5

The Ethics of Intertemporal Values and Valuations

We can see the scale and nature of the risks and the dangers of delay; we can see how to respond and what to do in the sense of how to set off in a good direction; and we have the economics to examine the problem, particularly if we think, with some imagination, of the dynamics, the learning, and the market failures. But how much should we do, and how fast should we do it? And who should do what and when? These are difficult questions, but they must be examined. To analyze this set of questions, because of the nature, scale, and longevity of the consequences of our actions, we must consider the underlying values and ethics both within and beyond the standard perspectives embraced in economics. This is the focus of the final two chapters of this part of the book.

Flowing directly from the previous two chapters on public policy and on economic modeling, the present chapter focuses on key ethical debates concerning climate change action that have played out within public economics: namely, over intertemporal values and valuations, often summarized under "discounting." My argument will highlight the deficiencies in—indeed, the misleading nature of—many of the conceptual approaches, assumptions, and perspectives commonly adopted in this important part of the economic analysis of climate change. Many discussions and analyses of intertemporal and intergenerational issues have been marred by serious analytical errors, particularly in applying standard approaches to discounting. The errors arise in part from paying insufficient attention to the magnitude of potential damages, in part from overlooking problems with market information, and unfortunately, in some cases, from inadequate understanding of the basic theories of discounting.

Accordingly, this chapter can be read as a normative accompaniment to the previous chapter's discussion which focused on the positive issues concerning the scientific and economic modeling of climate change. Together, the biases and confusions on both the positive and normative sides of the modeling and analysis have contributed to the stock of bad or misjudged arguments against strong and urgent action on climate change which this book seeks to challenge.

This chapter should also be considered in the context of chapter 6, which goes beyond the standard ethical approaches used in economics to consider a broader range of ethical foundations from moral philosophy that inform and enrich policy analysis of climate change.

The present chapter contains a substantial amount of technical detail and mathematics. This detail is necessary to identify precisely both the key issues at stake on discounting and the important biases or errors in some of the literature. I have tried to put most of the mathematics into the technical appendix that follows the chapter. Nonetheless, I hope the nontechnical reader, even though skimming or skipping some of the more formal parts of the discussions, will be able to understand the broad flow of the argument.

The chapter begins (section 5.1) with basic definitions of the concepts associated with discounting and shows what an understanding of these concepts implies for thinking about the major issues involved, including, centrally, how well or badly off future generations may be as a result of our decisions. In section 5.2, I discuss attempts to read ethical positions from market behavior and show how badly such attempts can go wrong. Section 5.3 examines some arguments around appropriate long-term discounting for public projects. Sections 5.4 and 5.5 are focused on ideas around "pure-time discounting," in other words the discounting of lives, or discrimination, purely on the basis of date of birth, quite separately from issues of how well or badly off people in the future may be.

5.1 Discounting and the big picture

Many of the decisions concerning climate change relate actions we may take now to possible consequences in the future. Some of these consequences will be far in the future, and many or most of them will be uncertain. How should we evaluate such consequences, many of which are difficult to understand and could occur in a diverse range of possible circumstances, relative to costs or benefits which take place now? This sounds like the question, familiar to economists from the economic appraisal of projects, of how to discount an incremental benefit of one unit of some good in the future relative to one unit of that good now.

If this were a set of issues involving fairly minor investments and consequences, and the assumption of one good were not misleading, then I would not raise queries about language. In this context, however, I think it is the phrase "intertemporal values and valuations" that we should use, rather than "discounting." To insist on the term "discounting" runs the risk of shoehorning this part of the economics and ethics of climate change into a very narrow form. Indeed some have seemed to imagine that policy analysis on climate change can be reduced to a discussion about a single number, the "discount rate," a concept still narrower than discounting.

Let us begin by defining the relevant concepts and then place them in the context of the scale of risk, the scale of necessary change, and the arguments for action. On discounting, the conceptual place to start is the discount factor, although many discussions jump straight to discount rates. The former concept is logically prior: the rate of the fall of the discount factor (for a given good, at a particular time) is the discount rate (for that good, at that time). The discount factor is the value of a unit of a given good at some date in the future, *t*, relative to the value of a unit of that good now. The discount factor, and thus the discount rate, will usually depend on both the good and the time being considered. The value of an extra unit of a good is generally assessed by looking at the impact on a social welfare function of the availability of an extra unit of that good. This is essentially the definition of a shadow price.[1] The impact on social welfare is the full effect (working through the economy) of an extra overall availability of the good.[2] It will depend on the functioning of the economy, including who gains and who loses from the extra availability, and the social weightings attached to those gains and losses. Shadow prices can be very different for different availabilities of goods, different times, different model structures, and different distributional weightings.

Accounting for net social costs and benefits from an action or project has to take place in terms of some unit of account. We call this unit of account the numeraire, and it is often a particular good (or sometimes a type of income). The discount factors and rates for the good at different points of time are those relevant for that good; they will vary with choice of good. The difference between discount rates for two different possible numeraire goods is equal to the rate of change of their relative shadow prices (see the technical appendix).

The concepts of discounting, discount factor, and discount rates have great usefulness; and we do not have to fall into the common errors described below. But in using them in the context of climate change, all too many economists and others fall into analytical traps they barely recognize. This is yet another example of the dangers in the economics and ethics of climate change of focusing on standard narrow formulations and losing sight of the big issues at stake. We have to look in a careful way, using all the economics we can muster, into intertemporal values in imperfect and uncertain economies where the range of possible circumstances is very large and future levels and distributions of well-being are strongly influenced by the decisions we take now. The Economics 101 treatment of intertemporal issues using models with perfect markets, small changes, and one good at each point of time does not come close to a reasonable representation of the issue at hand. It is profoundly misleading.

In my view, the key choices in policy for climate change are the strategic ones among radically different emissions paths. We should examine different emissions paths in terms of the consequences and risks they bring. When we do this we recognize that there is a very powerful case for strong and urgent reductions of emissions flows, with an aim of radically reducing the probabilities of the larger temperature increases and the great dangers they bring, which are likely under current trajectories. Could those strategic choices be reached or guided by describing the shadow prices, discount factors, and the like associated with current trajectories and gradually making investments guided by those prices? I would suggest that such a process, largely or solely guided by marginal signals, would be very unlikely to generate change at the pace necessary.

Formally speaking, in models of maximization, with "well-behaved" functional forms,[3] there is a duality between prices and quantities. At a maximum or optimum, the set of shadow prices associated with the optimum would imply that any perturbation from that optimum would make a loss if evaluated at those prices. In this sense the shadow prices guide us to an optimum, and each optimum has an associated set of shadow prices. If functional forms and model structures are sufficiently "well-behaved" at a point away from the optimum, if choices are guided by shadow prices associated with that point, then the incremental decisions will move us toward the optimum.

However, the context of climate change is much more difficult than that—and not only technically (probably there are many "badly behaved" relationships).[4] The scale of the risks involved, the uncertainties, the scale and nature of necessary change, and the dangers of delay make an approach based on investing guided by local shadow prices, in a process of many marginal steps, extremely difficult to carry through, and potentially very misleading. Issues of this scale and temporal sensitivity cannot be convincingly represented as the sum or integral of a sequence of marginal changes in a static model or in a framework where the clock is conceptually stopped while a process edging toward some optimum plays itself through. We have a complicated dynamic process, with major and disruptive change, with learning and discovery, and where risk is of the essence.

The sensible way forward, in my view, is a strategic analysis of possible actions and consequences along the lines just described and embodied in the approach of the preceding chapters. Implementation will indeed require a strong reliance on markets, entrepreneurship, and private investment, but the strategy has to be drawn on a bigger scale. This is a familiar mathematical and practical point. Mathematically, prices and marginal valuations will depend on the overall path. Practically, investors will find decisions very difficult if they cannot have some confidence in the overall direction of travel and the policies associated with it.

