1 Assuming Rationality

Policies inspired by behavioural science have been implemented throughout the world for many years on an ad-hoc basis, but substantive efforts to create a broad behavioural public policy approach is a relatively recent endeavour. The main intellectual catalyst to these efforts was a series of publications written, in large part, by some of the world's leading behavioural economists in the first decade of the twenty-first century (Camerer *et al.*, 2003; Thaler and Sunstein, 2003; 2008). Those writings outlined conceptual frameworks for how behavioural economics can underpin mostly soft forms of paternalism or, in other words, non-mandatory behaviour change. Regarding policy influence, the non-mandatory emphasis within these policy frameworks was important, particularly in the Anglo-American world, where the political climate was set against further regulation and enforcement, at least according to government rhetoric (Behavioural Insights Team, 2010). Equally important in an era of austerity was the promise that many of these behavioural interventions would be financially inexpensive to implement.

In terms of creating a dedicated behavioural public policy unit, the British were the first to embrace this new approach to policy at the central government level (Halpern, 2015). In 2010, soon after becoming the prime minister, David Cameron established the Behavioural Insights Team with a view to recommending policy proposals informed by behavioural science. Similar initiatives have now been established, or at least considered, in several other countries, including Sweden (McDaid *et al.*, 2014), The Netherlands (National Institute for Public Health and the Environment, 2011), France (Oullier and Sauneron, 2010) and Denmark (Economist, 2012), as well as in much of the English-speaking world, most notably in the United States, where the

Social and Behavioural Sciences Team has tested and proposed a raft of policies, from promoting retirement security to improving college access, health coverage, health status and energy efficiency. The approach has a growing influence in the developing country context also, with, for example, the World Bank focussing on this topic in its 2015 World Development Report and establishing its own behavioural insights team, the Global Insights Initiative. Quite apart from the work undertaken by these new behavioural public policy units, the ad-hoc implementation of behavioural science-informed policy interventions around the world continues apace.

The ways in which behavioural economics and, more broadly, behavioural science might be used to inform policy is a theme that we will return to later in this book. However, in order to understand properly what behavioural public policy might have to offer, some knowledge of the most robust findings of behavioural economics is warranted. Chapter 2 will describe some of these findings, which were a response to pre-existing assumptions of rationality in mainstream economics. A good behavioural public policy analyst ought to have some knowledge of these origins.

THE ORIGINS OF ECONOMIC RATIONALITY

Economics as a formal field of study did not exist until the latter part of the eighteenth century. Up until that time, mathematicians assumed that when faced with a choice between two options, with each option offering a probability of winning an amount of money, a rational individual would choose that which offered the greatest expected value. For example, the expected value of a lottery that offers a 50% chance of winning \$100 and a 50% chance of winning \$0 is \$50 (i.e., 0.5*100 + 0.5*0), and the expected value of a lottery that offers a 25% chance of winning \$160 and a 75% chance of winning \$0 is \$40 (i.e., 0.25*160 + 0.75*0). Most seventeenth-century mathematicians would have therefore assumed that a rational individual would prefer the former lottery over the latter, if the prices of both lotteries were reasonable and equal.

In 1713, in a correspondence with the mathematician Pierre Rémond de Montmort, however, Nicolas Bernoulli questioned whether the assumption of expected value maximisation was always appropriate (Bernoulli, 1738; Zabell, 1990). Bernoulli devised an ingenious game, known as the St Petersburg paradox, to illustrate his point. The game involves the tossing of a fair coin and the participation of an individual who is informed that he will be paid a prize on the landing of the first head. Bernoulli used ducats as his currency of choice, but the example works with all denominations. Let us thus assume that the individual is informed that he will be paid $\$2^n$ for his participation in the game, where n is the number of tosses required for the first head to land. The individual is then asked how much he is willing to pay to play the game.

According to the principle of expected value maximisation, the individual should be willing to pay everything he owns, because the expected value of the game is infinite. To see this, note that the probability of the first head landing on the first toss of the coin is 0.5, in which case the individual is paid $\$2^1 = \$2$. The probability of the first head landing on the second toss is 0.25, in which case the individual is paid $\$2^2 = \$4$, and if the first head lands on the third toss, which it will with a probability of 0.125, the payoff is $\$2^3 = \$8$. The expected value of the game is calculated by summing all of the payoffs, weighted by their related probabilities of occurrence, associated with the first head landing on any particular toss of the coin. Numerically, this is given by:

$$\begin{aligned}\$(1/2)2 + \$(1/4)2^2 + \$(1/8)2^3 + \$(1/16)2^4 + \ldots \\ = \$1 + \$1 + \$1 + \$1 + \ldots = \$\infty.\end{aligned}$$

Bernoulli recognised that people are likely to be willing to pay only quite modest amounts of money to play the St Petersburg game. Indeed, Allais (1990) stated that the psychological value of the game is generally less than \$20, and even that may be an overstatement, with one study showing that most people are unwilling to pay more than \$4 (Schmeidler and Wakker, 1990).

