Climatic Change (2009) 96:275–297
DOI 10.1007/s10584-009-9641-z

The economics of climate change: a post-stern perspective

Geoffrey Heal

Received: 7 November 2007 / Accepted: 8 June 2009 / Published online: 23 July 2009

Abstract What have we learned from the outpouring of literature as a result of the Stern Review of the Economics of Climate Change? A lot. We have explored the model space and the parameter space much more thoroughly. The Stern Review has catalyzed a fundamental rethinking of the economic case for action on climate change. We are in a position to give some conditions that are sufficient to provide a case for strong action on climate change, but we need more work before we have a fully satisfactory account of the relevant economics. In particular, we need to understand better how climate change affects natural capital—the natural environment and the ecosystems comprising it—and how this in turn affects human welfare.

1 Introduction

In November 2006 the U.K. government published *The Economics of Climate Change: The Stern Review*, written by a team led by Nicholas Stern (2006). The publication of the Stern Review provoked an unprecedented outpouring of papers on the same topic, including an entire issue of *The Economists' Voice* and large collections in the *Review of Environmental Economics and Policy* and in *World Economics*. And many of them were by very distinguished colleagues, so the Stern Review provoked not only quantity but quality too. What have we learned from all of this? Are there any emerging conclusions? In particular, what do we have to assume to make an economic case for prompt and significant action to reduce greenhouse gas emissions? This, it seems to me, is the really controversial issue,

G. Heal (✉)
Columbia Business School, Columbia University, 616 Uris Hall,
3022 Broadway, New York, NY 10027-6902, USA
e-mail: gmh1@columbia.edu
URL: www.gsb.columbia.edu/faculty/gheal

and the one that matters from a policy perspective. There is an amazing disjunction between economists and natural scientists on this issue: most natural scientists take it as completely self-evident that the consequences of climate change justify significant actions to mitigate the buildup of greenhouse gases, whereas there is a range of opinions on this matter among economists, with the conventional wisdom being until quite recently very different from that in the scientific community. Who is missing something important here—the economists or the scientists?

What we have learned from the recent debate, and what it takes to make a case for action on climate change, are the issues on which I focus. The recent debate has clarified many important issues, and we are in a position to give conditions that are sufficient to provide a case for strong action on climate change, but we need more work before we have a fully satisfactory account of the relevant economics. In particular, we need to understand better how climate change affects natural capital—the natural environment and the ecosystems comprising it—and how this in turn affects human welfare. I take some first steps in this direction.

2 Welfare economics and climate change

Let me begin with the basic economic theory of climate change. The first topic I want to spend time on is the discount rate, but before that there is a simple, important and interesting point that Duncan Foley has recently emphasized (2007). The emission of greenhouse gases is a massive negative external effect—the Stern Review refers to it as possibly the greatest market failure in history. Foley's point is that with such a large uninternalized externality, the business as usual scenario with no action on climate change obviously cannot be Pareto efficient, so if we move to correct the externality it must in principle be possible to make a Pareto improving (or "win–win") change to the world economy. If we do this then there is in aggregate no net cost to correcting climate change: the gains must outweigh the costs so that the gainers could compensate the losers and still gain. We can all come out ahead—whether we do is a matter of institutional design, on which many people are now working. The numbers in the Stern Review support this point, indicating that the gains from action on climate change greatly outweigh the costs, but the point would be valid whatever the numbers.

3 Discount rates and the environment

Now to discount rates. As anyone who has spent even a short time on this issue must be aware, one of the controversial issues is the choice of a discount rate. By this we mean the pure rate of time preference (PRTP), to be distinguished clearly from the consumption discount rate (CDR). The PRTP is the δ in the expression $\int_0^\infty u(c_t)e^{-\delta t}dt$ where c_t is aggregate consumption at time t, u is a utility function showing strictly diminishing returns to consumption and we are summing discounted utility over all remaining time. The other discount rate concept, the CDR, is the rate of change of the present value of the marginal utility of consumption, that is, the rate of change of $\frac{e^{-\delta t}du(c_t)}{dc_t}$. For the case of a single consumption good—and we will turn to the case of multiple goods later—it follows from well-known arguments going back to Ramsey (1928) (see Heal 2005 for a review) that this is equal to the PRTP plus the rate of change of consumption times the elasticity of the marginal utility of consumption:

$$\rho_t = \delta + \eta\,(c_t)\,R\,(c_t) \tag{1}$$

where ρ_t is the consumption discount rate applied to consumption at time t, $\eta\,(c_t) = -\frac{cu''}{u'} > 0$ is the elasticity of the marginal utility of consumption and $R(c_t)$ is the rate of change of consumption at time t. (Here $u' = \frac{du(c)}{dc}$ and $u'' = \frac{d}{dc}u'$.)