From this point of view an intense focus on discount factors and rates, particularly if made the central issue, can divert attention from the big strategic issues that are at the core of the decisions and challenges. The discount factors and rates will matter to investment programs, public and private. They will be important to implementation and to the calculation of important social costs, including that of carbon. But we must be careful to avoid missing the forest of the big decisions by looking only at the trees of valuing marginal increments. This is surely a case where we should start with an examination of the overall strategy and see marginal valuations, shadow prices, within that strategic framework. Both intertemporal and intragenerational values will influence that strategic framework, but they are also endogenous to it.[5]

Unfortunately, in discussions in economics concerning the ethics of climate change, the role of discounting is not only overdone but badly done. The misleading formulations and mistakes matter not only in terms of diversion from the big strategic picture but also in influencing the important role that discounting should play within the implementation of a strategy. The mistakes are closely linked to the attempts to shoehorn the very nonstandard set of issues in economics and ethics that are raised by climate change into the standard framework of marginal adjustments. The misleading formulations and mistakes matter more generally, beyond the subject of climate change—for example, in framing long-run decisions on infrastructure and on the environment. The rest of this chapter and the technical appendix are devoted to a more detailed examination of the key issues involved in intertemporal values and discounting in the context of climate change. Some of the issues are technical, but it is crucial to get the underlying concepts clear and the technical analysis right. Lack of clarity on basic concepts and techniques and misleading applications on flawed foundations, including on discounting, have caused great confusion in discussing climate change and have sometimes led the subject to deeply misleading conclusions.

5.2 Attempts to base discount rates on markets

In order to illustrate the limitations of standard approaches to discounting that try to base the specification of key ethical values on market information in the context of climate change, let me look at two propositions, each of which is false or deeply misleading in this context. First, “The question of discounting and intertemporal values for climate change is basically the same question as the choice of the discount rate, or hurdle rate of return, against which all investment projects should be assessed.” Second, “We can learn most or all of what we need to know for intertemporal valuations and the choice of discount rate in this context from capital markets.”

Regarding the first of these propositions, it should be clear from the definition of discount factors and rates (depending as they do on the good in question and on the time being considered) and from the magnitude of potential change that anyone who speaks of "the discount rate" in the context of climate change, as if there were one rate, constant over time, which should be applied to all investment projects, including investing in reducing emissions, should go to the bottom of the class. To formulate the question in this way is to make a collection of serious errors. The detail of an appropriate treatment of the issues and the nature and implications of the common errors are important. More technical aspects are contained in the technical appendix that follows; here the focus is on three particular issues that arise frequently in discussions and analyses of climate change: attempts to base discounting on market rates (this section); arguments around the use of the "opportunity cost of funds" (section 5.3); and arguments for pure-time discounting (section 5.4).

We have already gone much of the way to explaining why it makes little sense to try to derive implied social values for policy on climate change from market observations of rates of interest or return. Such attempts generally involve a whole series of mistakes. These are set out explicitly here because it is remarkable how tempting such a procedure has sometimes appeared to be in the context of climate change.[6] (As noted in chapter 4, however, some—notably Bill Nordhaus and Marty Weitzman—are changing their approach in the light of an enhanced focus on just how big the effects of climate change could be.)[7]

Interestingly, 6% has occasionally seemed a popular[8] choice for "discount rates." Supporting arguments—mistaken, as I shall argue—are sometimes based on medium- to long-term returns in rich countries (primarily the US) on risky financial assets such as shares. Assuming such a rate were to be applied over 50 and 100 years, it would mean a unit of benefit being valued 50 years from now 18 times lower than it is now, and 339 times lower than now after 100 years. To assume such a rate comes close to saying "forget about issues concerning 100 years or more from now"—an ethical conclusion which is so strong that its validity requires examination. It is important to show how mistaken this approach can be.

The assertion that discount rates appropriate for application to long-run climate issues can be derived directly from observed market rates of interest or return can include some or all of the following errors. Some derive from implicit underlying modeling assumptions concerning the nature of growth or decline across many dimensions of goods and people, and others from basic problems with the functioning and existence of markets.

Problems arising from underlying modeling assumptions

• Extrapolations from past rates of return involve the assumption that the past circumstances, including economic growth rates, for which those rates of return had applied will continue in a fairly similar manner into the medium- and long-term future. That assumption cannot be remotely plausible when the analysis must cover the real possibility that in some circumstances the effects will be so severe (as described in chapters 1 and 4) as to generate great damage to livelihoods, severe dislocation for many, major conflict, and substantial loss of life. In circumstances where future generations may be substantially worse off than ourselves, there would be a powerful case (and see below) for discount factors (for the associated consumption good) of less than one (extra units at some time in the future are seen as *more valuable* than now) and thus, over some periods, for negative discount rates.

• The possible devastation of the natural environment by climate change indicates the possibility that the discount rate with environmental services as the good under examination[9] could be negative. Suppose we have postulated a high discount rate, possibly market-derived, however mistaken that may be, with aggregate consumption as the numeraire; we must then allow the possibility of a rapidly increasing relative price of environmental goods, potentially at a rate much faster than the discounting. In practice this increasing relative price is likely to be forgotten.

• There is a potential for sharp deterioration in the distribution of income as a result of climate change—with the poorest being hit earliest and hardest. Thus taking a numeraire for discounting as aggregate consumption, while ignoring changes in income distribution, is likely to be misleading.

Problems arising from functioning of markets

• Capital markets, particularly for the long term, have a whole host of imperfections concerned with asymmetric information across various parties (including problems of moral hazard and adverse selection), market manipulation, limited ability to carry risk by different parties, and so on. In these circumstances, the standard arguments that they reveal relevant marginal rates of transformation or substitution over time are weak.

• There is no substantial financial or other market that applies to collective decision-making over a century or two. Markets deal mostly with individual decisions over relatively short-term scales; the high end of timescales is perhaps two or three decades for mortgages or four or five decades for pensions. Thus the markets that might give us guidance (even if we forgot about the many other serious problems with the argument) are not there.

5.3 Long-term rates in public decisions

Suppose we still tried to derive discount rates for social evaluation in this context, notwithstanding these serious errors or problems—what would we find if we tried to get as close as we could to a relevant discount rate? While I think it provides only limited guidance, for the reasons already described, the closest one we could find in the markets would probably be the "riskless" real rates on long-run government bonds. These are the longest-term among the options that might be available to individuals. Note that it is the riskless rate (nothing, of course, is completely riskless) that is relevant here, since in most of the specific formulations in terms of an overall societal objective, risk is usually handled separately by taking mathematical expectations of social welfare across the range of possible outcomes, taking account of their perceived probability.

Such rates for the UK and US have generally been around 1.5% or so over 50 years, hugely different from 6%.[10] Interestingly, as Weitzman points out,[11] the reason this may be so far below the long-run 6% on shares is the substantial riskiness in the shares, particularly in terms of "weight in tails"; this is an important argument, interesting in its own right, but which appears implicitly to accept a riskless rate around 1.5%.

Interestingly, a recent study by Giglio and colleagues has shown that long-run discount rates for individuals might also be around 1.5%.[12]

They examined the discounting implicit in comparisons between freehold prices (ownership in perpetuity) in the UK and Singapore and long-term tenancies for property ownership.

A common and longstanding line of argument on discounting in project appraisal is to use the "opportunity cost of funds." This could be based on long-term rates for government borrowing and could then lead to an appeal to the types of numbers, e.g. 1.5%, just discussed.

Another argument would be to base this opportunity cost on social rates of return to investment on the grounds that, through the production process, a unit of investment can be turned to $(1 + r)$ next period if r is the rate of return. In this argument the discount factor would be $1/(1 + r)$ and the discount rate r.

