CHAPTER 13

The Discount Rate

1 Introduction

The most important price in any CBA in the public domain is invariably the discount rate that is used. For in virtually any CBA in the public sphere, the costs and benefits involved do not usually accrue instantaneously at the point in time at which a particular project in question is carried out. They are usually spread over many years. For example, both the costs and the benefits that are expected from some new transport projects are expected to accrue over several decades. In other projects, major costs are incurred at the end of the project's operation. For example, a major cost in nuclear energy will be the costs of decommissioning and of disposing of the accumulated radioactive waste when the facility in question has reached the end of its useful life after two or three decades.

But the discount rate that society ought to use is an exceptionally value-loaded price. For how far society ought to count a pound's worth of costs or benefits accruing in some distant period of time as having the same utility value to society as a pound's worth of costs or benefits accruing today is largely an ethical question. And it is a crucial one.

For example, with a discount rate of only 3%, a £100 benefit or cost accruing in 20 years time will be worth only about £55, and if the discount rate used is 7% it will be worth less than £26. Thus the choice of the discount rate makes an enormous difference to a CBA of almost every public project. Yet it might appear that this is a purely technical matter, and that the discount rate is just another price and is fixed on the market

W. Beckerman, *Economics as Applied Ethics*,
DOI 10.1007/978-3-319-50319-6_13

like any other price. But, in fact, this is not the case. True, there is a discount rate in the market – or, rather, there are hundreds of them in the market. But the trouble is that these are unlikely to correspond to the rate that the public authorities ought to use for public projects if the objective is to maximise society's welfare. And it is not fixed in the same way as are ordinary prices. Why is this the case? The next section discusses this, and the subsequent sections consider what can be done about it.

2 The Market Rate of Discount

In the absence of personal considerations – such as 'sympathy' or friendship or family ties – a utility-maximising individual who had absolutely no preference at all between this year's and next year's consumption would not lend £100 this year for a certain repayment of only £100 next year. For if, say, 5% was the rate of interest he could get in the market for riskless loans, he could invest the £100 this year at 5%, and get £105 back next year. A promise of £100 next year would only be worth £100/1.05 this year, that is, £95.23. In other words, £95.23 is the *present value* to him of a certain receipt of £100 next year. It is all he would need to invest in order to receive £100 next year at the market rate of interest open to him. Five per cent is thus the rate at which he discounts the future.

What determines the market rates of interest open to him – that is, the rate at which he will discount the future? The theory of the determination of market interest rates is part of macro-economic theory and lies outside the scope of this book. But a brief simplified version of it could run something like this. At any point in time the general level of interest rates is determined largely by macro-economic considerations, such as those familiar to any student of economics and that are represented in what is known as the IS/LM framework. This framework encapsulates several variables, including the relationship between national income and the total amount of savings, people's preference for holding money rather than bonds, monetary policy, expected inflation rates, international capital flows and fluctuating and uncertain expectations concerning the future profitability of investment, and so on. As Keynes emphasised 80 years ago the rate of interest is virtually a current phenomenon and its value at any point in time is a poor guide to likely future rates of return on investment.

Nevertheless, the interest rate resulting from the interaction of these variables will still bear some relation to people's willingness to abstain from current consumption – that is, to save – and to the extent to which firms

can find it profitable to borrow funds in order to finance investment. As regards the former, although variations in the overall income level in an economy will tend to bring total actual savings into equality with total investment, the level of income at which this equality will be obtained will depend on how willing people are to save at different levels of income and interest rates. If, for example, at any given level of interest rates and income, people were relatively very reluctant to save, the level of income necessary to bring forth a level of aggregate savings that equalled total investment would have to be much higher than if people were very willing to save a lot at any given level of interest rates and income. It might even correspond to a level of income that exceeds the productive capacity of the economy, with a resulting tendency towards inflation or a suboptimal balance of payments.

3 Private Versus Society's Rate of Time Preference

Given the macro-economic reasons why the rate of interest is not a price like any other price, and because it is unlikely at any point of time to reflect accurately the future long-term rate of return on investment, it would be unwise to rely on it as a guide to the rate of discount that society should use in allocating resources between consumption and investment or between different investment opportunities.

In addition, there is reason to believe that, even in long-term 'equilibrium' (insofar as it makes sense to talk about 'equilibrium' in this context), market rates of interest will be higher than the rates that would be appropriate for society to use in its public investment decisions. It has been argued earlier that people's choices did not always promote their welfare, so that it might be reasonable to disregard their choices in some cases. But as far as the discount rate is concerned there may be good reason to disregard the market rate of interest (and hence of discount), *not* because people's choices will not correspond to *their own welfares*, but because *their* welfares may not correspond to *society's* welfare. There are various reasons why this may be the case.