What do these two discount rates mean? The PRTP δ is the rate at which we discount the welfare of future people *just because they are in the future*: it is, if you like, the rate of intergenerational discrimination. Note that there are at least two reasons why we may wish to value increments of consumption going to different people differently: one is that they live at different times, which is captured by δ, and the other is that they have different income levels, which we discuss shortly.[1] A PRTP greater than zero lets us value the utility of future people less than that of present people, *just because they live in the future rather than the present*. They are valued differently even if they have the same incomes. Doing this is making the same kind of judgment as one would make if one valued the utility of people in Asia differently from that of people in Africa, except that we are using different dimensions of the space–time continuum as the basis for differentiation.

There may be an argument for discounting future benefits if we are uncertain that the future will exist, if we are worried about an Armageddon at some point within our timescale. Stern uses this argument to justify a very low pure rate of time preference. Civilization might, for example, be destroyed by a meteor hitting the earth. This point seems to be valid conceptually but for simplicity I am neglecting it.

That an increment of consumption is less important to a rich person than to a poor person has long been a staple of utilitarian arguments for income redistribution and progressive taxation (see Sen 1973), and is almost universally accepted. This is reflected in the diminishing marginal utility of consumption, and the rate at which marginal utility falls as consumption rises is captured by $\eta(c_t)$. Equation 1 pulls together time preference and distributional judgments: the rate at which the value of an increment of consumption changes over time, the CDR ρ_t, equals the PRTP δ plus the rate at which the marginal utility of consumption is falling. This latter is the rate at which consumption is increasing over time $R(c_t)$ times the elasticity of the marginal utility of consumption $\eta(c_t)$.

Note that if consumption were *falling* rather than *rising* over time, then the second term in the expression for ρ_t would be negative and the CDR could in principle be negative, that is the value of an increment of consumption could be rising over time rather than falling. We would not be discounting but doing the opposite, whatever that is. It is not impossible that in a world of dramatic climate change and environmental degradation, consumption might fall at some point. It is even more likely that some aspects of consumption would fall while other continue to rise—recognizing this requires that we treat consumption as a vector of different

[1] We could also value them differently for all manner of other reasons - differences in nationality, ethnicity, and proximity either physically or genetically. In general we don't do these things, at least explicitly, which to me makes it strange that we do explicitly discriminate by proximity in time.

goods that can be affected differently by climate change. For an early recognition of this point see Fisher and Krutilla (1975), who comment that increasing scarcity of wilderness areas may drive up our valuation of them. A more detailed analysis in the context of a growth model is in Gerlagh and van der Zwaan (2002), who make the interesting point that with limited substitutability between environmental and manufactured goods and the growing scarcity of environmental goods, there is likely to be a version of Baumol's disease—an ever larger portion of income being spent on non-manufactured goods.

Let's follow this line of thought and disaggregate consumption at date t into a vector $c_t = (c_{1,t}, c_{2,t}, ...c_{n,t})$ of n different goods. Utility is increasing at a diminishing rate in all of these and is a concave function overall. In this case we have to change Eq. 1 for the consumption discount rate. Now there is a CDR for each type of consumption and we have n equations like Eq. 1, with a CDR for each good i equal to the PRTP plus the sum over all goods j of the elasticity of the marginal utility of consumption of good i with respect to good j times the growth rate of consumption of good j:

$$\rho_{i,t} = \delta + \eta_{ii}(c_t) R(c_{i,t}) + \sum_{j \neq i} \eta_{ij}(c_t) R(c_{j,t}) \tag{2}$$

where $\rho_{i,t}$ is the CDR on good i at date t, $R(c_{i,t})$ is the rate of change of consumption of good i at date t, and η_{ij} (c_t) is the elasticity of the marginal utility of good i with respect to the consumption of good j (see Heal 2005 for details: the most general framework of this type can be found in Malinvaud's classic 1953 paper (1953)). The own elasticities such as η_{ii} (c_t) are positive numbers, but the cross elasticities η_{ij} (c_t), $j \neq i$, are zero if the utility function is additively separable and can otherwise have either sign.

As an illustration consider the constant elasticity of substitution utility function

$$\left[\alpha c^{\sigma} + (1-\alpha) s^{\sigma}\right]^{\frac{1}{\sigma}} \tag{3}$$

Here we can think of c as produced consumption and s as an environmental stock that produces a flow of services—more on this later. In this case the cross elasticity of the marginal utility of consumption depends on whether c and s are substitutes or complements. For an elasticity $\sigma > 1$ they are substitutes and the cross elasticity is positive, and vice versa.

Let's test our intuitions on this. Take the case where the environmental good and produced consumption are highly complementary, so that indifference curves are near to right angled and the elasticity σ is close to zero. Then the cross elasticity is negative. This means that if the consumption of the environmental good is rising then this reduces the consumption discount rate on the regular good. Conversely if the availability of the environmental good is falling then this raises the consumption discount rate on the consumption good. These results make sense: because of the assumed complementarity, an increase in the amount of the environmental good will raise the marginal utility of the consumption good and so tend to lower the consumption discount rate, and vice versa. Of course, the own elasticity on the environmental good is positive so that if the availability of this good is falling then this will tend to make its own consumption discount rate negative.

