


LIFE ISN'T POINTLESS



3 How Our Behaviour Affects Our Wellbeing

To a psychologist it is self-evident that people are neither fully rational nor completely selfish, and that their tastes are anything but stable.

Daniel Kahneman

Individual Responsibility versus Collective Action

If social welfare is the objective – the thing that needs to be maximised – the next issue is Who is going to make that happen?

In traditional economics, a huge reliance is placed on the unfettered choices of individuals. As Adam Smith argued in the *Wealth of Nations*, people may be selfish; but, in order to get what they want, they do best by supplying what others need. The result of this process of voluntary exchange is an **efficient** situation where (on four assumptions) it would be impossible to make anyone better off without making someone else worse off. These assumptions are:

- no monopoly or oligopolies,
- no public goods (like roads or parks),
- full information for all,
- no one affects anyone else except through voluntary exchange (i.e. there are no ‘externalities’).

The classic example of a negative ‘**externality**’ is air pollution. The pollution just comes my way whether I agree to it or not. So an externality is something that just happens to me without my choosing it. Psychologists and sociologists will immediately realise that externalities comprise a huge fraction of everything that happens – from how we learn our values to whether we are robbed.¹ Coping with externality is essential if we want an efficient society.

There is also another problem with pure laissez faire (where individual action is unfettered). This is the issue of **equity**. Even if the framework is efficient, it may leave

¹ There is no problem with so-called pecuniary externalities, where my action affects market prices enjoyed by others. Layard and Walters (1978) p. 226.

some people much less happy than others. People differ hugely in their talents, in the wealth they inherit and, most importantly, in their inner capacity to live an enjoyable life.

Thus pure *laissez-faire* is neither efficient (the four assumptions are not satisfied) nor is it equitable. So we need both society and the state to be active in promoting the wellbeing of the people and the prevention of misery. And for this they need wellbeing science.

There is also a further reason. People themselves differ hugely from what is assumed in traditional economics. On the one hand, they are much less good at promoting their own best interests than economists assume. And, on the other hand, they do often care a lot about the wellbeing of others. So we have to begin this book with a fundamental look at some of the basics of human behaviour.

In this chapter we shall investigate two questions.

- To what extent do people pursue their own best interests?
- In what ways are people affected by others' behaviour beyond voluntary exchange – in other words, what are key examples of externality?

Human Decision-Making

Clearly, people do to some extent pursue their own best interests. That is one reason why the human race has survived. People are 'attracted' to things that are good for survival – and feel happier when they achieve them (eating, mating etc.). And they 'avoid' things which are bad for their survival. This 'approach/avoidance' system is a basic part of human nature. But things are more complicated than that.

In 2002, the Nobel prize for economics was awarded for the first time to a psychologist. Some would say 'About time', since economics is based on a psychological theory about how people behave. But Daniel Kahneman's message was different from that of the economists. As Kahneman put it, economists typically assume that people 'are rational, selfish and their tastes do not change'.² After decades of research with his friend Amos Tversky, Kahneman believed otherwise.

Together they founded a new subject – called **behavioural economics** by some and **behavioural science** by others. Many others followed them, especially the economist/psychologist Richard Thaler (another Nobel laureate) and the lawyer Cass Sunstein. In their book *Nudge*,³ they distinguished sharply between the people of traditional economic theory (who they called Econs) and real people (who they called Humans).

The two central issues are these: how effectively do people pursue their own wellbeing, and how do they impact on the wellbeing of others? There are at least five major ways in which humans do not maximise their own wellbeing:

- Addiction and lack of self-control,
- Unforeseen adaptation,

² Kahneman (2011) p. 269. This whole chapter draws heavily on this book.

³ Thaler and Sunstein (2008). See also Thaler (2015).

- Framing,
- Loss-aversion,
- Unselfishness.

Addiction and self-control

People regularly decide to give up smoking, alcohol, drugs or gambling but then fail to do so. Frequently, they decide to do it next year, and when next year comes they decide to do it the following year. This is a case of ‘dynamic inconsistency’. It is inconsistent because in the first year they have planned on the basis of their preferences in the first year; but, when the second year comes, their preferences for the second year have changed. More generally, as people look forward, many of them give a very high priority to the present over all ensuing periods: they use a very high discount rate between now and next year. But they use a much lower discount rate when they compare next year with the one that follows, and so on – and this is why they plan to give up next year rather than this year. Interestingly, neuroscientists have shown that decisions between this year and next are made in more primitive parts of the brain than decisions between more distant future years.⁴

