



MANY-SIDED DICE

THE GOVERNMENT USE THEM
TO MAKE DIFFICULT DECISIONS

18 Cost-Effectiveness and Policy Choice

The ultimate purpose of economics, of course, is to understand and promote the enhancement of wellbeing.

Ben Bernanke, former Chairman, US Federal Reserve

The ultimate purpose of wellbeing science is to help us increase wellbeing. Hopefully, readers of this book will by now have learned a little more about themselves, which they can use to improve their own wellbeing and that of others. But what about policy-makers, be they in central or local government, or in NGOs big and small? Are there steps by which wellbeing science could help them improve their contribution to human wellbeing?¹

In this chapter we ask the following questions

- How would policy-makers spend their money if they wanted to maximise wellbeing? How can we measure cost-effectiveness when benefits are measured in units of wellbeing?
- How does this approach compare with traditional cost-benefit analysis where benefits are measured in units of money? Which approach is better?
- If policy-makers wanted especially to reduce misery, how would they proceed?
- How far are central policy-makers using the new approach?

The Goal

First, policy-makers would need to be clear that wellbeing is their overarching goal. At present most organisations have **multiple goals**. But when they decide how much effort to devote to each of these goals, they are implicitly balancing one goal against another. Ultimately there can be no rationale for such decisions unless it is based on some overarching goal, against which the importance of the different subsidiary goals can be assessed.

¹ This chapter draws heavily on A. E. Clark et al. (2018) chapter 15. For a formal statement of the argument of this chapter, see Annex 18.1.

In Part I, we already set out the argument for wellbeing as the overarching goal. So, in this vision every organisation would be contributing in whatever way it could to the goal of maximising the **sum of future WELLBYs** (suitably discounted).²

For example, every nation's finance minister would say to each cabinet colleague 'When you make the case for your department's budget, please estimate how much each of your main proposed expenditures (new and old) will increase the wellbeing of the community. Tell us their effect on wellbeing and how much they will cost'. And the leaders of any organisation, large or small, would ask the same questions of its different branches.

Cost-Effectiveness Analysis

The decision on these proposals would then proceed as follows.³ Realistically, we can assume that any typical public organisation has an overall budget constraint – the total amount of overall expenditure it can incur on all its policies. Thus, for each policy the key issue is how much wellbeing it produces per dollar of cost – **the cost-effectiveness of the policy**. So policies would be ranked according to the (discounted) WELLBYs they will generate per dollar of net (discounted) expenditure.

Once they are ranked in order of their cost-effectiveness, we would choose those that were the most cost-effective until the available budget was exhausted. There would thus be a **cut-off level** (λ) of cost-effectiveness above which policies were approved and below which they were rejected. Thus, the criterion for approving policy would be⁴

$$\frac{\sum \sum W_{it}(1 - \delta)^t}{\sum C_t(1 - \delta)^t} > \text{Cut - off value} = \lambda \quad (1)$$

where C_t is the net cost in year t .

The appropriate **critical value** (λ) could be found by trial and error. An alternative approach would be to start with values of λ already used in more limited areas of life. For example, the WELLBY approach is similar in many ways to the approach already followed in some healthcare systems. As we have noted in Chapter 10, England's National Institute of Health and Care Excellence (NICE) evaluates a proposed treatment according to the number of QALYs (Quality Adjusted Life Years) that it produces relative to its cost. A treatment is only approved if the ratio of cost to benefit is low enough. Currently, NICE require the cost per QALY to be below around \$40,000. But this applies only to health expenditures and to the health-related benefits to the individuals who are treated in the healthcare system. The WELLBY approach

² Assuming for the moment that we ignore who gains and who loses.

³ For an early discussion of this whole approach, see O'Donnell et al. (2013). For more recent applications, see De Neve et al. (2020); and Frijters and Krekel (2021).