Whether produced goods and environmental services are substitutes or complements in consumption is not an issue that has been discussed in the literature, as with the few exceptions mentioned above people have worked with one-good models. There do however seem to be reasons to suppose that complementarity is the better assumption, with $\sigma < 1$. Dasgupta and Heal (1979), following Berry et al. (1978), suggest that in production there are technological limits to the possibility of substituting produced goods for natural resources. In particular we invoke the second law of thermodynamics (Berry and Salamon are thermodynamicists) to suggest that if energy is one of the inputs to a production process, then there is a lower bound to the isoquants on the energy axis. Similarly one can argue that certain ecosystem services or products, such as water and food, are essential to survival and cannot be replaced by produced goods. There are therefore lower bounds to indifference curves along these axes, implying if the utility function is CES that $\sigma < 1$.

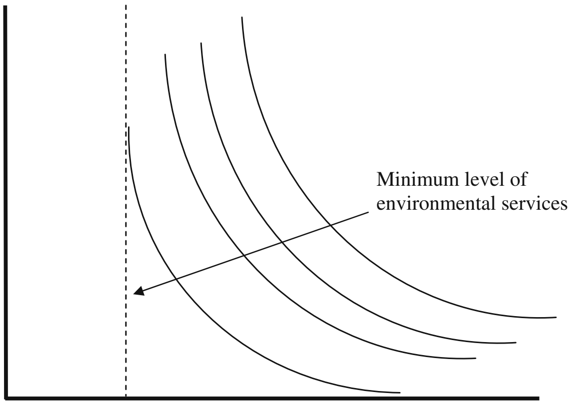

The figure illustrates this idea: it shows indifference curves for a two-argument utility function, consumption of produced goods and of ecosystem services, as in Eq. 3 above. There is a minimum level of ecosystem services needed for survival—think of this as water, air, and basic foodstuffs, all of which are ultimately produced from natural capital. For low welfare levels there is no substitutability between these and produced goods, so that indifference curves are close to right angled. At higher welfare levels where there are abundant amounts of both goods there is more scope for substitution. Taken literally, this implies that the elasticity of substitution is not constant but depends on and increases with welfare levels. This of course is not reflected in the CES function such as 3. A function with these properties is

$$\left[\alpha c^{\sigma} + (1-\alpha)(s-\varepsilon)^{\sigma}\right]^{\frac{1}{\sigma}} \tag{4}$$

which is simply the CES function we noted before, with the zero of the ecosystem service axis transformed by $\varepsilon > 0$. Utility is not defined for $s < \varepsilon$. Relative to the transformed origin $(\varepsilon, 0)$ there is still a constant elasticity of substitution σ but relative to $(0, 0)$ the elasticity is not constant. For $\sigma > 1$, every indifference curve, every welfare level, can be attained with only ε of ecosystem services, whereas with $\sigma < 1$ greater welfare levels require greater levels of ecosystem services (and of consumption goods).

One more theoretical point is worth making in this discussion of discount rates, and it relates to the connection if any between the discount rate and the rate of return on capital or the rate of interest ruling in capital markets. Going back for a moment to the case of a one-good model with just aggregate consumption in the utility function, as in Eq. 1, we can easily show that under certain assumptions the CDR equals the rate of return on capital (ROC):

$$\rho_t = \delta + \eta(c_t)R(c_t) = f'(k) \tag{5}$$

where k is an aggregate capital stock, $f(k)$ the aggregate production function and f' its first derivative or the marginal product of capital (see Heal 2005). What are the assumptions we need to make to get to this conclusion? We need the full panoply of perfect market assumptions—no market failures such as external effects or public goods, and an economy where agents have perfect foresight for the whole of future time and all markets clear at all dates, so that the economy follows an optimal path in the Ramsey sense over time. Taxes are also a drawback, as they drive a wedge between rates of return.

This relationship between the CDR and the ROC becomes more complex if there are several goods in the utility function, as in Eq. 2. In the case of just two goods as in the CES function in Eq. 3, we have that

$$\rho_{c,t} = f' + \eta_{cs}R(s) \tag{6}$$

In this case the relationship between the CDR on the consumption good and ROC f' depends on whether the environmental good and the consumption good are complements or substitutes ($\eta_{cs} < 0$ or > 0) and on whether the availability of the environmental good is growing or falling.

Karl-Goran Mäler has made an interesting point about consumption discount rates and natural capital: Stern and many others working in this area take the rate of growth of consumption as exogenous and match it to historical records. But there is evidence that many estimates of the rate of productivity growth are biased upwards, because the growth accounts on which they are based omit the depletion of natural capital. Hence growth forecasts based on these will be overly optimistic, biasing the consumption discount rate upwards.[2] In particular, Nordhaus's estimates of the exogenous rate of growth of total factor productivity in DICE will be too high.