This is clearly to use "free investible funds" as the numeraire in defining the discount factor.[13] Investment is not an obviously useful numeraire for this context of possible long-term dislocation, and where the questions at issue refer to major damage to individuals. One way of expressing the approach implicit in adopting this numeraire and hurdle rate would be to say that we can make "standard investments" reap the returns that market rates indicate, and "buy down" or "repair" any climate damage resulting from climate inaction.

However, this line of argument suffers from many of the same problems described above. Long-term rates of return on investment might be negative in an environment where capital goods could be destroyed or where the investment process itself could have strong negative externalities in the damage it might cause. Shorter-term market rates might be poor guides to these longer-term (and possibly negative) rates. Further, the future prices of "buying down" environmental damage may be much higher than we anticipate now. In other words, as noted above, the future price or cost of environmental services relative to standard consumption goods may be very high in a world severely damaged by climate change.

A recent high-profile paper[14] recommends using declining discount rates for the US, a procedure adopted in practice by the UK and France, in the context of long-term projects, including climate change and environmental issues. It also contrasts the discount rates based on using consumption and investment as numeraires. It treats the problem in large measure as one of uncertainty over discount rates. The broad conclusion to use declining discount rates is a step in a sensible direction. In bringing in uncertainty over discount rates, the authors contrast a 1% rate, such as we might see in long-term bond markets, and 7% market rates of return on (risky) private investment. They argue that if the two rates are equally likely in terms of being plausible long-term discount rates for public decisions, then an equivalent path of year-to-year discount rates giving the same expected net present value would show declining discount rates.

While this may be a result of theoretical interest, it does not, in my view, put the role of uncertainty in the right place. The uncertainty in this context is over *outcomes*. Any uncertainty over discount rates should be derived from uncertainty over outcomes, and particularly, in this context, from the possibility that future generations could be worse off.

In addition, the approach seems fundamentally misguided in regarding the different rates (here 1% and 7%) as similarly plausible "candidates" for long-term discount rates. Given the definition of a discount rate here and its derivation, from a discount factor concerning the relative price of a good in different periods, and the deliberately "riskless" concept (risk is handled elsewhere in the modeling), a 7% return based on risky equities has little plausible status in the argument. This is a second and still more fundamental reason why the problem is misposed in this approach, popular though it is.

Further, this approach loses the key points in this context concerning multiple goods and changes in their shadow values. And it fails to identify properly the issues around choosing consumption or investment as numeraire, as described above. In this case I would suggest that choosing consumption as the numeraire is more helpful, as the issues concern potentially dramatic effects on consumption and livelihoods and the focus should be these. It is more natural to choose investment as the numeraire in contexts where decisions focus on the allocation of investment within a public budget. Overall, this approach, while of some theoretical interest, has little conceptual relevance to the problem at hand.

There is an interesting set of papers focusing on the implications of volatile period-by-period real rates of interest for long-run rates.[15] These show that such volatility would point to long-run rates substantially below the average of the short-run rates. Farmer and colleagues provide illustrations using historical rates for fairly stable economies of 2% or below.[16] Thus even if one were to go the route, dubious in my view, of relying on market information for ethical discount rates, then rates around 1.5% on current and historical data would look plausible. I note that, for different reasons, *The Stern Review* (see chapter 2 and the technical appendix) had discount rates implicit in the consumption path at around this level.

5.4 Pure-time discounting

An important issue in intertemporal values is pure-time discounting. This concerns the ethical status of different lives. How should we evaluate impacts of our actions for those whose lives begin later than ours? We must recognize, of course, that many of those who would be radically affected 50 to 100 years from now by our current decisions are not abstract, possible future lives: a lot of these people are already with us. Generations overlap. How can we discuss these profoundly important ethical issues? In particular, are there any ethical foundations in moral philosophy for treating two people, with exactly the same consumptions (where a person consumes many goods at each point of time) in each period of their lives, differently in our social valuations simply on the grounds that one life starts later than the other? To place a lower value on the life that starts later is "pure-time discounting."

A clear and understandable set of reasons for discounting the welfare of future generations might follow from there being some probability of exogenous annihilation of the world (or at least exogenous to decisions on climate change). If such survival were an issue, we might weight the contribution of the social utility at time t by the probability of the world not being annihilated by then. If annihilation (think of it as a meteor, say) arrives[17] in period t to $t + \Delta t$ (where $t + \Delta t$ is Δt later than t and Δt is very small) with probability $\delta\Delta t$, then the probability of survival to t is $e^{-\delta t}$. Then a discount factor of $e^{-\delta t}$ weights the utility at time t by the probability of being alive; the corresponding "pure-time discount rate" is δ. That is a clear and understandable reason for a formulation that looks like "pure-time discounting," which might command wide agreement.

But discounting the welfare of future generations *beyond* that reason looks like discrimination simply by date of birth. We would not be doing it because of doubts about the future generations' existence: that has already been covered in the arguments just made. And it is quite separate from marginal valuations which depend on the level of consumption or wealth—to isolate the "pure-time" question we are assuming the levels are equal. At this point in the argument, to make the issue clear, we examine two individuals identical in every way apart from date of birth. We are concerned here with the discounting of welfare itself, in other words, discounting lives.

In this context, there seem to be three types of argument that might be available, beyond that of the probability of existence, to justify pure-time discounting: (1) moral behavior should prioritize those closer to us; (2) people actually do not seem to care about future generations as much as their own, and that tells us what their moral position actually is; (3) technically we get into problems of incompleteness of orderings or non-convergent integrals in expressions of objectives in standard forms in economics if we do not allow pure-time discounting, or if it is very low. None of these arguments, in my view, holds water as a case for pure-time discounting.

The first argument has often been associated with David Hume of the Scottish Enlightenment in the eighteenth century. (Of course, just because a great philosopher has taken a position does not itself make that position compelling.) In developing this first argument one could suggest that much of moral behavior is, or should be, based in and defined by family life and those closest to us, and that any understanding of good behavior must start there. A functional or evolutionary underpinning might be involved in the sense that societies where people devote themselves first to family might function better or survive better than those where they do not. But arguments concerning better functioning of a society or higher probabilities of survival of a group seem to have minimal functional or evolutionary relevance to the question of how far to imperil the whole planet—that is more like a one-shot game.

As it happens, a more careful reading[18] of Hume indicates that he was well aware of the problems of individuals' impatience and that he saw "governors and rulers" as being needed to overcome them, as these excerpts from "Of the Origin of Government" illustrate.[19] In discussing one's resolution to do the right thing ("prefer the greater good") 12 months hence, he notes that as one approaches that time, A new inclination to the present good springs up, and makes it difficult for me to adhere inflexibly to my first purpose and resolution. This natural infirmity I may very much regret, and I may endeavour, by all possible means, to free myself from it. I may have recourse to study and reflection within myself; to the advice of friends; to frequent meditation, and repeated resolution: And having experienced how ineffectual all these are, I may embrace with pleasure any other expedient, by which I may impose a restraint upon myself, and guard against this weakness. … Here then is the origin of civil government and society. Men are not able radically to cure, either in themselves or others, that narrowness of soul, which makes them prefer the present to the remote. They cannot change their natures. All they can do is to change their situation, and render the observance of justice the immediate interest of some particular persons, and its violation their more remote. These persons, then, are not only induced to observe those rules in their own conduct, but also to constrain others to a like regularity, and enforce the dictates of equity thro' the whole society.