⁴ The same δ is used for wellbeing and for cost, on the assumption that the rate at which expenditure produces wellbeing is constant over time.

would relate to all government expenditure and to the wellbeing effects on all the individuals affected. It would be the standard way in which all public expenditures are decided.

Relation to traditional cost-benefit analysis

A natural question is, How does this approach compare with traditional economic cost-benefit analysis, where benefits are measured in units of money (rather than of wellbeing)? The answer is that the results of existing **cost-benefit analysis** can be incorporated very easily into the wellbeing framework. For we know the impact that money income has on wellbeing, and we can always therefore convert any benefits measured in money into benefits that are measured in units of wellbeing. To be precise, if income is Y and X is a policy variable, the effect of the policy on wellbeing is given by

$$\frac{dW}{dX} = \frac{dY}{dX} \cdot \frac{dW}{dY} \quad (2)$$

where dY/dX is the effect of X on income, and dW/dY is the marginal utility of money.

This approach is often very useful, because for some policies it is easier to measure their effects initially in monetary units. This applies not just to direct effects on wages or other incomes (due, for example, to educational investment), but also to intangible benefits (like reduced journey times). These intangible benefits are in this case valued by what people's behaviour shows they would be willing to pay for them (by their **revealed preference**).⁵

But **willingness to pay** works only when people can show by their choices how much they value different outcomes. Sometimes they can do this, but very often they cannot. They can do it for things like transport, industrial production, education and some aspects of the environment. But many outcomes are not things that people can choose – they are things that just happen to people through outside influences – what economists call external effects. People fall sick, children get abused, elderly people get abandoned and people get mugged. We cannot learn about how much people value these experiences by observing choices. So how are we to evaluate policies like vaccination, or child protection, or family courts, or elderly care, or police protection? Measuring benefits in units of wellbeing is an obvious solution.

Critics might argue that, even though people can't show their values by their choices, we can ask them **hypothetical questions** about how much they would in principle be willing to pay to promote these goods? Unfortunately, however, it has been shown repeatedly that asking people hypothetical questions about how they value things produces nonsensical answers.⁶

So data on the happiness effects of activities may offer a better route to evidence-based policy making. But why not then translate those wellbeing estimates back into

⁵ Layard and Glaister (1994).

⁶ Kahneman, Ritov and Schadke (2000).

units of money? Thus, suppose we know the effect of a policy upon wellbeing. We could compute the equivalent gain in income that would increase wellbeing as much as the policy would. This **equivalent variation** in income could be computed by dividing the change in wellbeing by the marginal utility of income. Thus, if policy X improves wellbeing by dW/dX , the equivalent variation in income is given by

$$\frac{dY}{dX} = \frac{dW}{dX} / \frac{dW}{dY}. \quad (3)$$

This is the reverse operation to that in equation (2). The resulting aggregate benefit measured in units of money could then be compared with the cost.

There are, however, two overwhelming objections to this approach. First, it automatically makes changes in happiness less important if they occur to poor people – it treats a dollar as the same whether it belongs to a Trump or a tramp. To avoid this, the results could be analysed separately for different income groups, applying a different marginal utility of income to each group. This re-establishes wellbeing as the measure of benefit, so why not simply stay with it in the first place?

Second, we might not want to simply add the ΔW s, but rather to give extra weight to those with low initial happiness. If the monetary valuation procedure is followed, there is no way to do this, since the happiness level of each individual has become invisible. We may or may not want to give extra weight to those who are most miserable (see Chapter 2), but it is helpful to retain the ability to do so.

Note that throughout this chapter we assume that the total budget is determined by political considerations. We do not use the wellbeing approach to determine the total of public expenditure. The same thing happens with traditional cost-benefit analysis – projects are often rejected even if their monetary benefits exceed their costs. The reason is that there is not enough public money to finance all projects whose benefits exceed their costs. In consequence, the only projects that get through are those with a high enough ratio B/C , with the cut-off value being frequently higher than 1.