Unforeseen adaptation/habituation

Another source of bad decisions is adaptation that is not foreseen. Human beings are generally pretty adaptable. They adapt to hardship – and they habituate to good experiences. In other words, the initial impact of a change on wellbeing eventually wears off, either partially or totally.⁵ As we shall see in this book, **people adapt** (partly at least) to divorce, disability and poverty, and they also adapt to new houses, a new car, a new partner or a new child.⁶

If adaptation is **total**, the only way to change your wellbeing permanently is to experience continual change. Thus in ‘total adaptation’ we might have

$$W_t = b (X_t - X_{t-1}) \quad (1)$$

where X is anything that matters to you. So if you like X , your mood is higher when X is increasing. But, once X stops increasing, you return to your previous level of wellbeing. In other words, there is a ‘set point’ of wellbeing – neither good changes nor bad ones have any permanent effect.⁷

But if in fact we follow people through their lives, we find that their wellbeing often fluctuates a lot from one decade to another.⁸ So, total adaptation is relatively

⁴ Laibson (1998) and Rabin (1998).

⁵ Addiction often arises from this source. At first the substance raises wellbeing, but over time, you need more and more of it to get the same good effect.

⁶ See later chapters and A.E. Clark et al. (2018).

⁷ Of course, if you can keep X growing every year, you will be better off but you are then on a ‘**hedonic treadmill**’.

⁸ Headey (2006).

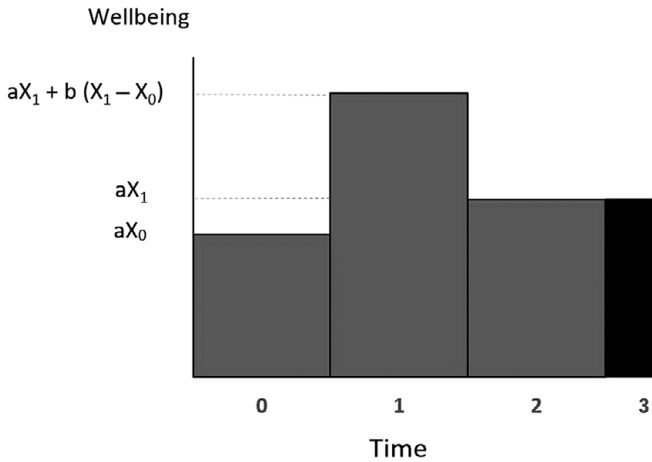


Figure 3.1 Partial adaptation of wellbeing to changes in X

rare – poorer people are, for example, less happy than rich people (other things being equal). But they also suffer less than you might imagine if you thought about how you would feel next year if you suddenly became that poor.

For the most common case in real life is ‘**partial adaptation**’. This is very common. When you move house, your next new house makes you feel great at first, but then gradually you take it for granted (you reduce your attention to it) and you return towards how you felt before the move.⁹ The same happens with a new car or a rise in income. A typical case of partial adaptation is

$$W_t = aX_t + b(X_t - X_{t-1}) \quad (2)$$

So suppose we start from a position where X has equalled X_0 for some time, with

$$W_0 = aX_0$$

Then X rises to X_1 and stays there from then on. In response, wellbeing in the first time period is

$$W_1 = aX_1 + b(X_1 - X_0)$$

But in the following period, there is no further stimulating increase in X . So

$$W_2 = aX_1$$

This sequence is shown in the Figure 3.1, where X is taken to be desirable.¹⁰ The key point is that the degree of adaptation is measured by the ratio of b to a . The higher b is relative to a the more adaptation occurs, and the further W moves back to its original

⁹ Odermatt and Stutzer (2019).

¹⁰ If more lagged terms appeared in equation (2), the decline to aX_1 would be more gradual.

position after its initial change. This kind of pattern applies to the effects of finding a partner, losing a partner, having children and many basic experiences.

If one looks deeper into the reasons for adaptation, it clearly has much to do with **attention**. When you attend to things, they become very important. When you do not, they become less so. As Kahneman puts it, ‘nothing in life is as important as you think it is when you are thinking about it’. Things that are very salient and obtrusive like money can easily grab people’s attention, although for much of the time most people don’t actually think about them. For example, most people think that people in California are happier than people in the mid-west region of the United States. This is because, when they make the assessment, they are thinking about the difference in the weather. But in fact there is no difference in happiness because neither group actually spend that much time thinking about the weather.¹¹ Nor after a while do most people spend much time thinking about their car, unless someone draws their attention to it.

