what works wellbeing

How cost effective is a workplace wellbeing activity?

A how-to guide to cost effectiveness analysis

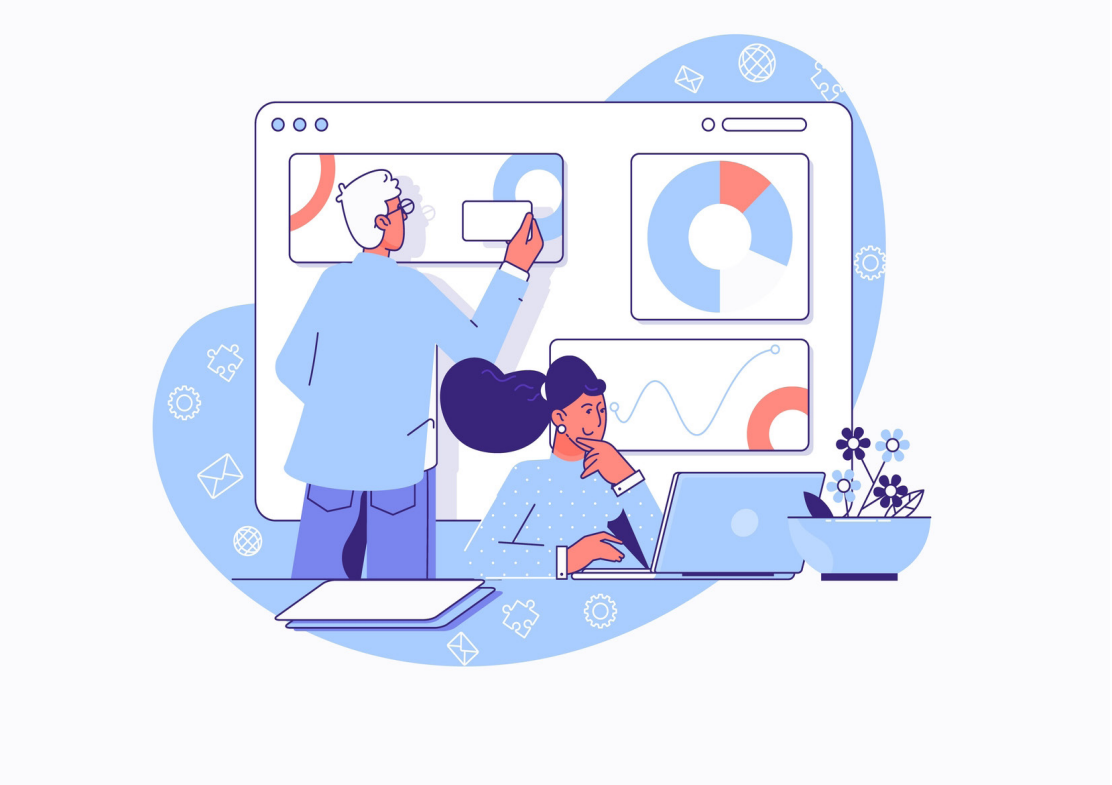

June 2020

The What Works Centre for Wellbeing helps people and communities to thrive by supporting decision-makers to understand what wellbeing is; how to measure it; and what works to improve it.
Find out more

We are an independent collaborating centre, bringing together decision-makers in government, local authorities, businesses, charities, funders, and academics, among others.
See our current partners

This guide is based on Bryce et al. (2020) and is produced by Magdalena Soffia, Senior Work and Wellbeing Analyst at the What Works Centre for Wellbeing in collaboration with Bryce at al.

How can this guide help you?

In 2019, the All-Party Parliamentary Group on Wellbeing Economics recommended that employers routinely measure workers' wellbeing and put policies and programmes in place to promote their wellbeing.[1]

If your goal is to improve staff wellbeing across an organisation or industry ,and have implemented an intervention in the workplace, this guide is for you. It will introduce you to the main actions needed to evaluate the cost-effectiveness of that intervention.