The alternative approach would be to allow wellbeing cost-effectiveness analysis to determine the total of public spending. But this would lead to much higher public expenditure. For example, if $W = 0.3 \log Y$, the value of a healthy life year (with $W = 7.5$) is in England \$750,000. But England's NICE does not sanction the expenditure of more than \$40,000 per additional life-year.⁷

Taxes and regulations

Going on, there are other important public policy problems besides how to spend a given budget total. There is the issue of how to structure the taxes. The approach here can be straightforward. If we envisage a self-financing tax change, we would simply

⁷ Moreover, if public money is raised by general taxation, we have to allow for the fact that, when other people pay a tax, that generates a positive externality for the rest of the population. So the equation $W = 0.3 \log Y$ overvalues the cost of reducing incomes generally across the population. A lower coefficient would imply an even higher optimal scale of public expenditure.

evaluate how this alters the happiness of each member of the population and aggregate these changes (assuming we are simply maximising the sum of wellbeing across all members of the population). We should also use wellbeing as the criterion for whether to introduce a new regulation or to abolish an old one.⁸

Five major issues

It is time to address some thorny issues. First, there is the issue of the **discount rate**. We discussed this in Chapter 15, where we suggested something like 1.5% per annum.

Then there is the issue of the **length of life**. If we know the quality of any additional life-year that a policy would produce, we would value the extra life-year by the quality of life it produces. But otherwise we would value changes in life-years by the average level of wellbeing (in the country in question).

Then there is the issue of how we treat the **birth-rate**. If we increase discounted WELLBYS by encouraging the birth of more children, would that count as a benefit? If it did, it would almost certainly be the most cost-effective way of increasing the sum of future wellbeing. For example, by switching expenditure from healthcare to child subsidies, we could surely increase the birth rate by more (in %) than we reduced the length of life. This would increase the number of future WELLBYS. But most people would not support the policy. We would therefore propose that, when evaluating the effect of policies, we ignore any effect on total WELLBYS coming from changes in the number of people born.⁹

Next, there is the issue of **whose wellbeing** counts. In principle, it should cover at least the whole of humanity. Every person is equally important. But for private ethics, there are some people who are easier for us to help than others. So, in practice, human society works through a division of labour. People take especial care of people close to them, and this also satisfies the need that humans have for a kind of affection they can only give to a small number of people.

But, when it comes to public policy or charitable activity, the **circle of concern** has to be widened to include people we do not know, including people in far-off parts of the world.¹⁰ Ideally, each government would choose to do whatever it could for the good of humanity. In practice, democratic governments inevitably feel that their main responsibility is for their own electorates. But this responsibility would hopefully include two other important considerations:

⁸ In practice, a regulation may also have a net cost to the organisation making it. To allow for this we need to rewrite the decision criterion in equation (1) as $\sum \sum W_{it}(1 - \delta)^t > \lambda \sum C_t(1 - \delta)^t$

On regulation, one obvious issue is ‘Do smoking bans improve human wellbeing?’ This has been studied using data on more than half a million Europeans since 1990. Odermatt and Stutzer (2015). The conclusion is that the ban increased the life satisfaction of those smokers who wanted to quit, without significant negative effects on any other group. See also Gruber and Mullainathan (2005).

⁹ In other words, we treat the number of people born as exogeneous. But if it changes, we still look at the average wellbeing of all those born. For an extensive discussion of this issue, see Meade (1955); Parfit (1984); and Broome (2004).

¹⁰ Singer (1981).

- the altruistic desire to help less happy nations and
- the necessity of collaborating with other countries to secure global public goods like fighting climate change and securing world peace.