3.1 Choosing a discount rate

Having set out some of the relevant theory, let's return to the issue of the choice of a discount rate, and see whether there are any non-controversial points everyone can agree on. The first point to note is that there is a big difference between the two discount rates we have spoken of—the PRTP and the CDR. The PRTP is exogenous to the economic problem, and its choice is an ethical act, a decision on the relative weights of different generations of human beings. The CDR, in contrast, is in part endogenous: it depends on δ certainly which is exogenous but also on $R(c)$, the consumption growth rate, which is clearly endogenous and the outcome of the operation of the economic system, and also on $\eta(c)$ which is again endogenous and depends on the level of consumption. (These statements are true whether c is a scalar

[2]Personal communication, Karl-Goran Mäler, 2008.

or a vector.) Of course the choice of a form for the utility function, and therefore of the value of $\eta(c)$, is also an ethical choice. How much less valuable is \$100 to a rich person rather than to a poor person? η answers this question, which is normative.

The bottom line, then, is that we can't select either the PRTP or the CDR without making ethical judgments. My own personal ethical judgment on the PRTP is that it should be zero, and in this I am in the same camp as many British economists—Frank Ramsey (1928) famously commented that "discounting future utilities is ethically indefensible and arises purely from a weakness of the imagination," Roy Harrod spoke in similar terms (1948), and the philosopher Henry Sidgwick commented (1890, page 412) that "It seems ... clear that the time at which a man exists cannot affect the value of his happiness from a universal point of view; and that the interests of posterity must concern a Utilitarian as much as those of his contemporaries, except in so far as the effect of his actions on posterity—and even on the existence of human beings to be affected—must necessarily be more uncertain." This is an ethical judgment not a theorem so you don't have to agree, but I personally find it difficult to see any reason for valuing future people differently from present people just because of their futurity.

Note that in choosing a PRTP we do not take into account the fact that future people may be richer or poorer than us—the CDR deals with that via the terms such as $\eta(c_t)\,R(c_t)$. In choosing δ the issue is quite simply whether we want to discriminate against future people. I have never seen a convincing explanation of why this is the right thing to do.

I think there is often a lot of confusion between δ and ρ: as we can see from Eq. 1 setting $\delta = 0$ *does not imply* that the CDR is also zero. In the one-good world, provided consumption is rising the CDR will be positive even if the PRTP is zero.

Operationally what is the difference between these two discount rates? When do we use δ and when ρ? To pick a CDR you need a consumption growth rate. So you can't use this rate when the consumption path is endogenous: you need a growth profile for the economy before you can pick the CDR. If you are working with a model where the consumption growth rate is endogenous, you should be using δ the PRTP: together with the details of your model this will determine the CDR. More generally, if you are working with a general equilibrium model of the entire economy such as a Ramsey model or a dynamic multisector general equilibrium model, in which case consumption is clearly endogenous, you have to pick a PRTP. The model will then give you a set of CDRs. On the other hand, if you are doing a partial equilibrium analysis where the time path of consumption is exogenous and not affected by the choices you are studying, then you should use the CDR to evaluate future costs and benefits relative to the present. This rate tells you precisely what you need to know in this case—the rate at which the value of a marginal increment of consumption is changing over time. Because it is partial equilibrium your project will make only marginal changes about the pre-existing path.

Most of the models that have been widely used to analyze the economics of climate change, integrated assessment models or IAMs, are dynamic general equilibrium models and so need to be supplied with a PRTP. They will then determine a consumption discount rate as part of the solution. So the DICE model developed by Nordhaus (1993), probably the most widely used IAM, needs to be given a pure rate of time preference to complete its specification. This means that numbers such as the return on capital are irrelevant when choosing the discount rate for DICE: this is a purely ethical choice. The model will solve for a CDR which will reflect the rate of growth of consumption, the elasticity of marginal utility, and the productivity of capital.

It is worth noting that the CDR is not constant: however this does not lead to dynamic inconsistency. The CDR is derived from a model in which the PRTP is specified, and if that specification is as a constant, then the underlying equilibrium of the economy is dynamically consistent and so will remain so after any small perturbations about it, and it is to the evaluation of such small perturbations that the CDR is applied.

3.2 Discount rates and the rate of return

Let's return to the relationship between the CDR and the return on capital, as this has been widely discussed in the literature. As we saw in Eq. 5, for the aggregate one-good model this is equal to the return on capital under certain assumptions. These are very strong indeed, and it is not clear to me that they are relevant in the context of climate change. On the contrary, I think it's rather clear that they are not. One assumption that we need to get this equality is that there is no market failure, yet as the Stern Review observed, climate change is probably the largest external effect and market failure in history. So is it appropriate to assume no market failure in evaluating a consumption discount rate for a model of climate change? This seems to be a rhetorical question![3] Also required to get Eq. 5 is an assumption of perfect foresight in all capital markets indefinitely far into the future. Again this is not an assumption that I personally find persuasive, particularly at a time of crisis in capital markets. These issues are important because both Nordhaus (2007) and Weitzman (2007a) have tried to argue for or against choices of the PRTP from observed values of the long-run return to capital. They take Eq. 5, put in what they think of as reasonable numbers for η and $R(c)$, and then judge the correctness of a value for δ by whether $\delta + \eta(c_t)R(c_t) = f'(k)$. There are lots of things questionable here. One is that as noted we need very strong assumptions for this equation to hold, assumptions that seem to be particularly out of place in a discussion of climate change. There is also the rather fundamental fallacy of deducing an "ought" from an "is," recognized as a philosophically unacceptable step for the last two and a half centuries, since the time of Hume.[4] These concerns seem to me to sink the whole endeavor of linking δ to f'. But suspending disbelief for a moment, you can argue that it is using Eq. 5 in the

