
https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

# 13 Income*

Money is like muck, not good except it be spread.

Francis Bacon

In 1974, Richard Easterlin published an important article. It was called 'Has economic growth improved the human lot?'.[1] In it he focused on two alleged facts that were not apparently consistent with each other:

(1) At a point in time, richer people are on average happier than poorer people, (though there is a huge overlap).
(2) Over time, increases in national income per head do not lead to increases in happiness.

This was the so-called **Easterlin paradox**. If richer people are happier than poorer people, you would think that, when a country becomes richer, it will also become happier. But, Easterlin claimed, this is not what happens. And his explanation was that, when national income increases, everyone increases the norm against which they compare their own income. The most obvious norm for comparison is the income of other people. If it is only their **relative** income that matters to people (rather than their **absolute** income), then that could explain why economic growth does not increase the nation's happiness.

In this chapter, we shall explore this hypothesis and much else besides.

- First, we shall look at individual happiness and how this is affected by a person's income. Similarly, we shall look at happiness and income across countries at a moment in time.
- Then we shall look at national happiness over time.
- Finally, we shall explore the role of relative income compared with absolute income and the policy implications of this difference.

* This chapter has benefitted greatly from help from Andrew Clark.

[1] Easterlin (1974).

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

## Differences Between Individuals

As between individuals, three central findings of wellbeing research are these:

(1) In every country, richer people are on average happier than poorer people.
(2) This difference is quite small and explains around 2% of the variance of wellbeing within the population.
(3) The effect of additional income gets smaller, the richer you are.

Figure 13.1 shows the position in the UK. Each bar shows the average wellbeing of people at each level of income. As you see, richer people are on average happier than poorer people. But there is also a huge overlap. This is shown by the thin lines around each average figure. The lines span from 1 standard deviation below the average to 1 standard deviation above (i.e., they show the range of wellbeing for the middle two-thirds of people at each level of income). As the figure shows, there are many poor people who are happier than the average rich person.

### The diminishing marginal utility of income

The diagram also shows something else. Extra income makes more difference to wellbeing at the bottom (left-hand) end of the scale than it does at the upper end of the scale. An extra $ of income produces a smaller and smaller amount of extra wellbeing the richer the person is. This old idea is now called '**the diminishing marginal utility of income**'; and before the behaviourist revolution (discussed in Chapter 2), it was a central belief of every economist. Indeed, it was why most economists favoured some redistribution of income – because, when $1 was transferred from a rich person to a poor person, the rich person lost less wellbeing than the poor person gained. So overall there was a net gain in wellbeing.

Until the last few decades, this was simply a speculative belief. But the new science now makes it possible to measure the quantitative effect with some accuracy. And, once we know how income affects wellbeing, we can compute by how much income inequality reduces average wellbeing, compared with a situation where everyone received the existing average level of income (see Box 13.1).

In investigating the effect of income on wellbeing, there are two issues:

(1) What is the functional form of the relationship?
(2) What is the actual size of the effect of income on wellbeing?

On the functional form, we can investigate this empirically, and it turns out that the semi-logarithmic linear relationship is a very good fit to the data in a whole range of surveys.[2]

Thus

$$W = \alpha \ log \ Y + etc \tag{1}$$

[2] Layard et al. (2008).

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

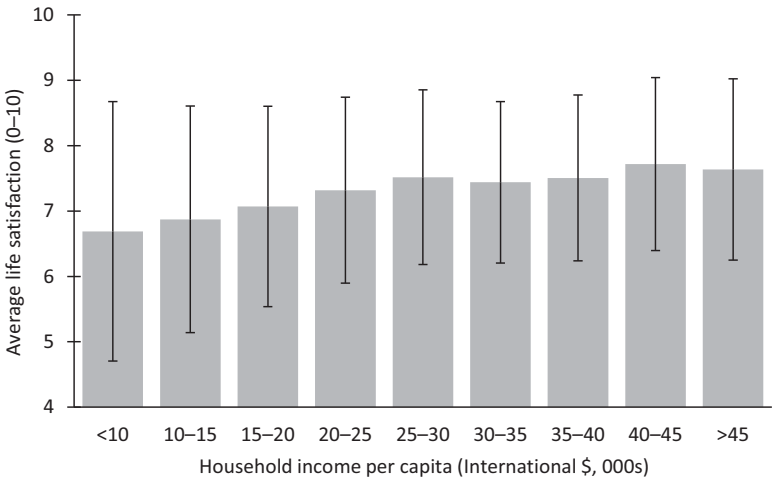


**Figure 13.1** Average life satisfaction at different levels of income (Britain) (Bracketed range includes 2/3 of each income group)
*Source*: Gallup World Poll 2017–19, Cantril ladder

where $Y$ is household income per equivalent adult. This fits in with the general law, known as the Weber–Fechner effect (discussed in Chapter 1), which says that, whether the experience is of light or sound, the perceived size of a change depends linearly on the proportional change in the thing that is changing. In a similar way, our feelings about a change in income depend linearly on the proportional change in income.