A generation later, Nicolas Bernoulli's cousin, Daniel Bernoulli, proposed an alternative to expected value maximisation in order to

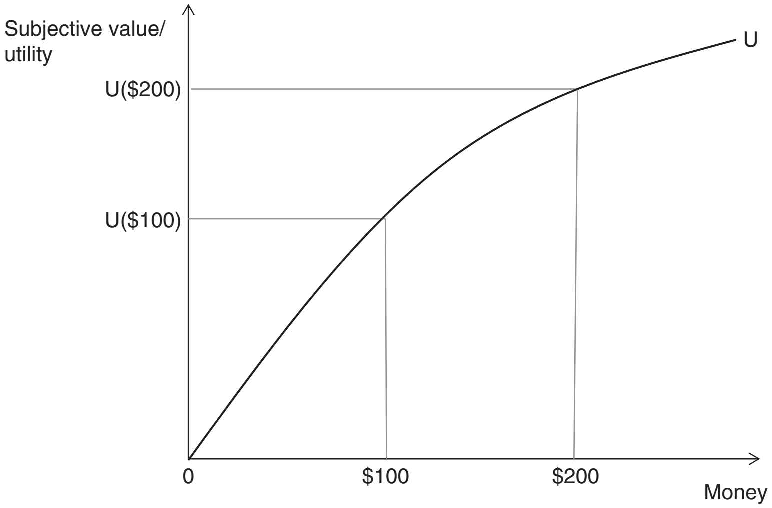

FIGURE 1.1 Declining marginal utility
Note: The declining marginal subjective value, or utility, curve is known as a concave utility function. It demonstrates that the utility enjoyed from a relatively large amount will be less than double the utility of an amount that is half as large. For example, the utility of $200 in Figure 1.1 is less than double the utility of $100. Similarly, the utility given by an additional, or marginal increase of, say, $5 on top of $200 will be less than the utility given by $5 on top of $100. In standard economic theory, the declining marginal utility curve is also assumed to apply to most goods in addition to money. A concave utility function implies that an individual dislikes taking risks, or is risk averse. That is, a person will sacrifice some of the expected value of a lottery in order to receive an amount of money for certain. For example, when faced with a fifty-fifty gamble of receiving $200 or $0, the individual would accept an amount less than $100, the expected value, in order to avoid facing the gamble.

accommodate the St Petersburg paradox (Bernoulli, 1738). He argued that the subjective value of money increases at a decreasing rate and that lotteries, rather than being evaluated in terms of their expected value, are evaluated in terms of their expected subjective value. This relationship between subjective value – or what is commonly referred to in the economic literature as utility – and money is illustrated in Figure 1.1, and has been a key assumption in the development of economic theory over the past two centuries.

The implication of Daniel Bernoulli's assertion was that a rational individual would seek to maximise expected subjective value rather than expected value, and by conceiving the concept of what is known as Bernoullian expected utility, he had laid the first cornerstone for later developments in rational choice theory.

A QUIET INTERLUDE

Following Bernoullian theory, the subjective value that people place on a good is meant to reflect their strength of preference for that good. For example, alluding to Figure 1.1, perhaps some people feel that $200 is not twice as good as $100, but, say, 1.6 or 1.8 times better. Subjective values that reflect strength of preference are known as cardinal utilities. Between the latter part of the eighteenth century and the end of the nineteenth century, some prominent thinkers, such as Jeremy Bentham and Francis Ysidro Edgeworth, attempted to develop techniques to measure cardinal utilities that could be meaningfully compared across different people. Essentially, the utilities were intended to be indicators of happiness and the measurement device was called a hedonometer, but they did not get very far in these attempts. The declining marginal utility curve was a central feature of utilitarianism, of which Bentham was the founding father. Thus, the Benthamite postulate that one should attempt to secure the greatest happiness for the greatest number was given a moral dimension, because the utility curve suggested that this could best be achieved by focussing society's efforts upon the relatively poor. However, by the beginning of the twentieth century, most philosophers and economists had reached the conclusion that an accurate quantitative measurement of interpersonal cardinal utility was impossible. For example, how could one accurately compare and quantitatively measure the enjoyment that one person experiences from listening to a Beethoven symphony to the pleasure that another person gets from eating a hamburger and to the pain that yet another person suffers from stepping on a pin?