First, the risk that savers take into account in deciding how much to save and invest is greater than the risk that society as a whole is facing. For example, one might invest one's savings in a company that fails on account of some superior product being developed by a competitor. In that case the individual who happens to have invested in it may lose all his money. The company that has been put out of business may disappear together

with the saver's savings. But society has not necessarily lost much. For example, the company's physical capital and the labour that had been employed in it will not have been wiped off the map in the same way that the saver's capital may have been. They may still be available for other productive uses. In this case, the risk of loss faced by the saver is greater than the risk of loss faced by society. Similarly, suppose that the company is nationalised without adequate compensation. The owners will lose out, but the assets still remain. Yet another ingredient of private risk that the ordinary saver might take into account is that he may not live to reap the rewards of his sacrifice of current consumption. It would be a pity to abstain from some current consumption and then get knocked over and killed by a bus the next day.

Second, in addition to these reasons for why, on the whole, private risk will exceed social risk, the amounts people will be willing to save or invest will be influenced by the existence of taxes. Individuals or companies who expect to pay taxes on the income they will derive from their savings and investments will need a higher rate of return on the investment than they would have done in the absence of the taxes. But from the point of view of society, these taxes are not a cost; they are purely transfer payments, even though they will usually involve some transactions costs.

Third, the market in financial assets is far from being a competitive market. In most countries a few very large financial institutions dominate the market and competition between them rarely takes the form of competitive rates of interest to borrowers. As a result, the market rate of interest is generally likely to be higher than it would be in a fully competitive market.

Fourth, there is probably a discrepancy between people's 'individual preferences' and their 'social preferences'. That is to say, people may prefer that society as a whole should discount the future less than they would do themselves in choices that affect only their personal allocation of resources. Indeed, there will be many policy options where an individual would be willing to contribute to some collective asset if he believed that others would do likewise – what Amartya Sen calls the 'isolation paradox'.[1]

Saving for the future is likely to be one such example of this. Individuals might not think it worthwhile making a sacrifice of their own consumption in order to save for future generations, but if they could assume that others were doing likewise they might be much more willing to do so. Hence, 'It is, for example, perfectly possible that in a society where no one saves anything, everyone might nevertheless be ready to vote for a political

proposal requiring each member of the society to save, say, 20 percent of his income for the sake of future generations'.[2]

Fifth, related to this it is likely that, at best, markets only reflect individual preferences over relatively short periods of time. They provide little information about people's preferences over generations, which is important in some areas of public policy, notably climate change.

Sixth, markets cannot properly reflect the interests and preferences of future generations, and some people would argue that society ought to do so, rather than just reflect the preferences of the current generation. As Sen says, 'If democracy means that all the people that are affected by a decision must themselves make the decision (directly or through a representative), then, clearly, there can be no democratic solution to the problem of the "optimum" rate of saving'. But he does go on to add that 'From this it does not, of course, necessarily follow that decisions made by some authoritarian body will be more "fair" to future generations than those made by the present generation voting collectively'.[3] This is another aspect of the general problem of social choice and the more particular value judgement concerning which is the society whose welfare we are seeking to maximise.

Thus, on balance, even in some ideal equilibrium state, market interest rates would generally be higher than the discount rates appropriate to public projects. Hence, there are two possible approaches that can be adopted in order to arrive at an appropriate social discount rate. One is to start with some assumed long-term equilibrium market rate and then to try to adjust it for the excess of private risk over social risk or for the influence of tax rates on the equilibrium market rate. But, of course, the vast amount of information that would be required in order to estimate the whole general equilibrium set of relevant equations is beyond our resources.

The alternative approach is to estimate what the appropriate rate of discount ought to be for society, given the sort of considerations that we think ought to be taken into account by society in fixing a discount rate. One of these will be the extent to which society ought to attach as much importance to the welfare of future generations as to the welfare of current generations. This is clearly an ethical judgement. Consequently, as the *Stern Review* recognises, it is not possible to avoid ethical issues when selecting a discount rate in this context.[4] This second approach has taken the form mainly of a focus on the two parameters in what is known as the 'Ramsey equation', in honour of the genius Frank Ramsey, who published the equation in a rightly famous article in 1928.