This gives us a direct measure of the marginal utility of income,[3] since

$$\frac{dW}{dY} = \frac{\alpha}{Y}. \tag{2}$$

So, a person's marginal utility of income is inversely proportional to her income: an extra $1 matters 10 times more to a poor person than to someone who is 10 times richer. If equation (2) is exactly right, the marginal utility of income will always be positive, even though it can be very low when a person's income is high; and there is some evidence that, in fact, marginal utility becomes zero at some point – the point of 'satiation'.[4]

## The size of the effect of income on wellbeing

However, we still need to know the size of the coefficient $\alpha$ To find it, we could estimate a cross-sectional equation of the following form

[3] Prior to the science of happiness, economists tried to infer this from the degree of risk aversion or other indirect means.
[4] Jebb et al. (2018). Using the Gallup World Poll, they estimate satiation at around $100,000 per equivalent adult in advanced countries, China and the Middle East/North Africa and around $50,000 in other regions.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

$$W_i = \alpha \log Y_i + \sum_j bjX_{ij} + u_i \quad (3)$$

where the $X_j$s are the other things that determine wellbeing – acting in this case as control variables. These controls should obviously include age, age squared[5] and gender, but what else is not completely clear. For there are at least five problems that arise when estimating equation (3).

## Problems

(1) If income is **measured** inaccurately, the estimate of $\alpha$ will be biased downwards.
(2) If we include controls that are themselves affected by income (and thus **'mediate'** its effect), we shall underestimate the total effect of income upon wellbeing.
(3) If, on the other hand, we omit variables that are positively correlated with both income and wellbeing, we shall overestimate the effects of income, by omitting these **'confounding'** variables.
(4) The relationship in equation (3) may not be properly identified, because there may also be a relation between income and wellbeing in which wellbeing is causing income. This is the problem of **reverse causality.**
(5) If people are affected by **other people's incomes**, we should include these in the controls.

We shall deal with these problems as follows.

(1) We measure income as **income per person** in the household (children being converted into 'equivalent adults'). This is a proxy for living standards. (Though consumption data would be preferable, they are not generally available in surveys where we also have data on wellbeing.) In most surveys, incomes are self-reported, but they are no less accurate than data on, for example, health.
(2) We shall show estimates **with and without** including mediating variables.
(3) To handle confounding variables that are omitted, we shall exploit the **panel** nature of the data, by including in each equation a person fixed effect $f_i$ . This will remove the effect of any omitted variables that do not vary over time. It gives an equation

$$\mathrm{W}_{it} = \alpha \log \mathrm{Y}_{it} + \sum_j bj\mathrm{X}_{ijt} + f_i + v_t + e_{it} \quad (4)$$

The estimate of $\alpha$ is now based on comparisons 'within-person' (rather than across persons). As we shall find, such panel equations tend to give very low estimates of $\alpha$,[6] but this is partly due to enhanced effects of measurement error[7] and partly because the timing of effects may not be properly represented by the form of the equation.

[5] This needs to be included because in advanced countries happiness tends to follow a U shape – higher in youth and old age and lower in middle age (see Chapter 14).
[6] See Table 13.2. [7] See any econometrics text.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

(4) To deal with the problem of reverse causality,[8] we shall show the results of including an element of income that is completely random and therefore **exogeneous.** This is people's winnings in a lottery.
(5) We shall normally include the log income per head of a person's **comparator** group (people of similar age, gender and region). This is to estimate the role of relative income as compared with absolute income. However, we leave reporting these comparator effects till later in the chapter.

## Results

We can begin with simple **cross-section** results for Britain, Germany, Australia and the United States. The data for the first three countries are from annual longitudinal household studies (UKHLS, SOEP and HILDA) and for the United States, they are from the annual cross-sectional BRFSS.

As Table 13.1 shows, the effects are not large. To take the US case, one additional point of log income, corresponding to nearly a tripling of income, will produce an extra 0.31 points of wellbeing (out of 10 points maximum).[9]

It is interesting to see how much of the variance of happiness is explained by income inequality. Remember that the partial correlation coefficient is

$$\beta = \frac{\alpha \text{ SD } (log\, Y)}{\text{SD } (W)} = 0.31 \times \frac{0.82}{1.55} = 0.16. \tag{5}$$

So the share of the variance of $W$ in the United States explained by income inequality is

$$\beta^2 = 0.16^2 = 0.0256 = 2.56\%.$$

This 2.56% compares with an $R^2$ of around 19% for the full set of influences shown in Figure 8.2. Yet, some economists claim that 'absolute income is the dominant fact determining wellbeing'.[10] That is clearly not correct. Income is one important influence but one among many.