It seems that, far from asserting the moral significance and ethical attraction of pure-time discounting, Hume was arguing the opposite. And, like Rawls and Sen, he emphasized the importance of seeking a greater objectivity and morality through a more remote decision-making process that promotes the "greater good," that treats the future with less apparent impatience or disdain; he saw the possibility (perhaps optimistically) that "civil magistrates, Kings and their ministers, our governors and rulers" might perform that function. It is clear that in his view decision-making for the collective good is very different from narrow individual decision-making. In his use of language such as "infirmity," "narrowness of soul," and "weakness," Hume seems to be counseling strongly against the use of heavy discounting (based on short-term preferences) and in favor of embodying much longer-term perspectives, and much lower discounting, in values and processes, thereby encouraging "better" individual behavior.[20]

The second argument, i.e., "that's the way people are," is also deeply flawed in my view, in its attempt to make a case for pure-time discounting in the context of climate change (and, indeed more generally). In trying to understand ethical issues and identify ethical criteria and responses, we often examine with one another key questions and principles. We do this to try to inform a discussion and to help public reasoning (in Sen's language).[21] It would seem strange to say that group decisions must be taken, implicitly or explicitly, by some sort of vote or diktat which is uninformed by any attempt to reason together on the issues at hand. If such a discussion is indeed opened, as it should be in my view, then we should try to identify what principles should and can be of help and what arguments might be considered. John Stuart Mill in particular has reminded us of the importance of discussion in shaping our views: public reasoning does not simply concern facts and mechanisms but also helps in shaping our understanding of values. That is a tradition for which Sen has argued powerfully in a number of contexts, including in *The Idea of Justice*.

In such a public discussion, I doubt whether discrimination by date of birth would be regarded as acceptable in the case of voting or treatment by the courts. Why would it be acceptable to give a vote of a 55-year-old more weight than that of a 20-year-old? Yet with pure-time discounting at, say, 2% per annum, the value of a life which begins 35 years later is counted at only a half of that of one beginning 35 years earlier.

Perhaps even worse is an argument that says this generation has an unfettered right to impose its own views on future generations and damage their environment in any way it thinks appropriate, having taken account of how much it happens to value their well-being. That would be the consequence of saying that the right thing is whatever current voters decide. Would we think it right that we knowingly harm children now in pursuit of current pleasure? If not, then why would we think it right to knowingly harm their prospects?

5.5 Some formal economic modeling and issues around "oversaving"

A third argument for pure-time discounting has tempted some economists. It arises in one class of mathematical formulations and models and is generally associated with the possibility for such models to produce "oversaving" in the sense of very high savings rates or recommendations for (possibly indefinite) strong postponement of consumption. It is inevitably technical, and the nontechnical reader may wish to skim through this section.

Specifically we find that in some of these models an infinite integral, or sum over time of utility at each point in time, diverges over time unless we assume a substantial pure-time discount rate.

To understand the issues we require a little technical detail. Other technical issues are covered in the technical appendix. We consider a simple neoclassical growth model with growth rate n, of population N, and exogenous labor-augmenting technical progress at rate α, which will imply a long-run steady-state growth rate of consumption per head of α. Welfare is specified as the sum or integral over time of total social utility at each point in time: $Nu\left(\frac{C}{N}\right)e^{-\delta t}$, where C is total consumption and N is population. For simplicity we shall assume that $u(x)$ takes the special form $x^{1-\eta}/(1-\eta)$. We examine some ways of thinking about η in the technical appendix. It can be seen as an indicator of aversion to inequality in the distribution of consumption. In this context $e^{-\delta t}$ is a pure-time discount factor; utility at time t is discounted by this factor regardless of consumption levels; and, in this context, δ is the pure-time discount rate.

It is then easy to see that the convergence of the infinite integral requires[22]

$$\eta\alpha + \delta > \alpha + n, \tag{5.1}$$

which I call for convenience the "convergence inequality."

This same inequality guarantees that $\lambda_t K_t$ tends to zero, where λ_t is the discount factor at time t and K_t is the capital stock. This is the "transversality condition" for optimum growth, i.e., that the shadow value of the stock of capital should tend to zero.[23] The intuition is that if the shadow value of the capital stock gets bigger and bigger (technically diverges), we would effectively be accumulating capital inefficiently or postponing indefinitely: the social value of capital at time T is the future utility stream it can yield, and, in a "well-behaved" problem, i.e., one where integrals show convergence, that present social value should decrease. It is an analogous argument to that which applies in a finite-horizon model: we would try to use up everything by the end of the period.[24]

The Ramsey rule (see discussion of equations (5.7) and (5.8) in the technical appendix) for optimality tells us that the long-run marginal product of capital and the long-run social discount rate, $\eta\alpha + \delta$, tend to equality. Note also that the marginal product of capital which maximizes steady-state consumption per head is equal to the rate of growth,[25] ($\alpha + n$). If the condition of the convergence inequality were to fail, then the Ramsey rule would take us in the long run to marginal products below and capital stocks above the levels that maximize long-run consumption per head. In other words, we would be inefficiently building up excess capital, in the sense that a lower long-run capital would give higher consumption per head: thus the only paths satisfying the Ramsey rule, a necessary condition for optimality, would be inefficient; and they would fail to satisfy the transversality condition. Thus if a path satisfies the Ramsey necessary conditions for optimality, it cannot satisfy the transversality conditions for optimality. Therefore the failure of the convergence inequality implies that no optimum exists. If $\eta > 1$, a higher α makes existence and convergence more likely because the discounting effect on the left-hand side of the convergence inequality is then larger than the effect associated with the "consumption-per-head-maximizing" marginal product of capital on the right-hand side.

In the case of divergence of the integral, the intertemporal optimization, here embodied in the Ramsey rule, is constantly pointing to postponement. As parameter values take us closer to the divergence boundary, we will see higher savings rates emerging as "optimum." Some take this as an ethical argument for strong pure-time discounting. In other words they use an "inverse optimum" approach: we do not see, or would not want to see, such high savings rates, therefore people or societies must have high pure-time discount rates—they have "revealed their values." Or, some have argued, such very high saving would penalize current generations heavily and therefore there may be something ethically wrong with the assumption of low or zero pure-time discounting. Thus, "oversaving," "divergence of integrals," "incompleteness of ordering" (see below), and nonexistence of an optimum are very closely related.

But these problems might be telling us something about the perceived structure of the model—for example, that the perceived long-run rate of technical progress should be high; the term $\eta\alpha$ on the left-hand side of the convergence inequality is discounting that arises from the combination of growth, α, and the inequality aversion parameter, η. Or these problems may be pointing us toward selecting a higher η. Or they may be suggesting that the whole model structure, including the framework of an infinite integral, is misleading. While horizons may be long, are they really infinite in the formal sense we use in the mathematics? And the centrality of the convergence inequality, which compares long-run social discount rates with long-run growth rates, should remind us that this approach is unlikely to tell us much about pure-time discounting: in the context of climate change we are really very unsure about what long-run growth rates of population might be, and we have little idea of what long-run rates of exogenous technical progress might be when the physical environment may change radically. Both rates could be negative. Indeed, the whole idea of exogenous growth in this context is implausible.[26] Letting a crucial subject such as "pure-time discounting" be dominated by the technical aspects of assumptions about very long-run growth rates, when we can have very little confidence on what growth rates might be relevant, or indeed little confidence on what the model structure for the very long run should be, seems misguided.

There are genuine arguments concerning why we do not postpone consumption an indefinitely long time into the future, including the possibilities of extinction and of technical progress (α could be the result of endogenous technical progress). Possible divergence or otherwise of integrals does not illuminate the issues greatly.