Intervention is any initiative or programme implemented with the intention of improving wellbeing. Ideally, this intervention should be implemented as part of a trial, where one group of people participate in the intervention while a similar group of people (the control group) do not.

Cost effectiveness calculator and guidance

For evaluations we recommend using the Cost Effectiveness calculator created by Bryce et al. (2020) along with the detailed step-by-step guidance and worked examples available in our website.

These resources are made for anyone who wishes to evaluate workplace wellbeing initiatives, including general managers or human resource managers in businesses, charities or public sector organisations, or officials or researchers undertaking a cost-effectiveness assessment on behalf of a workplace. You do not need to have any specialist expertise.

1 All-Party Parliamentary Group on Wellbeing Economics. 2019. A spending review to increase wellbeing: An open letter to the Chancellor. UK: APPG [https://wellbeingeconomics.co.uk/2019/05/24/a-spending-review-to-increasewellbeing-an-open-letter-to-the-chancellor-2019/].

What is Cost Effectiveness Analysis?

Cost-Effectiveness Analysis (CEA) is an estimation of how much an intervention costs relative to how effective it is at improving wellbeing.

CEA helps us to answer the following questions:

- Is what you spend on improving wellbeing worth it relative to the number of people helped and how much their wellbeing improves?
- If there are lots of different ways in which you can spend resources to improve wellbeing, which one is the most effective?

Although the CEA is done to monitor and evaluate the success of an intervention in place, learning about the general steps is also useful if you are at the stage of designing an intervention or assessing different programme options. This step-by-step introductory guide is particularly helpful to trial interventions implemented in an organisation, business or industry and aimed at improving staff wellbeing.

The new Cost Effectiveness Analysis

Cost-effectiveness evaluations measure the costs of interventions in monetary values and the wellbeing benefits in numerical units. The specificity of the approach used in this guide is that wellbeing is measured as life satisfaction.

Life Satisfaction is a widely used measure of individual wellbeing provided by the answer to a question similar to "Overall, how satisfied are you with your life nowadays?" Responses vary from 0 to 10 where 0 means "not at all satisfied" and 10 means "completely satisfied".

Since it is difficult to compare pounds with life satisfaction units, this guide follows the criterion used in national healthcare[2] where a treatment is regarded as cost-effective if the benefits of an extra year of life in perfect health (also called, quality adjusted life year, QALY) costs less than £25,000. Instead of calculating the cost of an extra year of life in perfect health, this guide will help you calculate the cost of improving one participant's life satisfaction by one point per year.

Conducting the analysis in five steps

For a cost-effectiveness evaluation, you need to follow five essential steps:

1. estimate the number of participants
2. estimate the costs of the intervention
3. estimate any changes to productivity
4. measure participants' wellbeing before and after the intervention
5. calculate the Cost-Effectiveness Ratio (CER).

2 National Institute for Health and Care Excellence (NICE).

Step one: gathering information about participants

The first data you need is the number of participants that took part in the intervention.

Participants (or intervention group) are the people in the workplace or industry who received the intervention. This does not include the people in the control group.

An ideal evaluation will also establish a control group to determine what would have happened without the intervention.

Control group consists of the people who took part in the intervention trial but were allocated to the group who did not receive the intervention. This group should be as similar as possible to the participant group (including working in the same workplace or industry), should not include those who refused to take part in the intervention, and their wellbeing should be measured at the same time and in the same way as the participants.

It is recommended that the number of participants (N participants) and the number of people in the control group (N control group) is at least 50 and that both groups are selected at random.

For more robust evaluations you can sample more participants and capture data on their key characteristics, like gender, age, job role, and so on. Since larger sample sizes improve the reliability of results and allow you to compare the cost-effectiveness of the intervention between different groups of workers.

Troubleshooting for step one

If you do not know the exact number of participants, do the CEA using an estimate of the number of participants.

If the number of participants is too small and you are not able to increase the sample, you can still conduct the CEA but making sure to interpret the results with caution, knowing that they are not statistically robust results.