Moving on to the issue of omitted variables, one possible solution is to use a fixed effects equation like that shown in Chapter 7 and thus exploit the panel nature of the data. As Table 13.2 shows, this produces smaller coefficients.[11] But we should probably ignore these time-series estimates due to the problems noted earlier.

A final way to handle omitted variables and reverse causation is to use data on lottery winnings. In the majority of studies, these are followed up over short periods

[8] De Neve and Oswald (2012) demonstrate that adolescent wellbeing affects subsequent income, using sibling fixed effects.
[9] These estimates hold other things constant. But some of these other things may be affected by income. To get a maximum estimate of the total impact of income we would hold nothing constant. In this case, the estimated cross-sectional coefficients tend to be very roughly double the estimates with controls of the kind we have shown.
[10] Sacks et al. (2013).
[11] This might suggest that the cross-sectional estimates may include an element of reverse causality.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

**Table 13.1** Effects of log income on life satisfaction (0–10) (pooled cross-sections)

| Britain | Germany | Australia | United States |
|---|---|---|---|
| 0.16 | 0.26 | 0.16 | 0.31 |
| (.01) | (.01) | (.01) | (.01) |

*Source*: A. E. Clark et al. (2018) Table 2.2; Britain, Understanding Society (1996–2014); Germany, SOEP (1984–2015); Australia, HILDA (2001–2015); United States, BRFSS (2006–14).
*Note*: Standard errors in brackets. 'Controls' include all those in Figure 8.1. Estimates omitting comparator income are very similar.

**Table 13.2** Effects of log income on life satisfaction (0–10) (individual fixed effects)

| Britain | Germany | Australia |
|---|---|---|
| 0.04 | 0.08 | 0.06 |
| (.01) | (.01) | (.01) |

*Source*: A. E. Clark et al. (2018) Table 2.2; Britain, Understanding Society (1996–2014); Germany, SOEP (1984–2015); Australia, HILDA (2001–2015).
*Note*: Standard errors in brackets. 'Controls' include all those in Figure 8.1. Estimates omitting comparator income are very similar.

and are therefore difficult to interpret. But in a remarkable study, Lindqvist and others followed up people who played the Swedish lottery for a further 22 years after that.[12] Among the players, the winnings could reasonably be counted as random. And the size of the winnings had a similar effect on wellbeing over all the 22 years. If we convert these one-off winnings into an equivalent income stream, the effect of a unit change in log income is to raise wellbeing (0–10) by 0.38 points.

One further issue: are the effects of log income different in **poorer countries** to the effects in **richer** ones? We can examine this, using data from the Gallup World Poll. The results are shown in Table 13.3. The coefficients are remarkably similar in countries at all levels of income. This is less surprising than might appear since (due to the diminishing marginal utility of income) an extra dollar in the hands of a poor person is worth 20 times more than in the hands of someone who is 20 times richer. We should also note that for high income countries the coefficient is somewhat higher than those shown in Table 13.1, which makes sense since in Gallup data the Table 13.1 countries do have coefficients somewhat below the high-income country average.

It is helpful to have in mind a benchmark coefficient for the effect of log income on wellbeing. In picking one, we should also bear in mind the danger that the Table 13.3

[12] Lindqvist et al. (2020). Obviously, lottery winnings may have different immediate psychological effects from other types of income – but less so if spent over a long period.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

**Table 13.3** Effect of log income on life satisfaction (0-10): By income of country (Pooled cross-section)

| High | Upper middle | Lower middle | Low | All |
|---|---|---|---|---|
| 0.37 | 0.43 | 0.45 | 0.35 | 0.40 |
| (.04) | (.04) | (.03) | (.04) | (.02) |

*Source*: Gallup World Poll 2009–19, Cantril ladder; individual data; regressions include as controls: unemployed, degree, partnered, health problems, age, age$^2$, country by year fixed effects; regressions by Ekaterina Oparina

coefficients are exaggerated through reverse causality. Thus we would suggest that a figure of 0.30 is a useful benchmark.[13]

This gives us straight away two vital pieces of information. The first is the marginal utility of income (meaning the change in WELLBYs for an extra dollar of annual income). As we explained earlier, this is given by $\alpha/Y$. So if annual income per head is \$30,000, the marginal utility of income is 1/100,000.[14] And 1 extra WELLBY is equivalent to some \$100,000 (spread across a group of people).

Second, we can measure the direct impact of inequality on average wellbeing. As Box 13.1 explains, in the typical country this is of the order of 0.12 points (out of 10) – that is how much average wellbeing would rise if average income remained the same but income inequality was abolished. It is a surprisingly small figure.