[3]We can readily extend the analysis of previous sections to reflect the externality associated with climate change. Suppose for the sake of a simple illustration that climate change is driven by the stock of capital, so that the change in the environmental stocks is $ds/dt = h(k)$ where $h' < 0$ and k is a capital stock. Then we can derive the following equation relating the return on capital to the CDR:

$$f_k = \delta + \eta_{cc}R(c) + \eta_{cs}R(s) - h'(k)\frac{p_s}{p_k}$$

where p_s and p_k are the shadow prices on the environmental stock and capital stock respectively.

[4]Hume was not the only C18th philospher to make this point. Less well-known but rather clear are Kant's comments on this point: "First, as concerns the source of metaphysical knowledge, its very concept implies that they cannot be empirical. Its principles .... must never be derived from experience. It must not be physical but metaphysical knowledge, namely, knowledge lying beyond experience." Kant (1950), page 13.

wrong way. Instead of going from f', η and R to δ, we should go the other way. The meaning of δ is clear and I think it is obvious what its value should be—zero. I think we know this with some confidence. But we don't know η and we probably have at best rough estimates of R and f'. If this equation were convincing and relevant there is a case for using it to go from δ to η, R and f' rather than the other way round.

There is of course another shortcoming in using Eq. 5 to make statements about this pure rate of time preference, and another reason why its relevance is at very best limited, which you can see immediately from Eq. 6. This is that it does not apply if there is more than one argument to the utility function, which is the case if it is important to distinguish between several different elements of consumption, because for example they will behave differently in response to climate change, as seems most likely to be the case. To develop this point we need to talk about natural capital and ecosystem services, a discussion we began in Section 3. As the World Bank recently emphasized, natural capital is an important part of the wealth of nations (World Bank 2006). Ecosystem services are some of the returns that come from that capital, and are of fundamental importance to human societies (Barbier and Heal 2006). Climate change will deplete natural capital and reduce the flow of ecosystem services. Both the Stern Review and the IPCC Fourth Assessment Reports (2007) are in agreement that this may in the long run be the most important and most costly aspect of climate change. That is why in Eq. 3 I have distinguished two types of consumption, one conventional consumption c and the other a stock of environmental or natural capital s that generates a flow of services, many of which may be essential for human survival. This is not a new distinction: it started with Jeffrey Krautkramer (1985), was developed by Scott Barrett (1994) and Heal (1998), and has been emphasized recently by Roger Guesnerie (2004) and by Thomas Sterner and Martin Persson (2007). On the assumption that s is falling over time because of climate change, we see from 6 that

$$\rho_{c,t} = f' + \eta_{cs} R(s)$$

Here $R(s) < 0$ but we don't know the sign of η_{cs}, which as noted above is positive if consumption and the environment are substitutes and negative if they are complements. (Arguments given earlier suggest that it might be negative.) This means that even in an implausible fully first best world with perfect foresight and all interactions captured by the market, we cannot be sure whether the discount rate on consumption goods should exceed or be less than the return on capital.

An issue that is sometimes raised is that if the PRTP is zero, then this leads to an inefficient allocation of capital if, as is surely the case, the rate of return on investment in the private sector is positive. Statements such as this are based on a confusion of the different discount rates: a zero PRTP does not imply that the rate of return on investment of any sort is zero. Far from it: the test discount rate to be applied to a small project in a first best context is the CDR, not the PRTP, which can easily be positive, indeed generally will be, even when the PRTP is zero. So this argument about the inefficiency of a zero PRTP is simply wrong.

Also wrong is the argument that observed savings rates are inconsistent with a PRTP of zero: the claim is that a zero PRTP would imply savings rates greatly in excess of those observed. This is just another version of the argument that the PRTP should equal the ROC: without that equality, one cannot deduce a high savings rate from a low PRTP.

