T to now. Humans tend to prefer things sooner rather than later. Hence, even with all else equal, utility today is preferred to utility tomorrow. The measure of this is termed the level of ‘pure time-preference’ and we will denote it by $\rho(t)$. Second, the reference situation tomorrow (in terms of resources, consumption etc) may not be the same as today and marginal utility from gaining or losing a unit will differ across time, quite apart from time-preference.

Assuming changes in expenditure are relatively small we can approximate utility changes using derivatives we have:

$$C = \text{Cost} = -u'(x) = -1$$

$$b = \text{Benefit at } T = Vu'(x_T)$$

$$B = \text{Benefit in today's utility} = \rho(T)b$$

So the project is worthwhile if:

$$V\rho(T)u'(x_T) \geq 1$$

This implicitly defines a discount factor $\delta(T) = \rho(T)u'(x_T)$ with the payoff of V at time T valued at $\delta(T)V$ today.

3.4. Example 2: Discount Rates and Climate Change. Suppose we can spend resources (or equivalently forgo consumption) today to mitigate the effects of climate in the form of benefits relative to the do-nothing scenario, at some point T in the future. Suppose, in the absence of this project, the economy would grow at rate g per year so

consumption at time T is e^{gT} times consumption today. Suppose pure time preference takes an exponential form so $\rho(T) = e^{-\rho T}$ and we have CES utility as before. Then:

$$B = V e^{-\rho T} e^{-\alpha g T} = V e^{-(\rho + \alpha g)T}$$

Comparing this with a standard exponential discount rate $e^{-\delta T}$ gives an implied discount rate $\delta = \rho + \alpha g$.⁴

3.5. Uncertainty. The models discussed above involve no uncertainty: all relevant values, e.g. the project's payoff, future consumption levels etc, are known with complete certainty. This is clearly unrealistic and it is useful to be able to consider situations which do involve (known) uncertainty.

The natural approach here is simply to replace costs and benefits with their expected values. If those receiving benefits (or bearing costs) are risk averse – as is likely – uncertainty will act to reduce benefits and increase costs (the certainty of £1m is worth more to most people than an evens chance of 0 or £2m). We illustrate by returning to our climate change example.

3.6. Example 3: Uncertainty and Climate Change. Consider our previous example regarding climate change. We will incorporate uncertainty a little indirectly here by assuming that direct costs and benefits remain certain but that we are uncertain as to the growth rate g in the absence of action (i.e. what the effects of climate change will if we do nothing).⁵

Uncertainty in the growth rate will affect our calculation of benefits because it will alter the level of consumption, and hence marginal utility, at which benefits in the future are evaluated. Remember, the benefits of mitigation are greater the more terrible the world when we do nothing (i.e. in the case of unmitigated climate change).

Formally, suppose the growth rate g follows some probability distribution given by G . Then:

⁴Note that this is, not coincidentally, the same as the real interest rate found in the Ramsey Kass Coopmans model.

⁵An alternative formulation of the climate change question would be to change the default scenario to full mitigation and the action scenario being do nothing. In this framework the benefit would be gained 'consumption' today while the cost would be the reduction in 'consumption' in the future. This cost would then be directly related to the growth rate in our main formulation of this problem.

$$C = -E(u'(x)) = -1$$

$$B = \rho(T)VE(u'(e^{gT}x))$$

$E(u'(e^{gT}x))$ is expected marginal utility at time T . With risk aversion this will be increasing in ‘uncertainty’ (e.g. a increase in the variance of the growth rate that preserved the mean). To say more we need to specify a functional form for u and G . Take u as our previous CES form and suppose growth is normally distributed: $g \sim N(\mu, \sigma^2)$. Then, recalling that the moment generating function for X , a normal random variable, is $E(e^{tX}) = e^{t\mu + \frac{1}{2}t^2\sigma^2}$, we have that:

$$E(u'(e^{gT}x)) = E(e^{-\alpha gT})E(x^{-\alpha}) = E(e^{-\alpha gT})E(u'(x)) = e^{-\alpha\mu T + \frac{1}{2}\alpha^2 T^2 \sigma^2}$$

Thus under uncertainty the discount rate is $\rho + \alpha\mu - \frac{1}{2}\alpha^2\sigma^2$. Recall that under certainty it was $\rho + \alpha g$. Taking the case where the growth rate under certainty, g , is equal to the expected growth rate under uncertainty μ , we see that the **impact of uncertainty (in the form of variance in the growth rate) acts to reduce the discount rate.**

To illustrate, take as a benchmark case $\rho = 0.5, \alpha = 1.0, g/\mu = 2.5$ so the discount rate is 3% in the case of certainty over the growth rate. Suppose however we aren’t sure how bad (unmitigated) climate change will be so we replace our certain growth rate of 2.5% with a normally distributed one with the same mean but a standard deviation (σ) of 2%. This would reduce the discount rate to 1%. Moreover, a standard deviation greater than 2.5% would result in an *negative* discount rate – in expectation, consumption in the future is more valuable than consumption today (even though on average we expect consumption to grow at 2.5% per year)!

The crucial point here is that uncertainty acts asymmetrically because of diminishing returns (equivalently risk aversion) in the utility function. Increasing uncertainty in the growth rate increases the chances *both* of being in a ‘great world’ where we are (relatively) well-off *and* of being in a ‘terrible world’ where we are (relatively) poor. In the former case our gain from extra consumption in the future (the result of climate change mitigation in this model) is reduced because we are already in a good situation but in the latter case the

gains are greatly increased because we are so badly off. Moreover, thanks to diminishing returns to consumption the gains in the bad case greatly outweigh the reduction in benefits in the good case. As a result increasing uncertainty increases the value of each unit of benefit in the future and thus reduces the discount rate.