**Box 13.1** The direct effect of income inequality on average wellbeing

We are interested here in the difference between average wellbeing as it now is and average wellbeing as it would be if everyone received the current average level of income. This difference equals[15]

$$0.3\left(\frac{\sum \log Y_i}{N} - log\ \bar{Y}\right).$$

Using a quadratic Taylor's series expansion of log $Y_i$ this becomes

$$0.3\left(log\,\bar{Y} + \frac{1}{\bar{Y}}(\bar{Y} - \bar{Y}) - \frac{1}{2}\frac{\sum(Y_i - \bar{Y})^2}{N\bar{Y}^2} - log\,\bar{Y}\right) = \frac{-0.3}{2}\frac{\mathrm{Var}\,(Y)}{\bar{Y}^2}.$$

In a typical advanced country, $\mathrm{Var}(Y)/\bar{Y}^2$ is around 0.8,[16] so the direct cost of inequality is 0.12 points (out of 10). There are also of course indirect effects through the pattern of human relationships and so on.

[13] The high-income country average in Table 13.3 is unaffected if we remove countries in the Middle East.

[14] This would mean that the **statistical value of a life-year** with wellbeing of 7.5 points is \$750,000 – higher than is usually allowed for in rich countries.

[15] This is analogous to the Atkinson (1970) income inequality index. If $W = \alpha log Y$, the Atkinson index of income inequality is given by $\mathrm{Var}(Y)/2\bar{Y}^2$.

[16] See A. E. Clark et al. (2018) annex, Tables D1–5.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

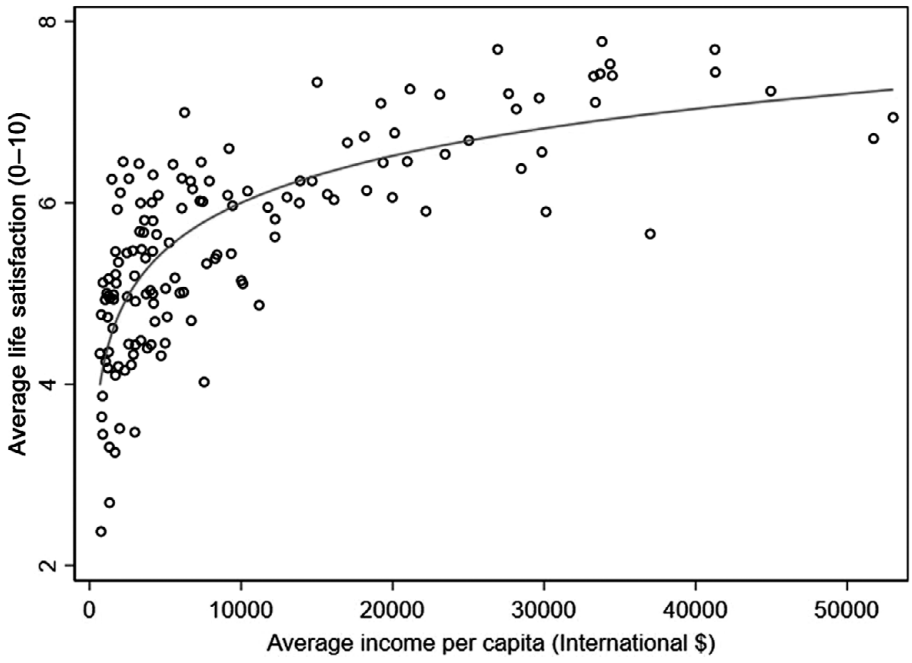


**Figure 13.2** Average life satisfaction and household income per head: across countries
*Source*: Gallup World Poll (2019)

## Differences Between Countries

We can turn now to differences in income between countries and how these are reflected in the country's average level of wellbeing. Figure 13.2 shows the scatter diagram of countries together with the line of best fit.

As with individuals, there is overwhelming evidence of diminishing marginal utility of income, and the line of best fit has been estimated using the logarithmic formulation. If no other variables are included, the effect of log income is very substantial. But, as we argued in Chapter 8, this is partly because high income is correlated with many other cultural variables including trust, social support, freedom and generosity. We do not know how far these other characteristics have been caused by national income. But suppose we are asking How much would it help if households were given more income? Then clearly we should be holding these other things constant. In that case, when estimated across countries, the coefficient of wellbeing on log income is 0.33 (se = 0.07).[17] This is conveniently similar to the coefficient across individuals.[18]

[17] See Chapter 8. With nothing held constant, the coefficient is rather over double this.
[18] This does not logically prove that social comparisons are unimportant. For example, it might logically be the case that $W_{ic} = b_1(log\, Y_{ic} - log\, \bar{Y}_c) + b_2(log\, \bar{Y}_c - log\, \bar{Y}_{world})$ and $b_1 = b_2$.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

## Time-Series for Countries

We turn now to the second part of the Easterlin paradox. His claim is that, as time passes, higher national income does not produce higher national wellbeing. In the cross-section of individuals we have found that

$$W = a\log Y + \text{etc.}$$

One would therefore expect that over time at the country level

$$\Delta\overline{W} = a\Delta\overline{\log Y} + \text{etc.}$$

But Easterlin says it does not.