So what are the implications of adaptation for the wellbeing of the individual? The key question is whether or not individuals foresee their adaptation. If people foresee the adaptation, the decisions they take will be optimal. But in many cases, people clearly overestimate the gain in wellbeing they will eventually get from a change – a higher income, a new car, a new house or even a new partner. In such cases, people will put too much effort into acquiring the new situation, and in some cases they will come to regret the effort they made.¹² The Harvard psychologist, Daniel Gilbert, has called this a case of ‘miswanting’.¹³

Framing

People’s behaviour is also hugely affected by how a decision is presented to them. Every advertiser knows this – to sell a product you associate it with something attractive, however irrelevant. And you make sure people hear about it as frequently as possible.¹⁴ Clearly advertising affects people’s tastes. In particular it often makes them feel they need something that they did not formerly need. So, not surprisingly, there is some evidence that advertising reduces human happiness.¹⁵

But the clearest case of framing is the effect of the ‘default’. For example, there are two ways in which people can be presented with the decision on whether to save for old age through a pension plan:

- (1) you are asked if you want to opt in to the plan
- (2) you are automatically enrolled unless you opt out (‘auto-enrolment’).

This can make a huge difference to what you do. For example, four US companies switched from opting in to opting out.¹⁶ Under opting in, only 25–43% joined the

¹¹ Schadke and Kahneman (1998); Loewenstein, O’Donoghue and Rabin (2003).

¹² Loewenstein and Schadke (1999) p. 90. ¹³ Gilbert and Wilson (2000).

¹⁴ On the power of mere exposure, see Zajonc (1968). ¹⁵ Michel et al. (2019).

¹⁶ Choi et al. (2006).

company's pension plan, but once enrolment was automatic (unless you opted out) 86–96% joined the plan. This is extraordinary. A huge decision depends simply on whether you have to push a button or not. It is clear that, on a matter as important as this, people have no clear preferences, even though most economists would assume they do. Instead, many people let the framing of the decision determine what they do.

Similar behaviour is found in relation to how you can pre-donate your organs after death. In countries where you have to opt in to donate, donation rates can be as low as 4%; but, when you have to opt out, the rates approach 100%.¹⁷ In all these cases, opting in appears to involve a psychic cost, but this cost disappears if someone else does it for you.

Because framing is so important, anyone wanting to influence people's behaviour must frame their approach to people carefully. This applies to government as much as to everyone else. This idea is central to the policy of **Nudge**, advocated by Richard Thaler and Cass Sunstein. They argue that, if we know how we want people to behave, we should if possible, make them do so by framing their choices right, rather than by dictating to them. This approach is known as '**libertarian paternalism**'.¹⁸

There are three key principles in framing the choices that people are offered:

- (i) Keep the message simple.
- (ii) If you want someone to do something, point out that lots of other people are doing it.
- (iii) Associate doing it with something positive rather than something negative.

The following experiment illustrates the last point. Two randomly selected groups (A and B) were offered the following choices:

Group A	Would you accept a gamble that offers a 10% chance to win \$95 and a 90% chance to lose \$5
Group B	Would you pay \$5 to participate in a lottery that offers a 10% chance to win \$100 and a 90% chance to win nothing

The majority of Group A refused to accept the gamble, while the majority of Group B agreed to participate in the lottery. And yet both these offers were identical (just check it out). The reason for the different answers was that Group A heard the word 'losing' while Group B did not.

Loss-aversion and the endowment effect

This leads directly to another key human characteristic: loss-aversion. Even for small things, people hate losing them by more than they would value gaining them. This is a

¹⁷ Johnston and Goldstein (2003) p. 1338

¹⁸ They particularly favour nudges where decisions are big, long-term and complicated.

deep psychological trait in all animals – their strongest impulse is to hold on to what they have. Just try removing food from your cat.

Richard Thaler calls this the ‘**endowment effect**’. So here is a remarkably simple experiment.¹⁹ Students are randomly divided into two groups. The first group is shown a particular type of mug and asked how much they would be willing to pay for it. The average answer is \$3.50. The other group are given the same mug for free. After a while they are asked how much money they would need to be paid in order to get them to part with the mug. The average amount is \$7. So the same mug is worth \$7 once you have it, but you are only willing to pay half that in order to get it. Thus there is no simple answer about the value of the mug, even to the same person: it depends whether she already owns it.²⁰ The value depends on the ‘**reference point**’.