If control group and participants cannot be selected at random, make sure that both groups share similar characteristics. If that is not possible, then interpret your results with caution, for differences between the wellbeing effects in participants and control group may be caused by differences in key characteristics.

If the intervention was rolled-out without a control group, you can still do a CEA evaluating the costs against wellbeing benefits for the participants only. However, you need to treat the results with caution because you are effectively assuming that those benefits would not have taken place without the intervention. Hence you risk overestimating the true wellbeing effects.

Step two: measuring costs

One of the minimum requirements to properly do CEA is to have data on the total monetary costs of the intervention for the organisation (C intervention).

Costs are any resource expended on the intervention that would not have been expended had the intervention not been implemented. This includes direct financial costs as well as indirect costs associated with delivery, opportunity costs associated with participation, plus the costs of any other resources. In this guide, all costs are expressed in GBP.

You need to decide which costs are important to you and which costs can be ignored. Generally, a large share of the total cost will be a result of the amount of time spent by participants engaging with the intervention and people delivering the intervention, this is time that could have been spent on other productive activities.[3] To express a person's time in monetary value, the default approach of this guide is to use their gross hourly wage (W).

Gross hourly wage is a measure of how much a person gets paid per hour before tax and other deductions. As this measure usually underestimates the true cost of a person's time to their employer, which might also include national insurance and pension contributions, we suggest increasing this figure by 25% every time you use it. The calculator does it automatically.

Step 2.1

Calculate the cost of time spent by participants by multiplying the number of participants, the total number of working hours spent by each participant, and their gross hourly wage, as it follows:

$$C_{participants} = N_{participants} * HRS_{participants} * W_{participants} * 1.25$$

Step 2.2

Calculate the cost of time spent by delivery personnel. This information may be readily available if you have already hired and paid the delivery personnel for their time. Otherwise, the costs of delivery personnel can be estimated in the same way as with participants, that is, multiplying the number of delivery personnel, by the number of hours spent by each of them, and their gross hourly wage:

$$C_{delivery\ personnel} = N_{delivery\ personnel} * HRS_{delivery\ personnel} * W_{delivery\ personnel} * 1.25$$

Delivery personnel are the people who spent time delivering the intervention, but who were not themselves participants. This may include people internal or external to the organisation, and may include volunteers or people being paid for their time.

3 The cost of people's time does not include the time spent by members of the control group, nor the time spent by participants and delivery personnel in monitoring and evaluation related activities (providing data, answering questionnaires, designing the evaluation, etc.), because this is time that is not being spent directly on providing and receiving the intervention.

Step 2.3
Estimate any other monetary expenses (C other) incurred in the intervention that would not have been incurred otherwise. This may include actual expenditure by the organisation delivering the intervention (e.g. travel) or the monetary value of other resources used such as equipment, premises, and so on.

Step 2.4
Lastly, calculate the net cost per participant by adding up all the partial costs estimated in steps 2.1 to 2.3, and dividing by the number of participants, as it follows:

$$C_{total} = \frac{(C_{participants} + C_{delivery\ personnel} + C_{other})}{N_{participants}}$$

Troubleshooting for step two

If participants spent different amounts of time engaging with the intervention, you must calculate participants costs separately for each group and then add them up to obtain the total costs of participants. The same applies if delivery personnel spent different amounts of time delivering the intervention.

If you do not know the exact salary, you can estimate it based on the average gross hourly wage of their occupational grade. This information can be found in the Office for National Statistics website[4] and is included in the Online Calculator.

If participants or delivery personnel belong to different pay grades or occupations, you can calculate time costs for each group of people separately.

If salary information corresponds to gross full-time equivalent annual salary, you must estimate the average hourly wage by dividing the annual salary figure by 2,080 (i.e.40 hours a week for 52 weeks in the year).

If the time value of some of the participants or delivery personnel is unknown (e.g. because they are unemployed, retired people, children, volunteers), use your own judgement to estimate the value of their time depending on the circumstances. For instance, if the time incurred could have been used in other productive activities, such as job search or volunteering, it may be appropriate to value people's time at the minimum national wage or the median wage. If engaging with the intervention is a leisure activity it may be appropriate to value participants' time at zero.