Any development towards numerical indicators of utility was perhaps hindered by a discursive style of political economy dominating the early economics discourse. Adam Smith, in the book that he is less well-known for among economists, had written at length on human psychology, and made statements pertaining to, for example, present bias, loss aversion and reciprocity (Smith, 1759). These behavioural phenomena will be discussed later in this book, but are at odds with several of the assumptions underlying twentieth-century rational choice theory and have caused some to label Smith a behavioural economist (Ashraf *et al.*, 2005). In this regard, however, Smith's statements were intuitive rather than empirically quantified, and his successors over the following century tended to intuit also. For instance, Ricardo stated that the value of a good is proportional to the cost of the labour taken to produce it, Malthus believed that population growth would inevitably lead to famine and John Stuart Mill and Karl Marx warned that wages would never rise much above subsistence levels, but empirical analyses to test competing claims were lacking. Consequently, the perceived usefulness of political economy for informing and assessing policy was called into question.

Alfred Marshall, one of the founders of neoclassical economics, set out to make the postulates of the discipline more testable and tested, and in contrast to the forewarnings of Mill and Marx, demonstrated empirically that wages were increasing over time as a consequence of greater productivity necessitated by competition. Although Marshall introduced mathematical rigour, he felt it important that economic texts ought to be accessible to the layperson. Perhaps unfortunately, Marshall's views on this matter went unheeded, with many economists keen to transform their discipline into something akin to a natural science, with the quest increasingly focussed upon developing neat models of internal consistency that were divorced from the human experience. Following Marshall, economics was revolutionised, with often almost impenetrable mathematical and applied empirical economics superseding political economy as the dominant force within the mainstream economics community.

Since most economists had turned away from the attempt to derive cardinal measures of utility by the time Marshall was writing, welfare economics proceeded to be built instead upon Vilfredo Pareto's criterion, which did not require cardinality, a development that Bruni and Sugden (2007) have called the Paretian turn. Pareto maintained that an improvement in the economic organisation of society would require at least one person becoming better off, without necessitating anyone to become worse off (Pareto, 1971). Therefore, all that was required was an ordinal measure of value – in other words, an indicator of improvement or deterioration from the existing situation – not a cardinal measure that specified the strength of those movements. Utility theory was not immune to the steady introduction of mathematics into economics. Between the 1920s and 1950s, a number of notable mathematicians and mathematically minded economists, including Frank Ramsey, Leonard Savage, John von Neumann, Oskar Morgenstern, Jacob Marschak and Paul Samuelson, contributed towards developing systems of formal logic that prescribed how people ought to choose when they are faced with risk or uncertainty if they want to maximise expected utility. This movement culminated in the specification of expected utility theory, or neo-Bernoullian theory, and lent itself to the development of instruments with which measures of cardinal utility could, it was proposed, be elicited. The interest in measurable utility was reborn (Camerer, 1995).

THE NEO-BERNOULLIAN FORMULATION

If Daniel Bernoulli's assumption that people should aim to maximise their utility is correct, then people ought to obey the formal axioms – the logical assumptions – of what is known as expected utility theory. When we say that people ought to behave in a particular way we are making a normative statement, as opposed to a descriptive statement, which relates to how people actually do behave. Expected utility theory is still today the dominant normative theory of decision making when faced with conditions of risk and uncertainty.

At Princeton in the 1940s, the mathematician John von Neumann and the economist Oskar Morgenstern specified most of the axioms of expected utility theory as a small part of their seminal work on game theory (von Neumann and Morgenstern, 1944). The axioms were almost immediately recognised as important in the economics community, not least because it proved easier to judge the intuitive plausibility of specific axioms than the utility representation – i.e., expected utility maximisation – that they imply. Expected utility theory was refined in the years immediately following von Neumann and Morgenstern's initial exposition (Marschak, 1950; Samuelson, 1952), and its crucial axioms are now thought to be ordering, continuity and independence (Camerer, 1995). Although some contend that there is no broadly accepted definition of behavioural economics (Heukelom, 2012), the discipline is commonly thought to focus upon the set of observations that show that people often systematically, and therefore seemingly deliberately, violate the assumptions of rational choice theory and the broader assumptions of standard economic theory, which we will consider later. The challenges to the axioms, particularly the independence axiom, were the origins of empirical behavioural economics, and thus a good student of behavioural public policy ought to be familiar with them.