3.3 A Sterner perspective

It's worth looking in more detail at the Sterner and Persson development of this point. They talk about the effect of changes in relative prices rather than consumption of produced and environmental goods, but the point is the same. If we consume both produced goods and the services of the environment, as in the utility function 3, then we can expect that with climate change environmental services will become scarce relative to produced goods and therefore their price will rise relative to that of produced goods (the "environmental Baumol disease" that Gerlagh and van der Zwaan refer to (2002)). Consequently the present value of an increment of environmental services may be rising over time, and the consumption discount rate on environmental services may thus be negative, precisely the point that we were making in Eq. 2 above. This could be the case even with a high PRTP, which is the main point of the Sterner and Persson paper. They also present an interesting modification of Nordhaus's DICE model to incorporate this point. They replace the standard utility function, which is an isoelastic function of aggregate consumption, by a CES function along the lines of Eq. 3 above, but modified to reflect a constant relative risk aversion:

$$\left[(1-\gamma)\,c^{1-1/\sigma}+\gamma s^{1-1/\sigma}\right]^{(1-\alpha)\sigma/(1-\sigma)}\Big/(1-\alpha)$$

(This formulation does not reflect the need for a minimum level of ecosystem services discussed above.) They assume that the supply of environmental services s is negatively affected by temperature according to the square of temperature, and that the share of environmental goods in consumption is about 20%, use these assumptions to calibrate the modified DICE model and and then run the model with the PRTP used by Nordhaus. Their runs show that even with such a high PRTP the presence of an environmental stock that is damaged by higher temperatures radically transforms the optimal emissions path of CO_2 and leads to a vastly more conservative policy towards climate change, with emissions both staying lower and falling faster. In fact it leads to a more aggressive reduction in greenhouse gases than recommended by the Stern Review.

4 Uncertainty

4.1 Uncertain discount rates

Martin Weitzman has argued (1998) that because there is so much argument about the right discount rate, we should think of the discount rate as being uncertain. He investigates what this implies, and makes an interesting observation, which is that when there are several possible discount rates and we don't know which is the "right" one, then in the long run we should work with the lowest of all the possible rates. Here is a simple example to illustrate this point. Suppose there are two possible discount rates δ_1 and δ_2 with $\delta_2 > \delta_1$. Suppose also that the chance that δ_1 is the correct rate is p. Then the discount factor to apply to a future date t is

$$pe^{-\delta_1 t}+(1-p)\,e^{-\delta_2 t}$$

and its rate of change, the instantaneous discount rate, can be written

$$-w_1\delta_1 - w_2\delta_2$$

where w_1 and w_2 are weights summing to one with the ratio w_2/w_1 going to zero as time goes to infinity.[5] Hence the long-run discount rate is δ_1, the lowest possible discount rate. If that is zero, then according to this argument we should use a zero discount rate to apply to the far distant future.[6]

My own judgment is that the right PRTP is zero, and I could use Weitzman's argument to justify the use of a zero PRTP in the long term—it is certainly one of the possible rates and clearly the lowest—but have never actually been tempted to do so. I think that this is because, although Weitzmans's result is undoubtedly technically correct, I am not totally certain of its philosophical foundations and implications. If we disagree over the PRTP, does it make sense to randomize across all the rates that are suggested? There may be assumptions under which this is the correct thing to do, but they are not part of Weitzman's paper and they are not obvious. It does seem that with disagreements about the PRTP then we have a social choice problem. Another issue that concerns me and makes me reluctant to put this idea into practice is that it gives a non-constant PRTP, one that changes over time and converges asymptotically on the lowest rate suggested. Any intertemporal plan constructed with such a rate will be dynamically inconsistent in the sense that if we follow it for a period and then stop and ask what is the best continuation from where we are, it will not be the plant that we originally adopted. I am not sure if this matters, but it certainly makes me pause and think.

4.2 Data and model uncertainty

Uncertainty is one of the dominant facts in any analysis of climate change. It's not that the underlying science is uncertain: the mechanisms through which the accumulation of greenhouse gases in the atmosphere warm the earth are in principle well-understood. The point is that there is some debate about whether they will lead to a temperature rise of 2°C or 6°, or possibly even more. There is also some uncertainty about the implications of any given temperature change for climate in the more general sense of patterns of precipitation, winds, humidity, etc.

Furthermore, even if we were to know accurately and in detail how the climate is going to change, we would still not understand fully the implications for social and economic activity (see Heal and Kriström 2002). An additional problem is that the type of uncertainty we face here does not fit the economists' traditional model of decision-making under uncertainty. The traditional model presumes that we have a known state space, a known (objective or subjective) probability distribution, and an expected utility function. In the context of climate change, we certainly don't have a known probability distribution: we have some information about the relative likelihood of different outcomes but usually not enough to form a full-fledged

[5] More formally $w_1 = pe^{-\delta_1 t}/\left(pe^{-\delta_1 t} + (1-p)\,e^{-\delta_2 t}\right)$ and $w_2 = pe^{-\delta_2 t}/\left(pe^{-\delta_1 t} + (1-p)\,e^{-\delta_2 t}\right)$ so $w_1 + w_2 = 1$ and $\lim_{t\to\infty} w_2/w_1 = 0$ implying that $\lim_{t\to\infty} w_1 = 1$.

[6] A referee commented that if we don't know the correct PRTP then a natural assumption is that this is continuouly distributed, in which case there is typically no lowest value.

probability distribution. In the terminology of Frank Knight, we are dealing with uncertainty rather than risk. In more contemporary terminology, we are in the world of ambiguity, where probabilities are not known. This raises some complex issues, but to start with we look at some of the more straightforward aspects of risk, where probabilities are assumed to be known.