If some of the facilities used to deliver the intervention are already owned by the organisation (e.g. a company boardroom) or will be used later on for other things not related to the intervention (e.g. mobile phones), it is recommended to include only the proportion of the cost corresponding to the amount of time that these assets were used in the intervention or the cost of rent/hiring that facility for the duration of the intervention.

4 Available at www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/occupation4digit soc2010ashetable14.

Step 3: measuring productivity benefits

Did the wellbeing intervention have any benefits to the organisation in terms of improved staff productivity (P) that can be quantified? If that is the case, productivity rises should be included as negative costs.

Productivity is the amount of output a person produces in their job over a given period of time, expressed in GBP. Wellbeing interventions may influence workers' productivity, for example by improving their skills, creativity, motivation, mental and physical health.

Productivity can be measured through different indicators, such as days of sickness absence of employees, amount of output produced by employees, improved customer satisfaction, etc.

The approach to quantifying productivity changes in GPB will vary depending on the measure of productivity you use. For instance, if you are measuring productivity in terms of sickness absence, you need to consider the total number of more or fewer sickness days recorded, and the average hourly wage of your participants.

An alternative is to estimate the percentage change in outputs produced by each worker and work out the monetary value of such output. You can find examples in the calculator user manual (Bryce et al. 2020) and the calculator will help you to impute some of these negative costs.

Since some indicators of productivity are difficult to quantify and will not be included, any estimate of the cost-effectiveness of an intervention is likely to be conservative from the employer's perspective.

Due to the difficulties with measurement, whenever productivity is measured and added, particular caution is needed to assess whether and to which extent the estimates is likely to overestimate effects.

To make sure that the productivity benefits observed are the result of the intervention, you need to measure average productivity among participants before and after the intervention (ideally, in the 12 months before and those following the intervention), and you need to compare this change to the one observed in the control group.

Step 3.1
Estimate the average (mean) productivity of the participants and control group at baseline and follow-up after the intervention.

Step 3.2
Estimate productivity changes from participants as the difference between their average productivity at follow-up, minus their average productivity at baseline, in percentage points.

Step 3.3

Estimate productivity changes in the control group as the difference between their average productivity at follow-up, minus their average productivity at baseline, in percentage points.

Step 3.4

Lastly calculate the net productivity benefits as the difference between the percentage productivity changes in participants (Step 3.2), minus the percentage productivity changes in the control group (Step 3.3), and then multiplying by the number of participants, their gross hourly wage, and the number of hours productivity benefit was sustained:

$$P_{total} = (P_{change\ participants} - P_{change\ control\ group}) * N_{participants} * W_{participants} * HRS_{productivity}$$

Step 3.5

Recalculate the net cost per participant estimated in Step 2.4 including the net productivity benefits (Step 3.4) as negative costs:

$$C_{total} = \frac{(C_{participants} + C_{delivery\ personnel} + C_{other} - P_{total})}{N_{participants}}$$

Troubleshooting for step three

If there is no reliable information on productivity changes, you can simply skip this step. Measuring productivity changes will improve the robustness of your evaluation, however, if the evidence on productivity changes you have is weak, you risk hugely overestimating the true cost-effectiveness, therefore it is better to leave them out.

If there is only information about productivity changes for participants, you may work with the assumption that productivity changes to the control group were zero, which risks distorting your CER. Overall, if you are evaluating the intervention without a control group, it is recommended that you do not include productivity benefits in your calculation.

If you only have evidence of productivity changes for less than 12 months after intervention, then only include that amount converted to hours. The Online Calculator requests the number of months the productivity gain was sustained and converts that figure into hours. For instance, if the effect lasts for 12 months, then multiply by 2080 hours. If it lasts for 6 months, multiply by 1040 hours.[5]

5 Note that if the user puts in a time period of more than 12 months, the Calculator assumes that the effect was sustained for just one year, to err on the side of caution.