We should note also, in the context of worries about "infinity" causing problems in models, that we do not show a willingness to pay an infinite sum to reduce the probability of a large number of deaths. Neither does such a willingness seem to be demanded by a consequentialist code of ethics. Indeed, it would make many decisions which involve risks of death very difficult to assess and manage. Generally, we know that there are many mathematical and other paradoxes associated with various formulations of the idea of infinity, and they do not necessarily throw great light on intertemporal values.[27]

A related argument is based on an axiomatic approach to intertemporal welfare evaluation.[28] Peter Diamond examines utility streams in an infinite horizon framework.[29] He shows that the assumption of *both* "equal treatment" (in the sense of a time reordering of a finite number of utilities being neutral) *and* "sensitivity" (higher utility at any time point increases welfare) is inconsistent with a preference ordering which is complete and continuous. He is considering an individual assessing different infinite streams of instantaneous utilities, but the result naturally carries over to the case of social welfare functions and the social evaluation of the utilities of a stream of future generations.[30] In the case of social evaluation, the assumption of equal treatment (also called "finite anonymity") rules out pure-time discounting, and sensitivity corresponds to the assumption that a Pareto improvement increases social welfare.[31]

The absence of completeness of an ordering arises from concerns over the ranking of the infinite tail of the utility stream, which behaves like a divergent integral in this model, when we have both "equal treatment" and "sensitivity." Essentially, trying to compare divergent integrals is trying to rank infinities of the same order, and that can lead to incompleteness (unless we are dealing with different orders of infinity such as aleph 0, 1, 2, etc., which is not the key point here).[32]

Some might be tempted to say that we have to abandon the assumption of equal treatment and go for pure-time discounting (discrimination by date of birth), thus getting convergent integrals, if the discounting is strong enough, and thereby a complete ranking. The first response to this argument is to note that, even if it were a strong one, it would not justify a large pure-time discount rate: a very small pure-time discount rate would suffice to give convergence of integrals in the models. But more importantly, the argument appears arbitrary. If a collection of assumptions is inconsistent, and you seek a consistent set, which one(s) do you drop? Amartya Sen has made this point over many years in his discussion of social choice theory,[33] and it is a clear theme in his recent book *The Idea of Justice*.[34] And as Diamond argues, "The goal is to answer a policy question. If a good criterion answers the question (compares the relevant alternatives), that is the end of the story. If the criterion does not answer the question, one needs further thought."[35] For example, incompleteness relative to some parts of the domain of choice does not invalidate a criterion.

Geir Asheim makes an interesting point in this context when he shows, with a "condition of immediate productivity" (positive returns on investment), that the (strong) Pareto and finite anonymity (pure-time-discounting) axioms imply that, generation by generation, utility should be nondecreasing, in other words should be sustainable in an often-used sense of that term (and see next chapter).[36] In the context of climate change and a potentially destructive environment, we cannot assume that in all circumstances all investment can be productive. Capital and other stocks can be destroyed, and thus returns on investment may be negative.

The line of argument based on investigating the consistency or otherwise of sets of axioms for social choice is interesting and valuable. But it does not tell us that we should discriminate by date of birth simply on the grounds that some sets of assumptions lead to divergent integrals and incomplete orderings.

Interestingly, many economists who have thought carefully about this issue have rejected pure-time discounting. Frank Ramsey (a philosopher and mathematician, as well as an economist) described pure-time discounting as "defective imagination." Roy Harrod and John Maynard Keynes agreed.[37] So too have some of the great economists of recent times such as Bob Solow and James Mirrlees, and the philosopher and economist Amartya Sen. They have seen no strong reason why we should discriminate across generations by date of birth and thus have seen no reason for pure-time discounting. Allowing for the probability of existence gives rise to a similar formality, but it is a different ethical issue. Of course, counting heads of proponents, even if very distinguished, is not necessarily a good way to evaluate an argument, but it does suggest that low or zero pure-time discounting is a considered position and not some capriciousness, regal edict or an assertion without explanation.

It is, however, the arguments that matter. I have tried to show that lines of argument commonly used in attempts to justify pure-time discounting as a moral position generally founder when scrutinized carefully. It is those who try to argue that high pure-time discounting is somehow "pragmatic" that do the "asserting," without sound or reasoned justification. Apart from invoking a weak position often wrongly attributed to Hume, they generally appeal to one of two arguments: (a) people just do it (i.e., exhibit impatience); and (b) some models "go wrong" in terms of strange conclusions if we make, together with other ethical assumptions and some structural assumptions in the model, the further assumption that pure-time discounting is zero or low. The former surely has weak ethical status, since it is rarely a considered response to an ethical question. The latter, as we have argued, should be seen as saying that there may be something wrong or missing from the models, or problems with other ethical assumptions, rather than necessarily implying strong statements about pure-time discounting; and we have gone further and shown that the assumptions concerning growth and decline and how they are endogenous to our choices may be critical assumptions for scrutiny. So too would be the heavy emphasis on taking literally the mathematical formulation of infinity. And if the "problem" that "goes wrong" is incompleteness, that may be something we can or should live with.

Low or zero pure-time discounting, on the other hand, is derived from consistent application of ethical principles that are transparent and widely used and explained in broader contexts than climate change. Equality of treatment and nondiscrimination, under the law and in connection with human rights, are basic to many constitutional or legal structures. John Harsanyi embodies the idea in his "impersonality" principle, as does Rawls in his "Veil of Ignorance."[38] Low or zero pure-time discounting simply applies the idea to date of birth.[39]

Putting the analysis of the one-good growth model together with the axiomatic treatment of Diamond, Basu and Mitra, Sen, and Asheim, we see that they are telling very similar stories. The convergence inequality and its derivation embody the essence of most of the relevant lessons from infinite-horizon growth analyses. They (the inequality and its derivation) (1) give us the intuition behind how nonconvergence of infinite integrals can arise; (2) warn us that convergence depends on long-run parameters such as exogenous growth rates of population and productivity, which in the case of climate change are unlikely to be exogenous and could be negative; (3) warn that if we look for optimum growth, these are the same parameter constellations that would prevent existence of an optimum; (4) explain how the nonconvergence or nonexistence of an optimum arises through consumption postponement; (5) explain that the case of zero α can oversimplify intuition in a misleading way, because with zero α, the convergence inequality becomes $\delta > n$ and the focus, particularly if n is assumed zero, comes to be on δ the pure-time discount rate; with nonzero α, both η and α are relevant. And remember that a long-run growth rate in consumption per head could arise from endogenous technical progress; it does not have to be exogenous here, nor does it have to be positive.

Thus the convergence inequality or condition and its derivation go to the heart of a number of issues and illustrate, among other things, the point that if a model produces difficult or inconsistent results we have to look at its whole structure and not just one parameter (here δ). And we should question (as we have emphasized in the context of climate change) the assumptions of a one-good growth model, exogenous technical progress, exogenous population growth, and the absence of uncertainty. We must ask, moreover, whether the particular expression of infinite horizons, in the mathematical sense often invoked, may be the tail wagging the dog. It is the practical policy challenge of what we should do over the next several decades that is at issue. The consequences for the next few centuries really do matter, but that does not mean that our modeling parameters and ethics should be driven by the peculiarities of the convergence of infinite-horizon integrals.

5.6 Conclusions

We have seen that the context of climate change, in particular the potential breadth of the effects across the whole of economic and social life, the potential scale, and the length of the time periods involved, require us to take care to define the intertemporal and intergenerational issues in corresponding breadth, scale, and time period. There is no serious alternative to framing the problem as the management of immense risks, to making our discussion of the ethics explicit, and to putting at the center of that discussion the key issues associated with the scale of risk, including potential large-scale environmental destruction for future generations, conflict, great poverty, major differential impacts across countries and groups at different points of time, and substantial loss of life.

Lest I be misunderstood, let me be clear that I think that narrow aggregative models, be they in growth theory or integrated assessment models (IAMs) of the standard kind often used in this area (and discussed in chapter 4), can have a contribution to make. But in this context of great risk and potential disruption of economic and other structures, the contribution is modest and only one part, and not necessarily the major part, of an assessment of policy on climate change. Further, as I have argued in the preceding chapter,[40] IAMs have generally grossly underestimated risks from climate change and thus should not be seen as a central case; most of them embody extreme cases or "what might happen if we are very lucky."

There have been some very interesting and important discussions of discounting in the economic literature which do not commit the mistakes described here and which explore important questions.[41] It is not the purpose of this chapter to review all such arguments. I have tried to focus attention on those intertemporal issues where difficult and important ethical perspectives should enter most strongly into economic debate. And there are, of course, valuable approaches to intergenerational issues which go outside the standard social welfare approach, as I shall discuss in the next chapter.