Another clear example of loss-aversion is people’s response to uncertainty, where Kahneman was awarded the Nobel prize for what he and Tversky called ‘prospect theory’. So here is one simple experiment. You are asked whether to accept the following gamble:

- 50% chance of winning \$150 and
- 50% chance of losing \$100.

Someone who accepts this gamble can, on average, expect to win \$25. Yet the majority of people decline the offer.²¹ This cannot be explained by assuming that there is a smooth relationship between happiness and income, regardless of what the person’s income was until now.²² Instead, when people value their future income, they are comparing it with their present income. And they dislike losses more than they like gains. In other words, as Figure 3.2 shows, if a person started with an income of y_0 , she would lose more happiness if she lost \$100 than she would gain if she won \$150.

Loss-aversion goes way beyond our attitude to income, and it affects behaviour in all sorts of ways.²³ For example, professional golfers are more likely to sink a putt if it is needed to avoid a bogey (a loss) than if the same shot would give them a birdie (a gain). Taxi drivers work longer hours on bad days, even though there is less work. Homeowners insist on selling their houses for a higher price when it involves a loss. And investors are more likely to sell stocks that have increased in value (since they were bought) than to sell stocks that have gone down: they hate to ‘realise’ losses. None of these attempts to avoid losses would be in the interest of a rational maximiser of expected outcomes.²⁴

¹⁹ Kahneman et al. (1990).

²⁰ Another striking example is when students are randomly given either a coffee mug or a chocolate bar (both of which cost the same). They are then offered a chance to trade. One would expect that 50% would prefer what they did not have. In fact it was 18%. List (2003).

²¹ This is not of course true of traders who make hundreds of bets. For someone making this bet many times over would be almost certain to win a lot.

²² As Rabin (2000) shows, a person with a smooth utility-of-income function who rejected a 50:50 chance to win \$200 or lose \$100 would also have to reject a 50:50 chance to win \$20,000 or lose \$200. This shows that a smooth utility-of-income function cannot be the reason for loss-aversion.

²³ Camerer et al. (1997); Odean (1998); Genesove and Mayer (2001); Kahneman (2011).

²⁴ This is one example of separate ‘mental accounting’ – we judge things in compartments, not in their overall context.




Figure 3.2 Loss-aversion

Many studies have attempted to measure the scale of loss-aversion, as measured by the slope of the wellbeing/income relationship for losses relative to that for gains. The typical finding is a ratio of 2:1. For example, suppose there was

- a 50% chance of winning \$ x and
- a 50% chance of losing \$100

and you were asked what value of x would make you just willing to undertake the gamble. The typical answer is \$200.²⁵

A similar size of effect can be found at the level of the **whole economy** – this time in relation to actual wellbeing experienced rather than anticipated wellbeing. When incomes fall in a slump, the fall in national happiness is double the increase in happiness that occurs when incomes rise in a boom.²⁶ Once again the ratio is 2:1. Loss-aversion explains many forms of behaviour that standard economics cannot explain but that wise people have always noticed. One of these is wage stickiness in the downward direction. As Keynes pointed out, people fiercely resist a cut in the money value of their wage. If prices are falling enough, this may appear irrational since a person might gain in real income even though their money wage fell. But people do not see it that way, and a combination of loss-aversion and ‘money illusion’ stops wages falling and exacerbates the slump. This simple observation is central to Keynesian economics, but it is inconsistent with pure rationality.

²⁵ See, for example, Kahneman (2011).

²⁶ At least double – see De Neve and Ward (2017).




Figure 3.3 Loss aversion with diminished sensitivity
Source: Adapted from Kahneman (2011)

More generally, it is loss-aversion that makes many economic reforms so difficult to achieve. Even if 95% of the population would gain, a reform will often be defeated because of the 5% who lose. Examples abound in every country. In Britain in the 1930s, mass unemployment caused little political protest until the government reformed the unemployment benefit system in a way that benefitted most unemployed people but hurt a few.²⁷

Before leaving loss-aversion, let's note one further fact: **diminished sensitivity**. One aspect of this is standard in economics – as you gain more and more, you value each additional gain less and less. But the same is true of losses – as you lose more and more, each extra loss causes less and less additional pain. For example, suppose once again that we look at gains and losses of money. The corresponding changes in wellbeing are shown in Figure 3.3.

This has remarkable results. Suppose people have to choose between the following alternatives:

- a 50% chance of losing \$100
- the certainty of losing \$50.