Step 4: calculating wellbeing benefits

How much has wellbeing improved since the start of the intervention? As a minimum requirement you need to measure participants' wellbeing level immediately before the start of the intervention (baseline data), and at least once after the intervention started.[6]

Wellbeing is a broad term that describes the extent to which a person has a good life. It can be measured in lots of different ways, preferably, at an individual level. This guide uses life satisfaction as a standard measure but other measures can also be used as long as they are converted into life satisfaction units.

As when calculating productivity gains, the advantage of having a control group is that you can determine how much of the observed wellbeing benefits for participants would have occurred in the absence of the intervention. This is called the counterfactual change.

Counterfactual change is the change that would have occurred without the intervention. It is measured through any changes in wellbeing experienced by the control group.

The method to calculate the net wellbeing effect (that is, taking into account the counterfactual) is called the difference-in-difference, and equals the change observed among participants minus the change observed in the control group. To make sure that the changes occurred in both groups are comparable, baseline and follow-up wellbeing measures should be collected at the same time for both participants and control group.

Also, to make any changes between the participants and control group comparable, you should capture their individual wellbeing using the same question to both groups and at each point in time. Ideally, you will ask participants about their Life Satisfaction (LS) because this is considered a reliable indicator of wellbeing.[7]

Step 4.1
Estimate the average (mean) life satisfaction of participants and control group at baseline and follow-up.

Step 4.2
Estimate the changes to the participants as the difference between their average life satisfaction at follow-up, minus their average life satisfaction at baseline, in LS units.

6 In theory, you can start to evaluate wellbeing benefits even before the intervention is finished.

7 Layard, R., 2016. Measuring wellbeing and cost-effectiveness analysis using subjective wellbeing. What Works Centre for Wellbeing: Measuring wellbeing series, Discussion Paper 1.

Step 4.3
In the same way, estimate the changes to the control group as the difference between their average life satisfaction at follow-up, minus their average life satisfaction at baseline, in LS units.

Step 4.4
Calculate net LS effect per participant as the difference between the changes in participants (Step 4.2), minus the changes in the control group (Step 4.3) as indicated below. Note that, when using the calculator, the user only has to actually do step 4.1, the rest is calculated automatically.

$$LS_{total} = LS_{change\ participants} - LS_{change\ control\ group}$$

Step 4.5
How long are the benefits expected to last? The calculator provides results based on assumed durations of one, two and three years, so the user can make a judgment as to which one is most realistic. To estimate the duration of the effect and whether it diminishes or increases over time you can gather additional wellbeing measures and repeat steps 4.1 to 4.4 at each point in time.

Troubleshooting for step four

If the wellbeing measure is not Life Satisfaction (e.g. happiness, job satisfaction, General Health Questionnaire, etc.), you may be able to convert it into Life Satisfaction units using the exchange rates provided in Bryce et al (2020). The calculator will do the conversion automatically.

If there is no control group to calculate the net Life satisfaction effect, you may work with the assumption that the change in LS that would have occurred without the intervention is zero, therefore your net LS benefits would be the same as the benefits to the participants alone. However, making this assumption risks overestimating the true impact of the intervention so you need to interpret the results with caution.

If some of the participants or the people in the control group did not answer questions about their wellbeing at the time of the follow-up, their wellbeing data should be removed from the analysis, and the net wellbeing benefit calculated over the number of participants that completed the intervention.

If the follow-up wellbeing measure corresponds to less than a year after the intervention, you can work with the assumption that the observed effect will be sustained for the year.

If you have more than one follow-up wellbeing measure suggesting that the intervention continues to have effects that last beyond a year, then these should be included but values should be adjusted through discount rates to take account of how far into the future they are sustained.[8] However, if you use the calculator you do not need to worry about discount rates. You will have the option of inputting up to three post-intervention evaluation points and the results will then show what this does to the CER depending on different assumptions about the persistence of effects.