There is an analog of Eq. 1, $\rho_t = \delta + \eta(c_t)R(c_t)$, which addresses uncertainty in the growth rate $R(c_t)$. Assume the growth rate at time t is a random variable drawn from a distribution that is identically and independently normally distributed at each time accoeding to $N(\mu, \sigma^2)$. Then the equivalent to 1 is just

$$\rho_t = \delta + \eta(c_t)R(c_t) - \frac{1}{2}\mu^2\sigma^2 \tag{7}$$

so that the CDR is reduced as both the mean and the variance of the growth rate increase.

Another point about uncertainty that has been made by a number of authors is that climate change and many of its consequences are irreversible.[7] Melting of ice sheets and glaciers, extinction of species, and destruction of corral reefs, are all irreversible, at least on a timescale relevant to human societies. Furthermore, as there are many things we do not know about climate change and its social and economic consequences, there is a real chance that we shall learn over time about the costs and benefits of climate change. In such a situation, there is a real option value associated with preserving the current climate. This is an argument for conservation, but not one that we can easily evaluate in quantitative terms.

4.3 Risk aversion

There is of course a more conventional risk-aversion argument for mitigating climate change. Uncertainty about the consequences of climate change means that we bear a risk, and risk-bearing is an activity to which most of us are averse. We pay insurers to bear our risks for us. I present next some simple yet suggestive calculations from Heal and Kriström (2002) indicating the importance of risk aversion. These calculations were addressed to the question: what cost is it worth incurring to avoid the risk of climate change? Within a simple framework we can carry out calculations that illustrate the issues involved, and how discount rates, risk aversion and probabilities interact.[8]

Denote society's income in the absence of climate change by I and the benefits derived from this income by utility $u(I)$. The expected utility after climate change is $\Sigma_j p_j u(I_j)$.[9] Climate change occurs, if it occurs at all, in year C. Denote by $\delta \leq 1$ the weight given to costs or benefits at date $t+1$ relative to those at t, so that δ^{t-1} is the weight given to those at t relative to those at 1. Then $(1-\delta)\times 100$ is the discount rate as a percent.

[7] See Fisher and Narain (2002), and for a survey see Heal and Kriström (2002).

[8] The calculations that follow are taken from work by Geoffrey Heal and Yun Lin.

[9] If there is climate change, then income drops from I to I_j with probability p_j, where clearly $I_j \leq I$ and $\Sigma_j p_j \leq 1$. This can be weakened to allow an income drop on the average, in order to include a possibility for income increases. With risk aversion, people would still be willing to pay to avoid the change. Thanks to Mark Machina for pointing this out.

Suppose that it possible by incurring a cost from now to the date C at which climate change might occur, to rule out this occurrence. What cost x is it worth our while incurring, from now to C, in order to ensure that the climate does not change at C? What is our willingness to pay to avoid climate change? The number x that we seek is the solution to the equation

$$\sum_{t=1}^{C} \delta^{t-1} \left[u(I) - u(I-x)\right] = \sum_{t=C+1}^{T} \delta^{t-1} \left[u(I) - \sum_{j} p_j u(I_j)\right] \quad (8)$$

The left hand side is the loss of utility in incurring the cost x from now to the time C of climate change, with future losses discounted back to the present: the loss each year is $[u(I) - u(I-x)]$, and we sum this, discounted, over all years up to C. The right hand side is what we would lose each year, in expected value terms, if climate change were to occur, summed from its occurrence at C to a distant date T, and again discounted to the present. The expected annual loss is $[u(I) - \Sigma_j p_j u(I_j)]$. This sum on the right is therefore also the benefit of avoiding climate change. The maximum we should be willing to pay is the value of x at which these two are equal: hence the equation. The date T is the maximum time horizon that we consider relevant to these calculations.

As a concrete illustration, we can think of x as the extra cost of moving as fast as possible to energy based on non-fossil sources, such as solar, geothermal or biomass. As these technologies develop, this cost will decline: we assume that it is zero by the time at which climate change would occur, which in the illustrative calculations is taken to be 50 years hence. Obviously there are some heroic assumptions here. Climate change is taken to be a discrete event. Preventive expenditures are assumed to be constant. But nevertheless the numbers may be indicative.

Below we present values of x for some illustrative parameter values and indicate their sensitivity to the assumptions. What we should be willing to pay, x, is expressed as a percent of the income level[10] I, which is taken to be 10. The calculations are only illustrative: we do not know enough about the costs or probabilities of climate change to make presenting a best estimate of x a useful exercise. The key conclusion is that for some parameter values that must be within the set considered possible, one might wish to spend up to 8.13% of national income on avoiding climate change. For other parameter values that are also possible, the number may be 0.1%. Even this is a big number in absolute terms. The most critical parameter in these calculations is an economic parameter, the PRTP. The index of risk aversion η is also very influential.