So what are the facts? The most striking fact supporting Easterlin is the story of the United States. As Figure 13.3 shows, average happiness has not increased there since the 1950s, despite rapid economic growth that was widely shared at least till the 1970s.

However, Figure 13.3 does not prove that in the United States higher income did not improve wellbeing. It might have done so, with other factors offsetting this effect. In any case, the experience of one country can prove little. So the first step is to see whether countries with higher long-term economic growth have had higher growth in wellbeing. It is important to stress the word long-term because there is no doubt that happiness rises in booms and falls in slumps (for all kinds of reasons we shall come to). But do countries with high long-term growth do better in terms of happiness?

### Long-term growth and wellbeing

To investigate this, the best available long time-series is for European countries, where the Eurobarometer survey has been conducted regularly in many countries since the

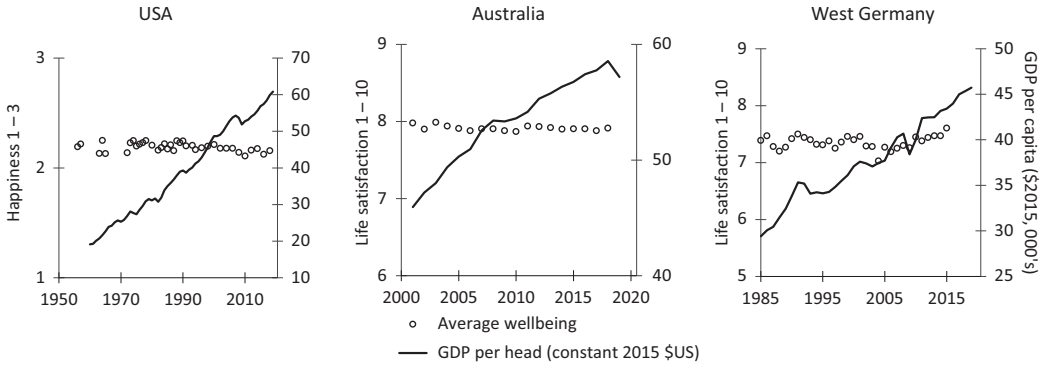


**Figure 13.3** Happiness and income per head in the United States, Australia and West Germany
*Sources*: United States: AIPO, NORC and GSS grafted together using overlapping data; Australia: HILDA; Germany: https://tinyurl.com/3jyjaw4x, based on SOEP.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

early 1970s.[19] To investigate the effect of long-term growth, it is of course necessary to control for booms and slumps (which drive wellbeing up and then down). This is done in the following equation, by including percentage unemployed ($u$) and annual percentage inflation ($\pi$), as well as GDP per head and country fixed effects.

$$\bar{W}_{ct} = a_1 + \underset{(.17)}{0.29}\ log\, \bar{Y}_{ct} - \underset{(.02)}{0.06}\, u - \underset{(.004)}{0.007}\, \pi + f_c.$$

As the equation shows, income is estimated to have a positive effect but with a huge standard error.[20] This is typical of multi-country time-series – reflecting in part the very different growth rates of wellbeing in different countries with similar rates of income growth.[21] Thus, the conclusion on the second part of the Easterlin hypothesis has to be 'sometimes yes and sometimes no'. But for citizens of the United States there is a particular challenge – of why they are no happier on average than people were in the 1950s.

## The Role of Income Comparisons and Adaptation

There is one obvious reason why national increases in income over time might produce lower changes in wellbeing than individual increases in income at one point in time. It is **social comparisons**.

Suppose that each of us has a comparator group with whom we compare our incomes and that much of our concern about income is focused on our **relative income** rather than our absolute income. Then a person's wellbeing depends positively on her well-being but negatively on the income of her comparators. The relationship could be[22]

$$W_i = a_1 \log Y_i - a_2 \log \bar{Y}_i + \text{etc.} \qquad (a_1, a_2 > 0) \tag{5}$$

where $\bar{Y}_i$ is the average income in the person's comparator group. This can also be written as the sum of the effects of absolute income ($Y_i$) and of relative income ($Y_i/\bar{Y}_i$):

$$W_i = (a_1 - a_2) \underset{\text{Absolute income}}{\underbrace{\log Y_i}} + a_2 \underset{\text{Relative income}}{\underbrace{\log \left(\frac{Y_i}{\bar{Y}_i}\right)}} + \text{etc.}$$

[19] The countries covered below are Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain, Sweden and the UK.

[20] Analysis by Ekaterina Oparina. Splitting the coefficient into trend and cycle makes no difference once unemployment is included. And including time dummies as well reduces 0.29–0.08. An identical equation to the one in the text, but using World Values Survey for all waves 1981–2019, gives the following coefficients: −0.11 (.20), −0.04 (.01), −0.02 (.00).