And what about the wellbeing of **other sentient beings**, besides humans? They must surely count. There is ample evidence that birds and mammals (at least) have feelings of pleasure and pain. For example, researchers have offered injured birds or mammals the choice between food that includes standard pain-killers and food that does not. Injured animals prefer the food with pain-killers. And, even more important, when they've taken the painkiller, they stop whimpering or calling out. This shows that the choice of the painkiller is not simply an automatic reaction to a wound but a reaction to an emotional feeling.¹¹

Finally, there is the issue of **equity**. As we said in Chapter 2, the starting point for public policy analysis could be the Benthamite approach of adding up all changes in WELLBYs, regardless of who they accrue to. But most people would probably wish to give extra weight to improving the wellbeing of those who are least happy. The problem is how to secure agreement on the weights. One obvious approach is to conduct a representative survey of the views of the public, and this is a high priority for future research. In the meantime, one natural approach is through sensitivity analysis, examining how far the results are reversed when different weights are used. Equally, when considering new policy initiatives, it seems natural to focus on those areas of life that account for the greatest amount of total misery.

Developing New Policies

The natural starting point in this search for new policies is to ask two very similar questions.

- What aspects of life do most to explain the inequality of wellbeing?
- What aspects of life do most to explain the proportion of people who have low wellbeing?

In practice, the answers to the two questions are very similar.¹² As we saw in Chapter 7, the answer to the first question comes from the following standardised regression equation:

$$\frac{W_i}{\sigma_w} = \sum \beta_j \frac{X_{ij}}{\sigma_j} + e_i \quad (4)$$

where β_j^2 measures the independent contribution of each variable X_j to the overall inequality (variance) of W .

¹¹ Singer (1995). To apply the WELLBY approach to non-humans remains however a major quantitative challenge.

¹² See A. E. Clark et al. (2018) and the comparison of our Figures 8.2 and 8.3. This result would be expected if, for example, all the right-hand variables were jointly normally distributed.

Table 18.1 What explains the variation in life satisfaction in adults over 25? (United Kingdom) – partial correlation coefficients

	β
Physical health	0.11 (0.01)
Mental health	0.19 (0.05)
Work (not unemployed)	0.06 (0.04)
Quality of work	0.16 (0.04)
Partnered	0.11 (0.03)
Income	0.09 (0.01)
Education	0.02 (0.01)

Source: See Figure 8.2. Standard errors in brackets.

Alternatively, we could focus on misery and ask what explains it. We could, for example, define misery as a level of life satisfaction below 6. Then the dummy variable for misery takes the following values

$$1 \text{ if } W < 6$$

$$0 \text{ if } W \geq 6.$$

If we run a regression equation of this dummy variable on our usual explanatory factors, the β_j^2 measure the independent contribution of each variable X_j to the presence or absence of misery.¹³

Empirically, it turns out that the pattern of β_j^2 obtained from estimating equation (4) are very similar to those obtained from the regression explaining misery, except that in the latter case the β_j s are all slightly smaller (because a binary variable is more difficult to explain).¹⁴ Given this similarity, it is enough to focus on equation (4) and search for policies in areas with high β , knowing that reducing the inequality in those areas would make a big impact on the prevalence of misery.¹⁵

In Table 18.1, we repeat the β s that we saw in Chapter 8. For each factor X_j , β_j^2 represents the share of inequality explained by the independent variation of X_j . Top is mental illness, with physical illness also important. Then comes the quality of work and personal relationships, and only then comes income. These are data for the UK, but similar rankings apply in other advanced countries.¹⁶

¹³ The variance of misery is $p_m(1-p_m)$, which for small values of p_m is close to p_m . If the right-hand variable is also a dummy variable (X) then (if it were the only independent variable) we could write $M_i = a X_i + b + e_i$, and $\bar{M} = a\bar{X} + b$. If instead, \bar{X} were zero, we could say that X had increased the proportion of people in misery by a p_x , where p_x is the proportion for whom X has the value of 1. By contrast $\beta^2 = \frac{a^2 p_i(1-p_x)}{p_m(1-p_m)}$. This is clearly related to ap_x .

¹⁴ See Figures 8.2 and 8.3.