They will prefer the gamble since it offers at least a 50% chance of losing nothing. To retain the hope of avoiding loss, decision-makers worldwide are prone to risk a lot, even though it might be more rational to accept a much smaller loss and move on.²⁸

The combination of loss-aversion and diminished sensitivity constitute what Kahneman and Tversky called 'prospect theory' and it was for this that Kahneman won the Nobel prize (Tversky having sadly died). Each prize-winner receives a personalised drawing, and the drawing on Kahneman's certificate was a more artistic version of Figure 3.3.

²⁷ Runciman (1966).

²⁸ There is also a well-documented tendency to exaggerate the **probability** of small risks, which is how insurers make their money. Tversky and Kahneman (1992).

Unselfishness

Finally, there is a key aspect of human behaviour that does not often appear in standard economics: **unselfishness**. People regularly do things for other people whom they will never see again and who could never return the favour. They return lost wallets to their owners and they are more likely to return wallets the more money they contain.²⁹ People give money in secret to charity; they dive into rivers to rescue drowning strangers. This behaviour can only be explained by the presence of a moral sense, partly learned and partly in our genes.³⁰ People of course vary in the strength of their moral sense. So what can account for this dual aspect of our human nature – the mixture of the selfish and the unselfish? You might think that at the individual level people who were more selfish would end up doing better in the Darwinian struggle for survival. And, if so, the genes for unselfishness would become less and less common. But this doesn't happen. Two things sustain the prevalence of cooperation.

The first is **social sanctions**. Your fellow citizens know that for many purposes cooperation is more efficient than competition. So they punish people who won't cooperate – such people get left out; people will not work with them. So, as the evidence shows, non-co-operators do worse on average than other humans.³¹ The same happens among chimpanzees. In an impressive recent experiment with chimpanzees, there were five times more co-operative acts than non-co-operative ones – and the non-co-operators were either shunned or physically punished.³² So one reason why people cooperate is to win the favour of others.

But there is a second reason: on average, pro-social behaviour makes us **feel good inside ourselves** – the 'warm glow'. In one experiment, students were randomly divided into two groups. One group were given money to spend on themselves, and the second group were given money to spend on other people. When asked about their happiness after the experiment, those who had given the money away were happier. This experiment has been repeated in four countries and a similar experiment with toddlers has produced the same results – they smile more.³³

Neuroscience confirms this mechanism. When playing the Prisoner's Dilemma game (when you can either 'co-operate' or not), people who co-operated showed more activity in the standard brain area where other positive rewards are processed.³⁴ And this was even before they knew whether the other player would co-operate. In other words, if you want to feel good, do good.

²⁹ Cohn et al. (2019). ³⁰ Ricard (2015).

³¹ For some evidence, see Schwartz (1970); and Lyubomirsky et al. (2005b). However, some cheats do also survive in the competitive struggle. Frank (1988) has a nice explanation of how in equilibrium both cheats and non-cheats can survive – by consorting with people like themselves. See also Fehr and Fischbacher (2003), who show how multiple equilibria are possible: if there are enough selfish people, even unselfish people will behave selfishly.

³² Suchak et al. (2016).

³³ Otake et al. (2006); Dunn et al. (2008); Aknin, L. B. et al. (2012); Aknin et al. (2013, 2015, 2019); Lane (2017). As Helliwell et al. (2018c) show on page 10, helping behaviour gives more pleasure to the helper when it is done for the sake of the other person than when it is done for self-oriented reasons.

³⁴ Rilling et al. (2002). See also Harbaugh et al. (2007).

Externalities

The conclusion so far is that Humans do not always maximise their own wellbeing. In cases where this produces a bad outcome, good advice, nudges and even government intervention might help them do better.³⁵ But there is another reason for government intervention that also arises from human nature: the role of **externality**. Let us look briefly at three forms of externality: the trustworthiness of others, the sources of our social norms and social comparisons.

The trustworthiness of others

We are hugely affected by the ethical character of those around us (see Figure 2.1). A libertarian might argue that we can choose who to have around us. But this is not quite right. Only some people can choose. If they choose the most trustworthy, the rest of the population end up interacting with the less trustworthy. What matters for the population as a whole is the total pool of characters they can interact with. So the virtues and vices of our fellow citizens are huge external influences on our wellbeing. This has always been given as one reason for public education.