4 The Social Rate of Discount

The Ramsey equation is generally written as

$$r = \delta + \eta g$$

There are differences in the terminology used to define the variables in this equation. We shall define them as follows:

r = the social rate of discount (i.e. the rate that society *ought* to adopt for public projects or policies)[5];
δ (delta) = the 'utility', or 'welfare', discount rate – that is, *'pure time preference'*;
η (eta) = the elasticity of utility with respect to consumption;
g = the expected future growth rate of consumption.

Thus *delta* reflects the extent to which a unit of utility accruing in the future is valued less today than an *equal* unit of utility enjoyed today. And eta reflects the assumption that, on account of the assumed diminishing marginal utility of consumption, the higher is expected future consumption the less utility is derived from a given further increase in consumption. Thus the two parameters together will reflect the extent to which expected future consumption will have less value from today's viewpoint. It will produce less utility and a unit of utility will have less value today. In the above equation, the two parameters, *delta* and *eta,* are often regarded as 'ethical' parameters, although, as argued below, the second of the two is impossible to estimate with any confidence; this is on account of practical difficulties, not ethical considerations.

4.1 The Ethics of Delta

Delta has two components. The first is society's rate of 'pure time preference' and corresponds to what, at the level of the individual, has been described as 'impatience'. The second component allows for the possible extinction of the human race. (To an individual this would correspond to an allowance for the realisation that one is not immortal, and abstract from the satisfaction one might derive from leaving bequests.) Combining these two elements, *delta* is the trade-off that society ought to make between a unit of utility (or welfare) accruing in the future and a unit of utility

accruing today. For the sake of the argument let us assume that for practical purposes we can ignore the second component of *delta* so that *delta* is simply equal to the first one. In that case, if one believes that a unit of utility enjoyed by a future generation ought to be valued as much as an equal unit enjoyed by the present generation, *delta* ought to be put at zero. If, however, society ought to value a unit of utility accruing to a future generations below that of an equal unit of utility accruing to the present generation *delta* should be positive. This is clearly a value judgement. And the more society ought to value current utility over future utility, the greater is *delta* and hence the greater the discount rate.

Economists who are under the spell of impersonal Utilitarianism would tend to believe that society ought to adopt a totally *impersonal* approach to the weight that ought to be attached to the utilities accruing to different people. So they would attach equal value to an equal unit of utility (but not consumption) accruing to a member of any generation. Hence – leaving aside the extinction possibility – pure time preference should be zero. This means that the value of *delta* in the Ramsey equation ought to be zero. One authoritative expression of this view is in the *Stern Review*, which states, for example, that 'we take a simple approach in this Review: if a future generation will be present, we suppose that it has the same claim on our ethical attention as the current one'.[6] What sort of ethical principle can justify this equality of valuations of welfare among generations?

It would obviously be the *impersonal* principle that is characteristic of most versions of classical Utilitarianism (Chapter 9). In this approach the goodness of any outcome is measured by the total utility resulting from the actions in question irrespective of who gets the utility or when. The *Stern Review* acknowledges that 'it is, of course, possible that people actually do place less value on the welfare of future generations, simply on the grounds that they are more distant in time. But it is hard to see any ethical justification for this' (p. 31).

The *Stern Review* appeals to some very eminent economists – from Ramsey and Pigou down to contemporary Nobel Laureate Solow – in support of the view that 'the only sound ethical basis for placing less value on the utility (as opposed to consumption) of future generations was the uncertainty over whether or not the world will exist, or whether those generations will be present' (Stern, 2006, p. 45). The quotations from Ramsey, Pigou and Harrod that are routinely quoted in condemnation of pure time preference as representing impatience or 'defective telescopic faculty' are as follows:

Ramsey: '. . . we do not discount later enjoyments in comparison with earlier ones, a practice which is ethically indefensible and arises merely from the weakness of the imagination . . . '
Or *Pigou:* ' everybody prefers present pleasures or satisfactions of a given magnitude to future pleasures or satisfactions of equal magnitude, even when the latter are perfectly certain to occur . . . this preference for present pleasures . . . implies only that our telescopic faculty is defective'.
Or Harrod: 'Time preference in this sense is a human infirmity . . . a polite expression for rapacity and the conquest of reason by passion'.