There is much economists can contribute to public discussion of intertemporal issues by applying their experience of the analyses of key intertemporal questions and the broad range of analytical tools at their disposal. But we must heed Sen's warning:

> If informed scrutiny by the public is central to any such social evaluation (as I believe is the case), the implicit values have to be made more explicit, rather than being shielded from scrutiny on the spurious ground that they are part of an "already available" metric that society can immediately use without further ado.[42]

In summary, the lessons we have learned from this chapter are:

- The concepts of discounting, discount factor, and discount rates have great usefulness in evaluating the future consequences of present decisions, but many economists have fallen into error in applying these concepts to the analysis of policy on climate change. One important example is to see discounting as, in large measure, exogenous (a search for "*the* discount rate"), when climate change can radically alter the circumstances of future generations and, in some possible scenarios, make them much poorer. That would surely radically alter the discount factor between those parts of the future and now.
- Other important errors relate to attempts to base discounting on market rates. Most markets and rates are not about the very long-term public decisions that are at the heart of much of the decision-making here. And most of the capital markets that might be thought to be relevant are very imperfect, among other things as a result of deep informational difficulties.
- In the appendix to this chapter, I argue that attempts to infer "governmental values" from public policy in this context do not take us far. A key problem is that such inferences require a specification of the model structures and constraints perceived by "the decision-maker" and results are very sensitive to such specifications. Thus, the argument is in places a little technical.
- Pure-time discounting involves valuing the welfare of people in the future lower than that of people today, simply on the grounds that they

are in the future. It is discrimination by date of birth and would get short shrift, and rightly so, if attempted in other parts of life.[43] We examined a number of possible arguments for such discrimination and found them all unconvincing. Given the weakness of such arguments, the adoption of strong pure-time discounting in economic models should be seen as the imposition or assertion of an arbitrary position. Equal treatment on the other hand has deep roots in moral philosophy.

- Where discount rates are used in modeling the economic benefits of climate policy, they should use consumption as the numeraire and adopt a pure-time discount rate close to zero (the small positive value reflecting the risk of planetary annihilation only).
- But we should not lose sight of the forest (the challenges of climate change) for the trees (questions around discounting): the key choices in policy for climate change are the *strategic* ones, among radically different emissions paths. As we saw when examining issues and strategies in earlier chapters, when we examine the consequences of different emissions paths there is a very powerful case for strong and urgent reductions of emissions flows.

Technical appendix

Appendix 5.1 Formal definitions of the discount factor and the discount rate

We start with the formal definitions of a discount factor and rate. If the discount factor for good i at time t is λ_{it}, then the discount rate, ρ_{it} at time t, is the rate of fall of λ_{it}, i.e., it measures how fast the discount factor is falling. In other words $\rho_{it}\Delta t$ measures how much less (proportionally) a unit of good i is valued at $t + \Delta t$ relative to its value at time t. In discrete time, if ρ_{it} were 0.1 per period, the unit of good i at time $t + 1$ would be 10% less valuable than at time t:

$$\rho_{it} \equiv -\dot{\lambda}_{it} / \lambda_{it}. \tag{5.2}$$

As we have emphasized, the "discount rate" depends on i and t. Equation (5.2) is sometimes said to define the "own" discount rate for good i.[44] When uncertainty is introduced, a further index for the "state of nature" will be necessary too; alternatively we can interpret goods in

This is a portion of the eBook at doi:10.7551/mitpress/10408.001.0001

different states of nature as different, so that the index *i* already carries that information (an umbrella when it is raining is deemed a different good from when it is dry).

How are own discount rates for goods *i* and *j* related? This is clear from the following, remembering that λ and ν depend on *t*:

$$\rho_{jt} - \rho_{it} = \dot{\nu}_{ij} / \nu_{ij} \text{ where } \nu_{ij} \equiv \lambda_{it} / \lambda_{jt}. \quad (5.3)$$

Thus the difference between the discount rates for goods *i* and *j* at time *t* is equal to the rate at which the relative valuation, or relative discount factors, or relative shadow prices[45] (we use the terms interchangeably for this discussion) are changing.

This type of argument has been familiar in economic theory for at least 60 years[46] and was generally recognized in the cost-benefit analysis literature of the 1960s, 1970s, and 1980s, in the phrase "the discount rate depends on the choice of numeraire."[47] In other words, if good *i* is chosen for numeraire, we will see a different array of discount rates from those which would arise if good *j* is chosen as numeraire. The choice of numeraire should not affect decision-making, but it will affect the expression of the social accounting.

We would expect λ_{it} to be influenced by the relative scarcity of *i* at that time. For example, if *i* were environmental services, it is possible that they would become more scarce, putting upward pressure on λ_{it} over time, and thus point to discount rates for this dimension being negative.[48] It would also be influenced by overall standards of living. If life was getting better, the overall λ levels might fall over time, but if conditions were getting worse they might rise.

Similar remarks apply when we recognize that incremental consumption may not have the same social value for different groups. Indeed, it is reasonable to argue that consumption increments for different groups will have different social marginal values, unless it is assumed directly (and that would seem arbitrary since the supporting ethical judgments are unclear) that they are all deemed to have the same value; or unless it can be assumed that an optimum set of lump-sum transfers is in place that sets the social marginal valuations to be the same, as a condition of optimality. The availability of such a set of lump-sum transfers is thoroughly implausible, for the usual information/incentive reasons that are standard in modern public economics.[49] Many would argue that a lower social marginal valuation should be applied to an increment to someone who is better off (e.g., under utilitarian or Bergson-Samuelson objectives).[50] The social value of good *i*, made available at time *t*, to household *h*, is $\lambda_{it}\mu^{ht}$, where μ^{ht} is the social marginal valuation of income to household *h* in time period *t*, or the welfare weight for *h* at time *t*. It would generally vary across households.

Narrow cost-benefit analysis, using standard marginal techniques, of small investment projects which cause minor deviations from a given growth path can be a very useful way of discriminating in a practical and reasoned way between different opportunities. In the context of climate change, however, we have seen that there are real possibilities that some (high-carbon) attempted investment or growth strategies could lead to immense change in the world's economic and social circumstances, including the possibility of rapid decline, wholesale destruction of the environment, radical change in income distribution, and the movement of people on a scale that could result in major, widespread, and extended conflict and loss of life.

In such a context, it could hardly be argued convincingly that discount rates could be treated as constant over time, for at least four important reasons. First, they would surely be sensitive to possible future income levels, particularly to possible decline, an outcome that could be associated with higher discount factors in the future and negative discount rates. Second, environmental services might collapse on key dimensions; thus it is difficult or impossible to make a case that the relative shadow price of such services and other consumption goods remains constant—discount rates will vary across goods (see equation (5.3)). Third, if world income distribution is radically affected, e.g., if the poorest are hit earliest and hardest by climate change, it cannot be argued that benefits or costs are likely to be spread in a fairly broad and stable way, and one which is similar for benefits, for costs, and for the raising of resources: these are the kind of arguments necessary if income distribution issues are put to one side. Fourth, widespread loss of life would require going outside the standard formulations of a fixed number of individuals and the evaluation of consequences only in terms of changes or goods accruing to those individuals.

Thus many attempts at presenting action on climate change as just another project to be compared with other possible projects, such as building water infrastructure in poor countries, fighting malaria, and so on,[51] are hopelessly flawed because they take no serious account in their method of cost-benefit analysis of the potential radical global changes involved in unmanaged climate change and their implications for intertemporal valuation.