If wellbeing was not measured both before and after the intervention, it will not be possible to conduct the evaluation or use the calculator. You can only gather information about its cost and compare it against other benefits such as improved work engagement, without being able to determine whether it is worth the money in terms of wellbeing gains.

8 See Wright L, Peasgood T, and MacLennan S. 2017. A guide to wellbeing economic evaluation, What Works Centre for Wellbeing.

Step 5: calculating the CER and making sense of results

Cost-effectiveness ratio (CER) is a measure of how much it costs to improve one person's wellbeing by a certain amount. In this guide, the CER is interpreted as the net cost of improving one person's life satisfaction by one point (on a 0-10 scale) for one year.

Step 5.1

Calculate the cost-effectiveness ratio (CER) as the ratio between the net cost per participant (as estimated in Step 2.4, or in Step 3.5 if you are including productivity changes) and the net LS effect per participant (as estimated in Step 4.4):

$$CER = \frac{C_{total}}{LS_{total}}$$

Step 5.2

Is your intervention cost-effective? An intervention is regarded as cost-effective if, over a year, it can deliver an extra point in life satisfaction costing less than 2,500 GBP. This is an indicative threshold which is comparable to the criterion used by the National Health and Care Excellence (NICE).

However, employers may identify higher or lower thresholds depending on their own willingness to pay for an improvement in wellbeing. Some may want their investment in the intervention to be fully recouped in productivity gains, therefore they would consider 0 GBP as their threshold.

CER will help inform decisions about whether interventions should be rolled out to more participants or other workplaces. The results may also be used to help select the lowest CER between multiple feasible intervention options.

Some interventions may be cost-effective if targeted at a particular group of people (for example, people of a specific gender, people in a specific job category, or people with specific health problems). The quality of the evaluation may be improved if you estimate CERs for different participants subgroups. If you think that different groups may benefit differently, you should run a separate trial for each subgroup from the start and follow the five steps of CEA for each group in turn.

From a societal perspective, an increase in workers' LS may have spillover benefits to the wellbeing of their family, friends or anyone else affected by the intervention.

Likewise, there might be wider social costs and savings external to the workplace, such as the time spent by participants in their own free time, the use of volunteers to deliver the intervention, free or subsidised use of premises and equipment, funding received from another source, or even savings for the NHS.

Note that the principles introduced in this report only consider the costs and benefits incurred privately from a workplace or employer perspective, therefore the results must be interpreted as a conservative estimate of the true benefits of the intervention to society.

A practical example: evaluating an intervention on goal setting for working age adults[9]

Oliver and McLeod (2018) ran a Randomised Control Trial to evaluate the effectiveness of an online version of a wellbeing intervention on goal setting for public sector employees across the UK.

There were 158 employees in the intervention and 149 in the control group. Employees spent three hours of their working time completing the online training modules. In addition, all participants spent 30 minutes on phone follow-up conversations with the researchers.

The intervention was mainly delivered by a researcher who spent around 75 hours of work-time adapting the intervention to working age adults and conducting telephone conversations with the intervention group. It is estimated that £300 was spent on travel by the researcher and £370 was spent on the design of a website to host the intervention.

Both participants and employees in the control group reported improved wellbeing twelve months after the start of the programme. Specifically, the average life satisfaction scores of participants increased (on a scale of 0 to 10) from 6.40 to 6.84, and the average life satisfaction of the control group fell slightly from 6.52 to 6.51.

Was this programme cost-effective?

Step one: how many people participated in the intervention? In the example, the number of participants was 158, and 149 people were in the control group.

Step two: how much did the intervention cost? In the example, the total cost of the intervention is not provided so you have to work out the cost by accounting for three types of costs: participation cost, delivery personnel costs, any other costs.

Step 2.1: in total, each participant spent 3.5 hours engaging with the intervention. The information on the employees' pay is not provided, however, since all of them are national administrative employees, it is assumed they all belong to the

9 This example uses a simplified calculation and will not give exactly the same results as the online calculator. For the original study see: Oliver, J.J. and MacLeod, A.K., 2018. Working adults' well-being: An online self-help goal-based intervention. Journal of Occupational and Organizational Psychology, 91:665-680. Thanks to Jeremy Oliver and Andrew MacLeod for sharing some of their unpublished data on life satisfaction.

same pay grade. According to the ONS data, the average gross hourly wage for 'National Government Administrative Occupations' is £12.48.