These sentiments may be perfectly justified critiques of the pure time preference exhibited by most individuals *in making choices affecting their own lives.* But they do not seem to have much bearing on time preference *over generations.* As Schelling wrote many years ago, 'That is because the alleged inborn preference for earlier rather than later consumption is exclusively concerned with the consumer's impatience with respect *to his or her own consumption'* (our italics).[7] Schelling went on to say that it is absurd to believe that the virtually universal preference for consumption during one's lifetime by oneself or some member of one's family as compared with consumption by somebody else in 200 years' time can be described as 'impatience' or 'myopia'. And Harrod's frequently quoted apparent disdain for the 'conquest of reason by passion' is very odd since Harrod was certainly familiar with Hume's even more widely quoted assertion that 'reason is, and ought only to be the slave of passions, and can never pretend to any other office than to serve and obey them'.[8] (As pointed out earlier, Hume's concept of 'passion' was a general term for emotion, attitudes and desires, and was not the same as the current connotation of passion in terms of extreme obsessions such as those often associated with romantic sentiments.)

The *impersonal* Utilitarianism espoused by Stern and others in the context of the discount rate can be contrasted with some version of what is sometimes known as 'agent-relative' ethics. This has been discussed in Chapter 10 so it would be superfluous to repeat the main points here. Suffice it to say here that the discounting problem is an extreme example of the general problem of 'moral distance' to which reference has already been made – that is, how far ought we to value the welfare of everybody else, including future generations, on a par with our own welfare. As argued in Chapter 18, there is no clear theory of intergenerational justice to guide us. And Utilitarianism is the only one of the main theories that is able to provide guidance as to the precise value of *delta* in the Ramsey

equation. For it requires pure time preference to be valued at zero! It is for this reason, perhaps, that impersonal Utilitarianism may appeal to those economists who die with the word 'maximise' engraved on their hearts, since it provides them with a precise figure to be inserted into a social welfare function extending over time that has to be maximised. But 'agent-relative ethics' or 'special obligations' ought not to be dismissed as lightly as is often the case in the sphere of policies that have very long-term effects.

4.2 The Ethics of 'Eta'

In standard micro-economics most individuals are assumed to have utility functions relating utility to consumption. The functions are generally assumed to be concave, thereby expressing the assumption that the higher their consumption the lower the marginal utility they would derive from a further increment of consumption. The greater the concavity of the function the faster would marginal utility fall off as consumption increases further. And at the level of society as a whole, if one adopts a 'welfarist' utility function such as the social welfare function discussed in Chapter 7, society's utility function resembles that of an individual, so it would attach lower marginal utility to higher levels of consumption. At any point in time, this means that society attaches less utility to a unit of consumption accruing to a rich man than to a poor man. And, if the degree of concavity is represented in the utility function by the coefficient *eta,* the greater is *eta* the greater is the fall in marginal utility attached to higher levels of consumption.[9]

Up to this point in the story, time – and hence discounting – does not enter into it. But if time is brought into the picture, the story changes. For if the overall level of consumption in society is expected to rise over time less value will be attached to future increments of utility. And the greater is *eta*, the faster will marginal utility fall off as consumption increases. Other things being equal, therefore, the more society would discount future consumption, so that less would be saved.

But what about a possible fall in consumption? If we go back to the individual, the greater the curvature of his utility function – that is, the greater is eta – the more utility he would lose from a fall in his consumption level. Hence, the greater is *eta* the more risk-averse he would be. For he would expect to lose more utility from a given fall in his consumption level. People differ, of course, with respect to their aversion to risk. At

some level of income some people may be gamblers, so that, over the relevant range, they will have convex utility functions.

This means that, when time is brought into the picture, the implications for discounting of different values of *eta* on discounting are mixed, and tend to conflict with each other. On the one hand, if, for example, consumption is expected to rise over time (i.e. a positive value for the coefficient 'g' in the Ramsey equation) the greater is *eta* the more would society discount future consumption levels. For people would expect to derive less pay-off in terms of utility from future increments of consumption. At the policy level, therefore, a lower level of investment would be justified.

However, if it is believed that there is a risk of some catastrophe – for example, caused by environmental disaster – then it might be necessary to carry out a much higher level of investment in order to reduce the danger of a such a disaster. This appears to imply a lower discount rate – that is, a lower value of *eta*. In short, the *eta* coefficient is expected to do too much work.[10] On the one hand, the more marginal utility is expected to fall off with higher consumption – that is, the greater is *eta* – the more should society discount the future. But the greater is *eta* the greater is risk aversion. So if it is believed that society does face some risk to future consumption levels, the greater should be the level of preventive investment.