The sources of our social norms

Our own tastes and norms are things we absorb from the outside world – mostly in childhood and adolescence. They affect what we enjoy – and how we treat others. Libertarians like the Nobel prize-winning economists Gary Becker and George Stiglitz argue that no one set of tastes is better than any other (in Latin, ‘de gustibus non disputandum est’).³⁶ Wellbeing science offers a different view on this. As we have seen, our preferences have huge effects on other people. Our genes certainly have some effect on our tastes and norms of behaviour. But we are also hugely affected by our own parents, our schools, our peers and society as a whole. One has only to think about how tastes in fashion or music change from decade to decade or our attitudes to gender, race and smoking.

Basic human needs may not change much. According to one school of thought, our social needs can be summarised under three headings: autonomy, competence and relatedness.³⁷

- **Autonomy.** We need a sense of agency – that we are in control of our lives.
- **Competence.** We need to feel able to do what is needed of us.

³⁵ Another issue is whether people do choose what they **think** will make them happy? Interesting research in this area shows that generally they do but not always. However, many of the contrary cases can be explained either by unselfishness or by people using a longer time horizon for their future happiness than the survey allows for (e.g., you take a less pleasant job to access a satisfying later career). See Benjamin et al. (2012); Fleurbaey and Schwandt (2015); Glaeser et al. (2016).

³⁶ Stigler and Becker (1977). ³⁷ Deci and Ryan (2000). Self-Determination Theory.

- **Relatedness.** We need to feel appreciated and to appreciate others – we need positive connections.

This may be true. But the specifics of what we desire depends hugely on the norms of the society in which we live.

Social comparisons

People are hugely affected by how they are doing compared with other people. They care about how their incomes compare with colleagues or how their children compare with other people's children. We discuss this further in Chapter 13. People also tend to imitate each other. This sort of herd behaviour is generally the result of safety-seeking behaviour among members of a herd.³⁸

Conclusions

In the traditional economic model, people have well-defined preferences and pursue them consistently and selfishly. According to this model, wellbeing is efficiently promoted, in a laissez-faire world, except for various problems. The biggest of these problems is 'externality', which is the way in which people affect other people without their agreement.

However, there are in fact many other problems. Humans are different from the traditional economic model in many ways.

- They often lack self-control (e.g., over drugs, alcohol and gambling), and what they plan for tomorrow they often fail to execute.
- They often do not realise how much they will adapt to change, and they put effort into new acquisitions that make less difference to them than they expect.
- They are hugely affected by how decisions are framed – for example, many will adopt a pension plan when it requires effort to opt out of it, but they would not choose it if they had to opt in to it.
- People are hugely loss-averse, which often makes desirable changes much more difficult.
- And, on the other side, they often help each other without expecting anything in return.

These complexities mean that government intervention or nudges are often needed to produce efficient outcomes. These arguments become even stronger when we consider the many pervasive 'externalities' affecting our experience:

- We benefit if we live in a trustworthy society.
- We get many of our norms and tastes from society.
- Because of rivalry, we lose wellbeing if others are more successful.

³⁸ Hamilton (1971).

Thus, the simple economic model explains a lot about the behaviour of companies and workers. But it is a very incomplete account of how people often behave. Clearly we need to understand why people behave as they do. But to know what behaviours should be encouraged, we also need to know how behaviour affects wellbeing. That is the central feature of the rest of this book.

Questions for discussion

- (1) If people fail to pursue their own best interest, should the government ever do more than nudge them?
- (2) What is the right mix of competition and cooperation?
- (3) Whose job is it to produce good citizens: (i) parents, (ii) schools, (iii) government, (iv) faith organisations, (v) others?
- (4) How useful is traditional economic theory for thinking about public policy? What are its main short-comings?

Further Reading

- Choi, J., Laibson, D., and Madrian, B. C. (2006). Saving for retirement on the path of least resistance. In E. McCaffrey and J. Slemrod (Eds.). *Behavioral Public Finance: Toward a New Agenda* (pp. 304–351). Russell Sage Foundation.
- Gilbert, D. T., and Wilson, T. D. (2000). Miswanting: Some problems in the forecasting of future affective states. In J. P. Forgas (Ed.). *Studies in Emotion and Social Interaction, Second Series. Feeling and Thinking: The Role of Affect in Social Cognition* (pp. 178–197). Cambridge University Press.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Allen Lane.
- Layard, R., and Walters, A. A. (1978). *Microeconomic Theory*. McGraw-Hill.
- Rilling, J. K., Gutman, D. A., Zeh, T. R., Pagnoni, G., Berns, G. S., and Kilts, C. D. (2002). A neural basis for social cooperation. *Neuron*, 35(2), 395–405.
- Thaler, R. H., and C. R. Sunstein (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press.