Hence, participants costs can be estimated as:

$C_{participants}$ = 158 employees * 3.5 hrs * 12.48 GBP * 1.25 = £8,626.80

Step 2.2: the intervention was delivered by one researcher only, who spent 75 hours delivering the intervention. The information on the facilitators' pay is not provided, however, the average gross hourly wage of 'Higher Education and Teaching Professionals' is 28.12 GBP according to the ONS data. Hence, delivery personnel costs can be estimated as:

$C_{delivery\ personnel}$ = 1 researcher * 75 hrs * 28.12 GBP * 1.25 = £2,636.25

Step 2.3: additional expenses included only travel expenses (£300) and website design (£370). Hence, other costs can be estimated as:

C_{other} = 300 + 370 = £670.00

Step 2.4: based on the information in steps 2.1 to 2.3 the net costs per participant of the intervention can be estimated as:

$$C_{total} = \frac{(8{,}626.80 + 2{,}636.25 + 670.00)}{158\ employees} = £75.53$$

Step three: were there any productivity benefits recorded by the participants and control group? The example does not provide any measure for productivity changes; therefore, this step can be ignored.

Step four: what were the wellbeing benefits recorded by participants and by the control group?

Step 4.1. The average life satisfaction of participants and control group at baseline and follow-up provided in the example were the following:

	Average LS (baseline)	Average LS (3 months)
Participants	6.40	6.84
Control Group	6.52	6.51

Step 4.2. The changes to the participants can be estimated as:

$LS_{participants}$ = 6.84 – 6.40 = 0.44

Step 4.3. The changes to the control group can be estimated as:

$LS_{\text{control group}} = 6.51 - 6.52 = -0.01$

Step 4.4. Lastly, the net LS effect per participant can be estimated as:

$LS_{\text{total}} = 0.44 - -0.01 = 0.45$ LS points over a year[10]

Step 4.5. What was the duration of the effect? The example only provides one follow-up measure a year after the intervention started, therefore this cannot be calculated. If there were more than one average LS follow-up measures, the calculator would estimate a different figure of net wellbeing benefits using discount rates and specific assumptions about the evolution of wellbeing over time.

Step five: from the example, the final CER would be calculated as:

$$\text{CER} = \frac{£75.53}{0.45 \text{ LS}} = £167.84$$

It is estimated that the goal setting intervention costs £167.84 to improve the life satisfaction of one employee by one point over a year - assuming that the wellbeing benefits apparent at three months persist for a year. Although we are not able to include information on any productivity gains, we do find that the intervention is cost-effective because the CER falls below the recommended threshold of £2,500.

Caveats

You will often find less than ideal examples out there. These examples teach us how difficult it may be to conduct evaluations properly, sometimes not being feasible at all. In practice, you are likely to find incomplete information on costs and wellbeing benefits.

Some of these information gaps can be worked out over assumptions which are necessary to make explicit at the moment of interpreting the results. For instance, if (as in the example above) average LS measures for participants were only available at baseline and three months after the start of the intervention, you can work over the assumption that the average LS for participants at twelve months remains the same as the average LS of participants at three months after baseline. In this case, you must inform that you assumed that wellbeing benefits apparent at three months persist for a year.

Another common difficulty you may find is when a parallel wellbeing intervention is taking place. How can you be sure that the observed changes in wellbeing are exclusively the result of the intervention evaluated? If you are still at the stage of planning a trial, make sure that it is timed such that your participants and/or control group are not receiving some other intervention that may affect their wellbeing.

10 In this example the data on wellbeing was only collected at baseline, 5 weeks and 3 months, and so we use the longest time period (3 months) and work with the assumption that these changes were sustained for a year.