But there is nothing particularly 'ethical' about these two ingredients of *eta*. There is no *ethical* reason why the utility functions of individuals or society as a whole should have any particular degree of declining marginal utility of consumption or any particular degree of risk aversion. It is true that the more concave is one's utility function the more one would fear a fall in consumption and hence the more one would prefer a more equal distribution over time of one's own consumption level. But there are two reasons why this hardly translates into an *ethical* judgement about the desirable degree of intergenerational equality.

To begin with, as argued in Chapter 15, it is not obvious that concern with equality at the level of society as a whole is an e*thical* matter, rather than a purely *instrumental* concern with the possible adverse effects of inequality in society. For, if pushed to explain *why* one should value equality, most – and possibly all – people would maintain either that it reduces poverty or that it makes for more harmonious social relationships. For example, it leads to less conflict between people or classes, better health, better labour relationships, and so on. In other words, equality may be believed to be *instrumentally* valuable. But it is not easy to see why equality is also *intrinsically* valuable.[11] It may be possible to justify such a view in the context of some theory of

intergenerational justice, but the difficulties that such a view would encounter are discussed in more detail in Chapter 18.

Second, insofar as greater equality at any point in time has the instrumental advantages enumerated earlier, the latter do not seem to apply between non-overlapping generations. For example, better relations or less conflict between people who exist at the same time cannot be claimed to apply also between non-overlapping generations.

In any case, how many proponents of inter-generational equality would object to continued economic growth and deplore the growth that has been achieved over the past, on the grounds that it would lead to, or had led to, greater equality between generations?

So if the value of *eta* is not an ethical question, for which there can be no objective answer, what other basis is there for attaching a figure to it? The *Stern Review* settled on *eta* equal to unity, apparently on the grounds that it corresponds to empirical estimates. But the *Stern Review* also asserts at various points that zero pure time preference would be supported by most people. However the empirical evidence for this view is very shaky. One study restricted to residents in Washington D.C. and Maryland showed that the average respondent would trade off 45 lives in 100 years' time against one life today, and another study in Sweden implied a trade-off over a similar period of 243 lives against one life today.[12] And it seems highly likely that a survey of the whole world's population would give an even higher trade-off. And even if reliable and relevant estimates could be made of people's discount rates, it would be even more difficult to identify how much these rates reflect the 'pure time preference' component of people's discount rates as distinct from the 'eta' parameter in the Ramsey equation.

5 Which Country's Discount Rate?

This highlights one of the ethical limitations on welfare economics that we have mentioned at many points in this book, namely the definition of 'the society' whose welfare we want to maximise. Societies are not homogeneous entities. Whatever burdens 'we' ought to bear in order to eliminate the risk of climate-induced worldwide catastrophe depends on who we mean by 'we'. For example, the 'pure rate of time preference' is likely to be much higher in poor countries, where the average expectation of life may be even below 50 years, than in rich countries where it is approaching 80 years.

The same applies to the social opportunity cost of investing in policies that provide a yield over the very long term. For example, Arrow *et al.* suggest that, if account is taken of the higher rates of return on investments in developing countries and an assumed value of 2.0 for the *eta* coefficient, the appropriate discount rate could be between 10% and 16%, in contrast to rates of only about 2–3% in developed countries.[13] In the same report Arrow *et al.* found that 'a review of World Bank projects estimated a real rate of return of 16% at project completion; one study found returns of 26% for primary education in developing countries'.[14] Other examples of very high rates of return in developing countries include estimated rates of return of between 25% and 39%.[15]

Even in developed countries the social opportunity cost appears to be much higher than that to which appeal is often made in discussions of climate change policy. In the OECD countries, until relatively recently, the long-term average yield on equities was over 5% (after corporate and other taxes) for many decades, which is comparable to a pre-tax rate of at least 7%. Even in Britain it has been estimated that the return on 'productive capital' over the long run has been about 4–4.5%.[16] To the extent that public investment directly or indirectly crowds out private investment, the opportunity cost of capital needs to be incorporated into the CBA, so the market rate of interest is relevant here.[17]

Furthermore, the use of different discount rates for different projects is likely to involve a general misallocation of resources. Hence, if it is believed that the ethically correct discount rate is much lower than the prevailing market rate, it ought to be applied across the board – that is, to all investment projects – after due differentiation to take account of externalities and so on. The net result then could be that society would have to sacrifice too much consumption in order to devote an enormous proportion of output to investment. Social welfare is hardly likely to be maximised by such policies. Misallocation of resources could also mean that even future generations will be condemned to a lower consumption level than might otherwise have been available. It may also reduce the ability of future generations to overcome future problems.