¹⁵ This is obvious when the factor is essentially a binary ‘bad’, e.g., mental or physical illness, lack of a partner or unemployment. When the factor is continuous, it is important **how** the inequality is reduced – for example, it would reduce misery if we increased low income but not if we reduced high income.

¹⁶ A. E. Clark et al. (2018) Table 6.3.

But the basic message of the wellbeing approach is clear. Policy-makers would not focus too heavily on economic issues. To support wellbeing, they would also give at least as much serious, evidence-based attention to

- mental health (treatment and promotion),
- physical health (treatment and promotion),
- the quality of work,
- support for families and
- community building.

Experiments

But this is only the beginning of the search for new policies. The next step is to identify specific policy changes that might be considered. Once a plausible policy option has been identified, it would ideally be the subject of a proper randomised experiment in the field. Where it is impractical or unethical to randomise across individuals, it is often possible to randomise across areas or across institutions (schools, hospitals, etc.). From such an experiment would come information on the short-run wellbeing benefits and budgetary costs of the experiment – and then it can ideally be projected into the longer term using a model.

The final result of such an evaluation would be an estimate of the change in WELLBYs per dollar of expenditure – or alternatively the number of people removed from misery per dollar. Or these estimates can be expressed the other way round – as the cost per WELLBY or the cost per person removed from misery. Table 18.2 is a crude illustration of the latter approach. It is a back-of-the-envelope calculation and we only include it to provoke thought and discussion and to encourage the reader to do better. In the table, we have taken four standard methods proposed in the UK for reducing misery. We then estimate what public costs would need to be incurred each year to ensure that there was one less person in misery. (The assumptions are in Annex 18.2.) The outcome is somewhat surprising. Better mental health care is the most cost-effective of the four policies, followed by active labour market policy and physical health care, with income redistribution the least effective. Because redistribution is so expensive and relatively few of those in misery are also poor, it is often more effective to spend public money on services in kind – helping people to help themselves.

Table 18.2 Average cost of reducing the numbers in misery, by one person

	£k per year
Poverty. Raising more people above the poverty line	180
Unemployment. Reducing unemployment by active labour market policy	30
Physical health. Raising more people from the worst 20% of health	100
Mental health. Treating more people for depression and anxiety	10

Source: A. E. Clark et al. (2017)

So in this analysis, the top priority is building up the social infrastructure (health care, skills development, employment, and community services).

Who Is Doing What?

So how many policy-makers worldwide now view wellbeing as their goal and act accordingly? Many express support for the idea, but many fewer are yet implementing it. Both the European Union's Council of Ministers and the OECD in Paris have requested their members to 'put people and their wellbeing at the centre of policy design'.¹⁷ They favour an 'economy of wellbeing', where wellbeing is the goal that the economy serves; but at the same time, wellbeing is valued for its positive effect on the economy.¹⁸ In China, President Xi has repeatedly stated that 'the wellbeing of the people is the fundamental goal of development'.¹⁹

But the country that has gone furthest in making wellbeing their goal is New Zealand. In 2019, the Labour government there announced its first Wellbeing Budget, which attracted worldwide interest. Its novel feature was to focus any new additional expenditure on things that increase wellbeing in a cost-effective way (mental health services, reduction of child poverty and domestic violence, Maori wellbeing and climate change). The rest of the budget was justified in terms of its effect on four pillars (physical capital, human capital, social capital and natural capital), which were seen as contributing to the sustainability of wellbeing – but that contribution was not quantified.

The New Zealand approach is one way in which change may come about. But one day governments (central and local) and NGOs may go further and evaluate their **whole** operation through the quantitative lens of wellbeing.²⁰ This is increasingly authorised in official manuals on evaluation,²¹ but most policy-makers have yet to use these tools. If they wanted to, they would need to establish their own Wellbeing Analysis Units that could scrutinise more and more of the policies they funded in terms of their effects on wellbeing.