[21] An analogous way to study the data is to estimate for each country the trend in wellbeing and the trend in log GDP per head and then regress one on the other (across countries). This is the approach in Easterlin and O'Connor (2020). See Annex 13.1.

[22] In this formulation $\bar{Y}$ has no effect on individual $i$'s choice of hours. There is some evidence that it does (see A. E. Clark et al. [2008]), which raises an even bigger policy problem – that people are working harder simply because others are doing so (the rat-race).

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

**Table 13.4** Effect of log own income and log comparator income on life satisfaction (0–10) (pooled cross-sections)

| | Britain | Germany | Australia | United States |
|---|---|---|---|---|
| **Own income** | **0.16** (.01) | **0.26** (.01) | **0.16** (.01) | **0.31** (.01) |
| **Comparator income** | **−0.23** (.07) | **−0.25** (.04) | **−0.17** (.06) | **−0.19** (.03) |

*Source*: A. E. Clark et al. (2018) Table 2.3; Britain, Understanding Society (1996–2014); Germany, SOEP (1984–2015); Australia, HILDA (2001–2015); United States, BRFSS (2006–14)
*Note*: Standard errors in brackets. 'controls' include all those in Figure 8.1.

If $a_2$ is substantial, comparators' income is a substantial force reducing our happiness.

So what is the evidence on the effect of comparator income? In the great majority of studies, it is negative and large.[23] But the findings sometimes depend on what group people are assumed to compare themselves with – very local (such as neighbours or colleagues) or people from the same region, age group and gender.[24] In the cross-sectional study reported in Table 13.1, we used the latter definition. The results are shown in Table 13.4.

In these countries, the effects of comparator income are negative and large. Thus, the effects of relative income are positive and large. Meanwhile, the effects of absolute average income, which is the sum of the two rows in Table 13.4, are small. If these numbers are anywhere near right, they provide an obvious explanation of the Easterlin paradox:

- When an individual has a higher income, holding $\bar{Y}$ constant, she is happier. This is mainly because her relative income is higher.
- But when the whole society becomes richer, $\bar{Y}$ rises and relative incomes do not change. (Some people may go up in relative terms and others down but the average of relative income remains constant.) So at the level of society the only effect of economic growth is the weaker effect of absolute income.

There is much other evidence that people care about relative income as well as absolute income. Some of it comes from neuroscience, led by Armin Falk of the University of Bonn.[25] His team organised an experiment where participants had to undertake a task while undergoing a functional MRI measurement of brain activity in the brain's reward centre, the ventral striatum. Those who successfully completed the task were given a financial reward, which was varied randomly. They were also told of the reward, if any, received by the person with whom they were paired. The findings were remarkable. The measure of activity in the ventral striatum increased by 0.92

[23] A. E. Clark et al. (2008).
[24] In local comparisons, the coefficient is sometimes positive. This is probably because neighbours' incomes are taken as a forecast of our own incomes ('light at the end of the tunnel'). Graham (2012); Ifcher et al. (2018).
[25] Dohmen et al. (2011).

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

**Table 13.5** Effect of own income, comparator income and own lagged income on life satisfaction (0–10) (with fixed effects)

| | Britain | Germany | Australia |
|---|---|---|---|
| Log own income | **0.06** (.01) | **0.19** (.01) | **0.06** (.01) |
| Log comparator income | **−0.09** (.06) | **−0.12** (.04) | **0.01** (.04) |
| Log previous 3 yrs' income | **−0.02** (.02) | **−0.08** (.01) | **−0.01** (0.01) |

*Source*: A. E. Clark et al. (2018) Table 2.4; Britain, Understanding Society (1996–2014); Germany, SOEP (1984–2015); Australia, HILDA (2001–2015).
*Note*: Standard errors in brackets. 'Controls' include all those in Figure 8.1. Note that comparator income has little movement except for people who move regions.

units for every €100 they themselves received and fell by 0.67 units for every €100 their pair received. So, relative income had double the effect of absolute income.

In another ingenious experiment, David Card (another winner of the Nobel Prize) and his colleagues examined the effect of knowing the incomes of your colleagues. It happened that the University of California, where he works, had recently put all faculty salaries online. But most people did not know about it. So Card informed a random selection of the faculty members that these data existed. He also measured the wellbeing of the treatment and control group before and after he did this. Those who learned about colleagues' salaries became on average less satisfied.[26] So relative income clearly matters.[27]

## Adaptation

However, some psychologists favour a different explanation of the Easterlin paradox – **adaptation**. According to this, people may enjoy an increase in income for a time but then they get used to the higher income and their wellbeing reverts to its former level. Unfortunately, not many studies have tested the effect of social comparisons and adaptation simultaneously. In Table 13.5 we do this, using the same panel data that was used cross-sectionally in Table 13.4. But this time we include a fixed effect for each individual, and we include not only social comparison income but also the lagged values of own income:

$$W_{it} = \alpha_1 \log Y_{it} - \alpha_2 \log \bar{Y}_{it} + \alpha_3 \log Y_{i,t-1} + a_4 \log Y_{i,t-2} + \text{etc} + v_t + f_i + e_{it}. \quad (6)$$

When this fixed effects model is run, the effect of one's own income is lower than in cross-section (as discussed earlier). But the negative effect of comparator income is

[26] Card et al. (2012).
[27] Similarly, Perez-Truglia (2020) found that, since Norwegian tax records became publicly available in 2001, the gap in life satisfaction between rich and poor people increased by 21%. Many psychologists, like Nicholas Chater and Gordon D. A. Brown, argue that people care more about their position in the rank order of income than they care about relative income. The broad implications of both views are similar.