6 Conclusions

Previous chapters have identified many ways in which value judgements enter into an appraisal of economic policies and projects. But there is one value judgement that enters into almost all of them. This is the discount

rate. For in the field of public policy, decisions have to be taken that will affect costs and benefits over a period of years stretching, in some cases, over generations. And the present value of costs and benefits that are expected to accrue in future years is highly sensitive to the discount rate that is used.

In normal everyday life, the rate at which an individual or a firm will discount the future will be closely related to market rates of interest. However, there are several reasons why this provides a poor guide to the rate at which society ought to discount the future. Consequently, attention has been concentrated on an ethical appraisal of how far society *ought* to value future costs and benefits relative to those that accrue today.

As discussed above, the usual framework for organising an analysis of the basic principles that should determine society's attitudes to future welfare by comparison with today's is based on the 'Ramsey equation'. This equation distinguishes between two crucial parameters. One is the extent to which a future unit of utility (welfare) is regarded as less valuable than an equal unit of utility accruing today. This is known as 'pure time preference'. Only if one is a Utilitarian can one put a figure on this parameter, namely zero. Otherwise one is in the no man's land of unquantifiable value judgements, where all that one can say is that if one cares very little about future welfare one would use what, in most practical applications, would be regarded as a rather high value for this parameter, and vice versa.

The other parameter is more complex and some attempts have even been made to justify a particular numerical value for it on the basis of surveys. But, even leaving aside the many limitations on the reliability of consumer surveys in this field, the problem is that this particular parameter reflects too many influences on human behaviour, some of which affect the parameter in opposite directions. For, on the one hand, it is related to the degree to which income levels in the future are likely to be higher than they are today and the extent to which this will lead to lower marginal utility of future increases in consumption. But this second parameter also reflects attitudes to risk – for example, the risk of some catastrophe in the future that could lead to a very big fall in consumption. Thus while this parameter is sometimes regarded as an 'ethical' parameter, the judgements that enter into it are as much judgements about what human preferences actually are as about what society believes they *ought* to be. It might have been expected that such judgements, like the selection of a

parameter to reflect 'pure time preference', could be illuminated by some theory of intergenerational justice. Unfortunately, as will be seen in Chapter 18, this does not provide a very useful solution.

However, as Chapter 18 also points out, justice does not exhaust the whole of morality. For example, the spontaneous manner in which most rich countries help poor countries whenever some natural disaster afflicts them owes little to any conscious theory of justice between nations. It seems that altruism or benevolence has become sufficiently deeply ingrained in the human psyche to become freed from the constraints of the smaller societies in which it was originally cultivated. This appears to apply also to concern with the welfare of future generations. So even though the vast majority of the world's population probably treat intergenerational distribution in an agent-relative manner (i.e. by giving priority to their families, friends and so on down the line), this does not seem to preclude a willingness to make some sacrifice for posterity, even when no advantage can conceivably be expected from it. One does not have to accept 100% *impersonal Utilitarianism* in order to make such a sacrifice.

The economic analysis of the discount rate has made a great contribution to bringing out into the open the ethical determinants of society's discount rate and hence to clarifying the ethical intuitions on which they depend. But, as emphasised at various points, ethics is about balancing conflicting intuitions and value judgements, and there is no way this can provide an objective scientific numerical estimate of where the balance lies.

Notes

1. Sen, 1984:123.
2. Sen, *ibid*:132.
3. Sen, *ibid:*121.
4. Stern, 2008
5. It is also often known as the 'social rate of time preference'.
6. Stern, 2006:31. The same sentiment is repeated on pages 45, 48 and 160 of this report. The view that one cannot logically attribute to unborn people any characteristic, whether it be brown hair or the possession of 'some claim' on anything or anybody (such as those embodied in 'rights') is discussed more fully in Chapter 18.
7. Schelling, T., 1995:396.
8. Hume, 1739: 2.3.3.5.

9. *Eta* is defined as the elasticity of utility with respect to consumption, For example, if the value of *eta* is taken as 2, a 10% increase in consumption would imply a 20% increase in utility.
10. See discussion of this in Beckerman and Hepburn, 2007.
11. See Raz, 1986, and Frankfurt, 2015.
12. Frederick, 2003.
13. Arrow *et al.* 1996:para. 4.2.1.
14. Arrow *et al.*, *ibid*:133.
15. Scott, *et al.* 1976:43–44.
16. Weale, 2009:4.
17. Arrow *et al.*, *ibid*:139.