Before describing the axioms, it is important to note that although the terms risk and uncertainty are often used interchangeably in the popular discourse, in economics they have distinct meanings. Von Neumann and Morgenstern's axiomatic framework was developed for decision making under conditions of risk, where probabilities are objectively known. For example, there is a 50 per cent chance that a fair coin will land heads up. Uncertainty, or ambiguity, refers to an event where the occurrence of a particular outcome falls within a range of probabilities. For instance, there might be a 15–30 per cent chance of rain tomorrow. Leonard Savage provided an axiomatic framework for expected utility theory under conditions of uncertainty, and developed what is known as subjective expected utility theory (Savage, 1954). Strictly speaking, subjective expected utility theory can be applied more broadly than expected utility theory in real world settings, because specific probabilities of events are rarely objectively known. However, Savage argued that when faced with mutually exclusive possible states of the world, individuals will themselves attach a specific probability to the occurrence of each state. Expected utility and subjective expected utility theory thus share the same crucial axioms.

To return to the axioms, ordering imposes two requirements on people; namely, that their preferences should be complete and transitive. Completeness is simply the requirement that people are able to express a preference between two or more goods. For example, if we select motor cars as the relevant goods, an individual should be able to state that he prefers a Mercedes over a BMW, or vice versa, or that these two types of car are equally preferable to him; that is to say, he is indifferent to the choice of car. Transitivity implies that if an individual prefers a Mercedes over a BMW, but prefers a BMW over a Jaguar, then he ought also to prefer the Mercedes over the Jaguar. A violation of transitivity is known as an intransitive cycle, which can have serious negative economic consequences for the perpetrator. For example, assume that an individual prefers a particular Mercedes over his own BMW, prefers the BMW over a Jaguar that he has noticed, but also prefers the Jaguar over the Mercedes. He would therefore be willing to swap his BMW plus pay a premium, say, \$x, for the Mercedes. He now owns the Mercedes, but he would be willing to swap this car plus pay a premium, say, \$y, for the Jaguar. He now owns the Jaguar, but would be willing to swap that car, plus pay a premium, say, \$z, for his original BMW. Therefore, in terms of car ownership he is back where he started, with the BMW, but has paid out \$x + \$y + \$z in the process. If he were to repeat this cycle, he may soon find himself without enough money to buy petrol. In the economic literature, this is known as a money pump, an economically irrational cycle that can lead to bankruptcy.

Continuity requires that if an individual is faced with three goods, there will be a specific, unique probability such that he will be indifferent between a gamble that offers a chance of the most and least preferable goods, and the intermediate good for certain. For example, if money is the good being offered, and, as is likely, an individual prefers \$20 over \$10 and \$10 over \$5, there will be a unique probability, say, p, where the individual is indifferent between receiving \$10 for certain, and a gamble offering p chance of \$20 and (1–p) chance of \$5. As we will see later, continuity is central to the utility elicitation instruments, but has not been subjected to much attention in the behavioural economics literature.

Independence, sometimes call separability (Broome, 1991) or the sure thing principle (Savage, 1954), is the most controversial axiom of expected utility theory, and implies that the intrinsic value that an individual places on any particular outcome will not be influenced by varying other possible outcomes on offer, or by varying the size of the probability of the outcome occurring. The implication of the independence axiom is that if an individual is asked to choose between two or more lotteries, then a common outcome that has the same chance of occurring across the lotteries will be deemed irrelevant to the individual when making his choice. The chance of the common outcome occurring is a sure thing irrespective of what is chosen, and the individual ought only to base his choice on the consequences that distinguish the options he faces.

Some of the main challenges to the independence axiom will be detailed in Chapter 2, but an indication that the value that people attach to different outcomes or goods is often dependent on the other possibilities that are available in the choice that they are presented with, even when there is no risk or uncertainty, is given by a phenomenon known as asymmetric dominance (Huber *et al.*, 1982). For example, a person is more likely to choose to buy a 28-inch television set for \$600 instead of a 24-inch television set for \$500 when a 26-inch set for \$650 is also included in the choice set. The 26-inch set serves as a decoy, affecting the value that people attach to the 28-inch set by making it appear better value for money than would otherwise be the case.

Lotteries with up to three outcomes can be plotted in a Marschak-Machina triangle (1989). The axioms of expected utility theory have implications for the shape of the indifference loci within the triangle, with the indifference loci depicting the preference