Such simpleminded cost-benefit comparisons also usually treat all such projects as separate stand-alone projects with no interactions. Water issues, diseases and health, and climate change are intimately interlinked, and it is a basic mistake in cost-benefit analysis to treat interlinked projects as separate when one is attempting to evaluate them. Such attempts to force a profoundly nonmarginal and systemwide set of issues into a narrow marginal framework, and in the process overlook the interlinkage of projects, appear to involve an ignorance of the basics of the theory and practice of cost-benefit analysis. Careful analyses of consequences should indeed be at the heart of strategic decisions, but the narrow forms of cost-benefit analysis sometimes used or proposed ride roughshod over the science concerning scale and interlinkage, the economics in terms of the basic theory of evaluation, and the philosophy relating to the underlying ethical issues.

It is interesting to note that toward the end of one of his papers on discounting, largely based on approaches involving simple social welfare functions and marginal effects, Dasgupta, who has written interestingly and wisely on the principles of project evaluation, concludes that the possible vast scale of losses makes such narrow cost-benefit analyses of very limited value.[52] I agree; and this suggests that focusing heavily on such narrow approaches risks diverting us from the analysis of intertemporal values and valuations in this context, where immense distributional issues and inequitable consequences can follow from our decisions.

Appendix 5.2 The Ramsey formulation

Let me now move from general formulations of underlying principles to a very particular structure which focuses on a one-good growth model. It involves a narrow approach to ethics and uncertainty and a highly aggregated approach both to consumers and goods. I do not wish to argue that it should occupy center stage in a discussion of discounting for climate change; but because it is so prominent in discussion of intertemporal ethics, it is important to set it out explicitly so that the problems in applying it in this context can themselves be made explicit. We shall see that notwithstanding its narrowness, the approach will also have some usefulness in analyzing parametric formulations of approaches to intertemporal values, and that it can illuminate some of the issues around pure-time discounting in infinite-horizon models. A core part of this story will be the Ramsey analysis or rule for optimum allocation between consumption and investment.

Many formulations of overall objectives in the modeling of intertemporal choices at the national level, in standard macroeconomic, general-equilibrium, or growth theory, use a simplified objective expressed as the maximization of the mathematical expectation of the integral of some function of aggregate consumption; see for example the following:

$$\mathrm{E}\int\sum_{h} u(C_{ht}, N_{ht}, t)\,dt. \tag{5.4}$$

The expectation operator E ranges over the space of possible outcomes, the integral is over time, and the summation is over individuals (or in many examples in the literature fairly aggregated subgroups in the population); u is a social utility function, C_{ht} is the total consumption (usually one-good but it could be a vector) of group h at a time t, and N_{ht} is the population of group h at time t. It is often assumed that C_{ht} is equally distributed among group h and consumption per head in the group is denoted c_{ht} (although other within-group distribution rules are possible).

While this formulation is fairly flexible in terms of the issues it covers, it does embody two important and narrow assumptions: it represents attitudes to risk by working with the expectation of social utility $u(\ )$, and it treats consumption at different points in time as separable.[53] The vector c_{ht} could be interpreted as allowing for the role of environmental services but all too often is in terms of one dimension only, aggregate consumption. While not embodied in (5.4) itself, applications often treat the probability distribution of outcomes at different times as independent—which is implausible because if climate turns out, for example, worse than expected,[54] then the same may be more likely to be true of later periods; further, some damages may be irreversible or long-lasting. See Stern (2013) and chapter 4 above for an elaboration of this idea.

In order to express in a simple way some of the common discussions of discounting based on this approach, let us further simplify the formulation of the social objective (5.4) as follows:

$$\mathrm{E}\int N_t u(c_t) e^{-\delta t} dt. \tag{5.5}$$

Now we have just one good, there are N_t households at a time t, with equal consumption[55] c_t, and "pure-time discounting" at a constant rate δ. Pure-time discounting here involves a discount factor, $e^{-\delta t}$, on future utility and attaches a lower weight to future generations or future utility of consumption entirely on the grounds that they are in the future. As discussed in the body of the chapter, it can be understood as discounting the lives, and thus utilities, of those born later simply on the grounds of date of birth, irrespective of what their consumption might be; it is quite explicitly discrimination by date of birth.

In this framework the discount factor at time t is

$$\lambda_t = u'(c_t)e^{-\delta t}. \tag{5.6}$$

Then λ_t is the marginal valuation of an overall extra unit of total consumption (distributed equally so that everyone at t gets an extra $1/N_t$); it is the partial derivative of (5.5) with respect to C_t. We take λ_0 as equal to one. Then λ_t can be seen as the shadow price of a unit of consumption at time t. In a simple optimum growth model, the optimality condition for the allocation between consumption and investment is that λ_t should be equal to the value of an extra unit of investment or capital.

The social discount rate, ρ_t, at time t is then the rate of fall of λ_t; that is,

$$\rho_t = -\frac{1}{u'}\frac{d}{dt}u' + \delta. \tag{5.7}$$

Just as the discount rate is the rate of fall of the discount factor, the social discount rate is the rate at which the shadow value of a unit of consumption falls. Optimality, or the Ramsey rule, would require that this should be equal to the social rate of return on investment or the rate at which the shadow value of an extra unit of investment falls; in simple cases this will be equal to the social marginal product of capital. The intuition is that on the margin, the return to allocating a unit of output to consumption or to investment should be the same.[56]

If $u(c)$ takes the special isoelastic form, with $u(c) = c^{1-\eta}/(1 - \eta)$ and $u' = c^{-\eta}$, equation (5.7) becomes:

$$\rho_t = \eta g_t + \delta, \text{ where } g_t \equiv \dot{c}/c \text{ at time } t. \tag{5.8}$$

Thus the social discount rate is equal to the elasticity of the marginal utility of income (η) times the growth rate plus the pure-time discount rate.

The Ramsey formulation has been widely used, although the narrowness of the assumptions used to derive it often get insufficient attention. We should note immediately in this context that we have to recognize that there are circumstances, with substantial probabilities, where g_t may be negative. Thus, as we have observed already, even in this narrow and often misleading framework we are quite likely to find negative discount rates.

The Ramsey formulation was discussed in appendix 2A of *The Stern Review*. That appendix also contained warnings about the unreliability of many of the simplifying assumptions: including that of using one good in relation to the important environmental services, uncertainty about future consumption, the possibility of decline for some individuals, endogenous population, and so on. These warnings appear to have been overlooked by many.[57] On the other hand, it has some usefulness in crystallizing some discussions of distributional judgments and some issues around infinite horizons, as we shall see below.

Appendix 5.3 Attempting to specify or infer distributional values

Let us return to the narrow framework of the one-good model and the expression for the social discount rate in equations (5.7) and (5.8): in particular we discuss possible ethical approaches to the specification of the "parameters" η and δ in equation (5.8). The formulation captures some issues of distributional judgments and pure-time preference in a clear and stark way: it focuses on relative consumption levels of future generations, distributional judgments (as reflected in (η)), and attitudes to discrimination between lives by date of birth, all of which are general issues beyond the narrow model.

The formulation is widely adopted, and thus it is important to see how discussion over the choice of η and δ might be articulated and which arguments might be relevant or robust. Choice of δ, pure-time discounting, has been examined carefully in the main body of the chapter. Here let us examine the specification of choice of η. If I were to take η as 2, then I would value an extra unit to person A, who has consumption one-fifth that of B, 25 times as much as an extra unit to B; if η were 1, then 5 times as much (see the specification of the isoelastic form of utility following equation (5.7)).[58] Such thought experiments are sometimes expressed in terms of a "leaky bucket," and this is a metaphor widely used in both economics and philosophy.[59] In the example of B with consumption 5 times that of A, we could ask whether we would make a marginal transfer from B to A even if four-fifths of it were lost from the bucket on the way. If the answer is "yes," then η is larger than 1. Such a discussion can directly inform a choice of η.