And they would take urgent action to improve the knowledge base. This requires, as we have said, literally thousands of experiments in which one policy is compared with a counterfactual for its effects upon wellbeing and upon costs. And, of course, all experiments ever conducted would measure wellbeing as one of the outcomes, whatever else they measured.

¹⁷ EU Council (2019).

¹⁸ This has also been pushed strongly by the Wellbeing Economy Alliance (based in Scotland) and by the Wellbeing Economy Governments which include Scotland, Iceland, Wales, Finland and New Zealand.

¹⁹ Speech on 18 October 2017 to the 19th National Congress of the Chinese Communist Party. There are 12 mentions of wellbeing in his speech.

²⁰ Local governments actively targeting wellbeing include Bristol, England; Jalisco, Mexico; Andhra Pradesh, India. NGOs can receive advice on this approach from the Happier Lives Institute, from Effective Altruism and from Give Well.

²¹ For example, HM Treasury (2020, 2021) in the United Kingdom.

The follow-up period for most experiments is quite short, even though the actual effects may be quite long-lived. To simulate the longer effects requires a model of how wellbeing evolves from year to year over the lifespan. So we need quantitative models of how wellbeing evolves over the lifespan – and of the claims which people in different circumstances impose on the public finances. Building these models is a priority for research.

We already know many of the coefficients that would apply in such models.²² But the better the knowledge base, the better the chances that policy-makers would use it. The wellbeing revolution would only happen through granular knowledge about the causes of wellbeing. Such knowledge would also become a central feature of modern social science.

Conclusions

If wellbeing were to be at the heart of policy-making, some major changes would be needed.

- (1) Every organisation would try in whatever way it could to generate the largest number of future WELLBYs (appropriately discounted).
- (2) Wherever there is a budget constraint, the available funds would go to those policies that generate the most WELLBYs (discounted) per dollar of expenditure (discounted).
- (3) Where traditional cost-benefit analysis measures benefits in units of money (rather than of wellbeing), these benefits could be readily changed into units of wellbeing by multiplying them by the marginal utility of money.
- (4) Because monetary cost-benefit is not able to capture more than a fraction of the benefits of public policy, it would be better to convert monetary benefits into wellbeing benefits rather than the reverse. Moreover, converting everything into money would sacrifice information since the marginal utility of income varies so much between people.
- (5) Policy-makers would not count as a benefit any effects of a policy change that affects the birth rate and through that the number of future WELLBYs.
- (6) Policy-makers would develop new policies in areas which are causing the largest numbers of people to live in misery (low wellbeing). This means areas with high β s. New Zealand has followed this approach. But, having developed the specific policies proposed, the next step would be to estimate their effect on total WELLBYs – subject to sensitivity analysis using differential equity-weights.
- (7) Thousands of experiments would be essential to evaluate possible specific policies. We would also need better models of the determinants of wellbeing over the life-course. The explanation of wellbeing would become a central aim of all the social sciences.

²² See, for example, A. E. Clark et al. (2018); Frijters et al. (2020); and Frijters and Krekel (2021).

Questions for discussion

- (1) How if at all does the wellbeing approach to public expenditure improve on traditional cost-benefit analysis. Can the two approaches be reconciled? How?
- (2) How can equity considerations be best incorporated in the design of policies aimed at wellbeing?
- (3) Would we ignore effects on the number of births?
- (4) What is your reaction to Table 18.2? How could you improve the analysis?

Further Reading

- Frijters, P., Clark, A. E., Krekel, C., and Layard, R. (2020). A happy choice: Wellbeing as the goal of government. *Behavioural Public Policy*, 4(2), 126–165.
- O'Donnell, G., Deaton, A., Durand, M., Halpern, D., and Layard, R. (2014). *Wellbeing and Policy*. Legatum Institute.
- Singer, P. (1995). *Animal Liberation*. Random House.