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

much more substantial than of lagged income from the previous three years. In any case, adaptation cannot be the main reason for the Easterlin paradox: in any community, most richer people have always been richer and poorer people have been poorer, and yet the richer people are on average happier.[28]

## Policy Implications of Income Comparisons

The analysis in this chapter has major implications. First, as we have said, diminishing marginal utility is a powerful argument for the redistribution of income to the point where it ceases to raise average wellbeing. This point has been understood for many years. But the role of social comparisons introduces a completely new argument in favour of taxation (not based on the diminishing marginal utility of income nor on the need to fund public goods).[29]

To see this, assume a world in which everybody is equal. It will still be the case that when someone works longer and therefore earns more, she increases the income level against which everyone else compares their income. This is a **negative externality**, and it leads to an inefficient scale of work unless something is done. What level of tax would suffice? Let us suppose that there are (N + 1) identical individuals each paid (for simplicity) a salary of one unit per hour worked – meaning that a person's earnings equal the hours she works. So suppose that individual wellbeing depends additively on own log income, log comparator income and the cost of work effort. Then, if there are no taxes,

$$W_i = \alpha_1 \log H_i - a_2 \log \bar{H} - C(H_i) \quad (C' > 0) \tag{7}$$

where $\bar{H}$ is the average income of the other N people, and the last term is the psychic cost of working $H_i$ hours.[30]

When anyone works longer to improve their own income, they also raise the average income $(\bar{H})$ against which others compare their income. In the absence of a tax, the individual chooses $H_i$ to maximise $W_i$, ignoring the effect this has on $\bar{H}$. So the **individual optimum** in the absence of a tax is found by setting $dW/dH_i$ equal to zero, which gives

$$\frac{\alpha_1}{H_i} = \frac{dC}{dH}(H_i) \tag{8}$$

The value of marginal income is equated to the psychic cost of obtaining the marginal income.

But the **social optimum** would also allow for the negative effect of increased $H_i$ on the wellbeing of the other N people. When $H_i$ goes up by one unit, $\bar{H}$ goes up by 1/N.

[28] This is not because their position in the distribution is temporary – it is mainly 'permanent'.
[29] Layard (1980, 2006).
[30] If the leisure time of comparators also had a negative effect, the distortion would be less. But there is evidence that people do not compare their leisure with that of others. See Solnick and Hemenway (1998).

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

This in turn reduces the wellbeing of each of the other N people by $(\alpha_2/\bar{H}_i)(1/\mathrm{N})$. But there are N people who are affected in this way, so the total loss of wellbeing is $(\alpha_2/\bar{H}_i)$. The **social optimum** takes this into account and is therefore given by[31]

$$\frac{a_1 - a_2}{\bar{H}_i} = \frac{dC}{dH}\ (\bar{H}_i). \tag{9}$$

This implies fewer hours of work.

So what is the optimum corrective tax?[32] Suppose it is structured as a linear tax with a constant marginal tax rate ($t$) and the yield is returned to the workers in lump-sum form. Then the **individual optimum** is now

$$\frac{\alpha_1(1-t)}{H_i} = \frac{dC}{dH}\ (H_i). \tag{10}$$

We now want to find that value of $t$ that ensures that the individual optimum is the same as the social optimum. This requires

$$\alpha_1(1-t) = \alpha_1 - \alpha_2.$$

So the optimum marginal tax rate is $\alpha_2/\alpha_1$ – the ratio of the pollution effect to the own income effect. Given the estimated values of $\alpha_2$, this could justify quite high marginal tax rates on grounds of efficiency. At the very least, this externality argument should be put against the traditional argument that taxation has an 'excess burden' – because it inefficiently discourages work.[33] If when people earn more they are imposing a cost on others, this should be taken into account in any argument about what is efficient.

Similarly, with cost-benefit analysis. If a project is paid for out of higher taxes paid by everybody, any loss of wellbeing from the tax will be partially softened by the fact that everybody else is losing as well.

It is sometimes said that this argument is less forceful if people only compare themselves with small numbers of other people. But, as our reasoning earlier about N showed, the argument applies whatever the size of the group.[34]

Another important implication of social comparisons is for us as individuals. People who make fewer comparisons are on average happier.[35] So we should train our tastes, as far as we can, to reduce $\alpha_2$. Libertarians argue that, if we don't do this, that is our lookout, and governments should ignore human failings. But this approach is inconsistent with an evidence-based approach to ethical questions.