There have been various attempts to assemble evidence from some public or collective decisions involving distributional judgments that might throw light on what value of η might summarize the values behind such decisions. To infer values from decisions in a formal way, we would generally have to consider an "inverse optimum" problem: for an observed decision, we ask what values (here η) in the objective function would be consistent with that decision. A great difficulty with this approach is that, for it to be usefully informative on the implied ethical position, we have to have a plausible description of what is in the mind, or collective minds, of the decision-maker(s). And in the plural case we have to argue that the decision of a collective, community, or nation can be plausibly modeled as if it were a single optimizer.

When we model policy problems as if they involve maximization of some objective, we specify the constraints and incentive structures assumed or perceived by the policymaker. For example, we have to make assumptions, in discussing the setting of taxes, concerning how people react to higher or lower taxes. All of us who have worked in this area of formal public economics know how difficult it is to write down plausible descriptions of reactions by individuals to tax or transfer changes, or in other words the structure of incentives or disincentives. And we are here in the "inverse optimum" problem trying to guess what is in the mind of the decision-maker concerning incentives, rather than carrying out an empirical analysis of how people react to taxes or incentives.

Stern and Atkinson and Brandolini have provided fairly extensive reviews of tax/transfer "evidence on η," more than three decades apart.[60]

They both conclude that there is a huge range of possibilities depending on the case studies used and assumptions made about perceived technologies, constraints, and incentives. There are many examples of proposed social values for which η would be close to zero: many people appear to think that a dollar of purchasing power or income has the same value wherever it may be. Al Harberger famously argued, in discussing cost-benefit analysis in seminars in Oxford that I attended in the late 1960s, that "a dollar is a dollar is a dollar," on the grounds that it can be moved around.[61] Atkinson and Brandolini point out that there are examples where an attempt to use an inverse optimum problem to "explain" income transfer policies could make it look as if η were negative: cases where policies essentially transfer resources from poorer to richer.[62]

Some discussions of the problem of inferring η have been based on individual or collective savings. Dasgupta, for example, appeals to certain simple aggregative savings models without technical progress and concludes that η should be at least 2.[63] But again we find acute sensitivity of the estimate to model assumptions. In savings models, for example, if we assume exogenous technical progress, then we will have a far lower "optimum" savings rate for a given η than if we do not.[64] Thus to "explain" a savings rate of 30–40% we could infer η equal to 1 or 1.5 with exogenous technical progress or η equal to 3–4 without.[65]

A further attempted route for inferring η has been associated with standard models of choice under uncertainty where, in expected utility frameworks, η is the index of relative risk aversion (constant in this case).[66] In that context we can again find η ranging from negative to very large: there are many who accept unfair gambles (thus a convex utility function over some range and, in this framework, negative η) and others who appear very risk-averse for some decisions (high η). We should note, however, that in any case the expected utility model appears to perform badly as a vehicle for understanding many individual decisions.[67] Thus the risk/uncertainty route does not offer much help, either, in attempts to pin down η via an appeal to inverse optimum approaches. Nevertheless it seems to remain attractive to some. Barro has recently claimed that analysis of the "equity premium" in portfolio analysis tells him that η is around 3.[68] So Barro goes for 3 and Dasgupta for at least 2.

Weitzman appeared confident that 2 was a reasonable specification for η.[69] Layard appears convinced that η is close to 1.[70] I stay with my conclusion of 1977, shared by Atkinson and Brandolini, that such attempts to infer η cannot take us far.[71]

Interestingly, Weitzman has shown that the equity premium can be understood in terms of strong weight in the tails on equities,[72] i.e., results are strongly influenced by assumptions on the underlying distribution of random variables. Such discussions illustrate that an estimate is very sensitive to assumptions about what is in the mind of the decision-maker, even if, in the case of uncertainty, we claim that the expected utility model, or other such model, is a good one. A further major ethical leap is involved if, in the case of uncertainty, we try to pass from individual decisions under uncertainty to social decision-making. The relevance of the former for the latter has to be reasoned and is not at all clear. That problem is less severe if the inference is from the modeling of government decisions.

In the models often used in the context of climate change, η can perform three functions: (1) for modeling choice under uncertainty, (2) for modeling intratemporal distributional issues, and (3) for modeling intertemporal distributional issues. This is surely too much for one parameter, itself located in one particular utilitarian structure. Treating the issues separately but consistently is an important subject for research, with applications beyond climate change. The multiple roles of η in analyses of growth, inequality, and risk aversion are illustrated by the effects of higher η in the context of growth and risk in some of the calculations associated with *The Stern Review* and subsequent discussion. An illustration of the interaction of these growth and uncertainty effects is provided in *The Stern Review*.[73] With underlying growth and some given assumptions on the riskiness of the future, a higher η gives higher implied (proportional) costs of climate change because it involves greater aversion to risk. But if the future is taken as less risky, a higher η gives lower implied costs of climate change, because its greater aversion to inequality puts less weight on benefits to future generations if it is assumed they are better off, and thus heavier discounting.[74]

In *The Stern Review* we focused on η equal to 1 as an "official position" (see the "Green Book," the Treasury's handbook on project appraisal).[75] That a number is "official" does not by itself add greatly to its credibility as an attractive ethical position. We did offer some sensitivity analysis to η in *The Stern Review*.[76]

There is a separate question here on what the status of a value for η might be in an ethical argument, even if we found that many distributional social decisions (whether intra- or intertemporal) could be "explained" by some given value of η. Would that level of η be compelling as a way of capturing values for climate decisions? That is not entirely clear. I note here only that estimates of η in inverse optimum problems are "all over the place" and thus offer little guidance.[77]

Some might argue that it would be better to avoid all the η-inferring discussion, whether or not it gives clear results, and ask directly about values in relation to the climate problem at hand. In my view, that is too sweeping. Distributional values do matter greatly and are not easy to characterize. Thus it sometimes helps to think these values through in a structured way in simple tightly defined circumstances in order to understand them better. One can then ask whether those thought processes can help us in setting values for a more complex problem.

I have already expressed the view that, for climate change with its huge range of possible outcomes, many of them potentially extremely difficult or catastrophic, the expected-utility one-good framework can have only a minor role in the argument. It can give some useful insights, but we should not overly focus on it. Chapters 1 and 2 describe a strategic approach to the problem that recognizes, in chapter 1, the magnitude of the scientific and ethical issues at stake and looks at many dimensions of outcomes, outcomes where people can be much poorer than now, where there is loss of life, where there are very different impacts across the income spectrum, where there is loss of biodiversity, and so on. And that approach includes the analysis of chapter 2 which shows how such immense risks can be radically reduced.

Distributional values involve difficult ethical questions, and with climate change they are severe, particularly because consequences could involve catastrophic outcomes and loss of life on a major scale. Such questions require open discussion. That ethical discussion should be explicit and should be set in the context of analytical frameworks appropriate to the problem. It should not be confined to narrow formulations of conventional economic models, designed for questions of limited scope, simply on the grounds that we are familiar with them.

This is a section of doi:10.7551/mitpress/10408.001.0001

Why Are We Waiting?

The Logic, Urgency, and Promise of Tackling Climate Change

By: Nicholas Stern

Citation:

Why Are We Waiting?: The Logic, Urgency, and Promise of Tackling Climate Change
By: Nicholas Stern
DOI: 10.7551/mitpress/10408.001.0001
ISBN (electronic): 9780262329200
Publisher: The MIT Press
Published: 2016

The open access edition of this book was made possible by generous funding and support from the author

The MIT Press

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu

This book was set in Sabon 10/14pt by Toppan Best-set Premedia Limited. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Stern, N. H. (Nicholas Herbert)
Why are we waiting? : the logic, urgency, and promise of tackling climate change / Nicholas Stern.
pages cm. — (The Lionel Robbins lectures)
Includes bibliographical references and index.
ISBN 978-0-262-02918-6 (hardcover : alk. paper)
1. Climatic changes—Economic aspects. 2. Climatic changes—Government policy. 3. Environmental policy—Economic aspects. I. Title.
QC903.S833 2015
363.738′74—dc23
2014039907

10 9 8 7 6 5 4 3 2 1