[31] Another way to arrive at this conclusion is simply to find the level of $H$, which, if everybody worked that long, would yield the social optimum. Differentiating equation (7) by $H$ and setting $dW/dH = 0$ yields (9) straight away.

[32] Boskin and Sheshinski (1978). They solve this problem for a world in which wage rates differ between people.

[33] In the economic jargon, the 'excess burden' compares the cost of the tax with the cost of a lump sum tax. The excess is due to the effect of the tax in making people substitute income for leisure. See Layard and Walters (1978) p. 87.

[34] This assumes that $\alpha_2$ is independent of N.

[35] For example, White et al. (2006).

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

A second argument in support of corrective taxation comes from unforeseen adaptation. If people work harder to increase their income but overestimate the effects of this on their own wellbeing,[36] that is another reason why some marginal taxation could be good for efficiency.

## Economic Fluctuations

We turn finally to fluctuations of income over the business cycle. What is certain is that wellbeing rises in booms and falls in slumps. Two processes are at work here. The first is **adaptation**. Morale is higher when income is high relative to previous income and low in recessions (when the reverse applies). And the second is **loss-aversion**. This time we are talking about ex post loss aversion (not the ex ante loss aversion that affects decisions). There is powerful evidence that the loss of wellbeing when income falls by a given amount is roughly double the increase in wellbeing when income rises by the same amount.[37] This has profound implications.

- First, it helps to explain the weak long-term relationship between income and wellbeing, since the years of income decline have such strong negative effects.
- Second, in terms of policy, it argues strongly for the importance of economic stability. In his famous Presidential Address to the American Economic Association, the neo-classical economist Robert Lucas argued that economic cycles were unimportant compared with the rate of long-term economic growth. Cycles could therefore be tolerated if they increased long-term growth. The implication of wellbeing research is the opposite: higher long-term growth ought not to be pursued if it leads to economic instability. Humans like stability and it is the job of policy to provide it.

## Conclusions

- The Easterlin paradox states that
  (1) in a given context richer people are on average happier than poorer people,
  (2) but over time greater national income per head does not cause greater national happiness.
- Statement (1) is certainly true. We reviewed a mass of evidence and concluded that, as a benchmark, a unit increase in log income raises wellbeing by 0.3 points (out of 10). The share of the within country variance in wellbeing explained by income inequality is 3% or less. So income is in no sense a proxy for wellbeing.
- Across countries, the effect of a unit change in log income per capita (other things equal) is also around 0.3 points of wellbeing.

[36] Loewenstein, O'Donoghue and Rabin (2003).
[37] One such study uses country data (De Neve et al. [2018]). Another uses individual data (Boyce et al. [2013]).

https://doi.org/10.1017/9781009298957.017 Published online by Cambridge University Press

- Over time, wellbeing has increased with income in some countries but not in others. Thus statement (2) is still the subject of ongoing research.
- From direct studies on individual data, it is clear that in most cases a rise in other people's income reduces your own wellbeing. This means that the effect of one person's income on that individual's wellbeing overestimates the effect of economic growth on the overall wellbeing of society.
- From a policy point of view, income comparisons mean that when a person earns more she imposes a cost on other people. This is a negative externality and one way to control it would be by corrective taxation. If the externality is as large as our estimates suggest, this could mean that quite high rates of marginal tax are efficient.
- The final issue is economic fluctuations. Wellbeing rises in booms and falls in slumps. One important reason is loss-aversion – people dislike a loss of income twice as much as they like a gain of equal size (both likes and dislikes being measured in units of ex post wellbeing). This may be a partial explanation of the Easterlin paradox. In terms of policy, it means that economic stability is enormously important and ought not to be sacrificed in pursuit of small increases in long-term economic growth.

## Questions for discussion

(1) Are the estimated effects of income plausible?
(2) Why do they not differ between rich and poor countries?
(3) Are national time-series estimates of the effect of income consistent with within-country estimates based on individuals?
(4) How important are social comparisons and what are their policy implications?
(5) How important is adaptation?

## Further Reading

Clark, A. E., et al. (2008). Relative income, happiness, and utility: An explanation for the Easterlin paradox and other puzzles. *Journal of Economic Literature*, 46(1), 95–144.

Clark, A. E., et al. (2018). *The Origins of Happiness: The Science of Wellbeing over the Life Course.* Princeton University Press.

Easterlin, R. A. (1974). Does economic growth improve the human lot? Some empirical evidence. In *Nations and Households in Economic Growth* (pp. 89–125). Academic Press.

Layard, R. (2006). Happiness and public policy: A challenge to the profession. *Economic Journal*, 116 (March), C24–C33.