A reasonable functional form for the utility functions $u(I)$, widely used in empirical studies of behavior under uncertainty, is the family of functions displaying constant relative risk aversion: the index of relative risk aversion (IRRA) for $u(I)$ at income I is $-Iu''/u'$. A reasonable range of empirical values for the index of relative risk aversion is from 2 to 6.

We work with three possible distributions of losses from climate change, called A, B and C and summarized in Table 1. In the first there is a 20% chance of a loss of

[10]Because of the choice of functions for which the IRRA is constant, the ratio x/I is independent of the value of I, so that we do not need to think hard in choosing a value for I.

Table 1 Alternative probability distributions

Loss	Probability		
	A	B	C
2%			0.24
5%	0.2	0.24	0.1
10%			0.01
15%	0.1	0.10	
20%		0.01	
25%	0.05		
0	0.65	0.65	0.65
E. Loss	3.75%	2.99%	0.99%

income of 5%, a 10% chance of an income loss of 15% and a 5% chance of a loss of 25% as a result of global warming. By implication, there is a 65% chance of no loss at all, and the expected loss is 3.75% of current income. This figure for the expected loss is generally consistent with the IPCC's estimate of the loss from climate change from its third assessment report, and rather more conservative than the estimates from its most recent report or from the Stern Review. The other two cases are even more conservative: the possible losses are lower and the probabilities are concentrated more at the low end of the distribution. In the most conservative case, the expected loss is fractionally under 1% of income, with a probability of 0.24 of an income loss of 2%, a probability of 0.10 of a loss of 5% and a probability of 0.01 of a loss of 10%.[11] Table 2 reports the values of x in the equation above for alternative combinations of the discount rate in percent, denoted δ, and the index of relative risks aversion.

As mentioned, the date for climate change C is assumed to be 50 years, and we take the upper limit of the sum of benefits T to be 1,000.

Table 2 reports the results of solving the equation for x for probability distribution A and a range values for the discount rate (from 1% to 5%) and for the IRRA (from 0 to 6). (For results for the other distributions see Heal and Kriström 2002).

4.4 Unknown probabilities

Some interesting papers have recently tried to tackle directly the point that we don't know the probability distribution governing the consequences of climate change. Claude and Marc Henry have looked at this issue from the perspective of ambiguity theory (Henry and Henry 2002), the theory of choice in the face of incomplete probability distributions (see also Claude Henry 2007), and Martin Weitzman has looked at it from the perspective of Bayesian updating of unknown parameters, including those of the probability distribution (Weitzman 2007b).

In each case the assumption is that we don't know the distribution of damages from climate change: in the case of the Henry paper, the lack of knowledge is of a very general kind, with some information about the relative likelihoods of different regions of the state space but not enough to form a probability distribution. Decision problems in this framework have been analyzed in an axiomatic framework similar

[11] For an interesting review of the available evidence on the probabilities of loss from climate change see Roughgarden and Schneider (1999), who take a range of expert opinions and fit a systematic probability density function to these.

Table 2 Willingness-to-pay for distribution A

IRRA	0	1	2	3	4	5	6
δ							
1	5.74	6.07	6.42	6.81	7.22	7.66	8.13
2	2.15	2.32	2.50	2.72	2.96	3.23	3.54
3	1.05	1.13	1.23	1.35	1.48	1.64	1.82
4	0.56	0.61	0.66	0.73	0.81	0.89	1.00
5	0.31	0.34	0.37	0.41	0.45	0.50	0.56

to that introduced by von Neuman and Morgenstern for expected utility theory by Ghiradato and Marinacci (2001). The central result of the Henry paper is that it in such a context it would be wrong to use standard expected utility theory and neglect the ambiguity. They see this as providing a limited degree of support for the much-debated "precautionary principle."

Weitzman assumes that we are uncertain about one of the key parameters driving climate change, the climate sensitivity parameter, and models the process of learning about this. I will attempt to summarize his analysis in the context of a simple example that he uses in his paper. This focuses on the parameter relating the change in atmospheric greenhouse gas concentrations to the change in global mean temperature. Climate scientists model this relationship as

$$\Delta T \approx \frac{S}{\ln 2} \Delta \ln CO_2 \tag{9}$$

where T is global mean temperature and S the climate sensitivity parameter. Weitzman assumes that the best estimates of S are uncertain. He works with an isoelastic utility function

$$u(c) = \frac{c^{1-\eta}}{1-\eta}$$

where $\eta > 1$. He uses a two-period model where consumption in the first period is normalized to unity and the growth rate of consumption from present to future is Y B ln$c$ where c is second period consumption. This growth rate depends on the climate sensitivity and for very large values of this parameter could be negative—if the increase in temperature is so large as to lower consumption (a possibility discussed in the context of discount rates above). In general the distribution of Y depends on that of S, with the left tail of Y corresponding to the right tail of S: large values of the climate sensitivity parameter lead to low or negative growth rates. The expected present value of an increment of consumption in the second period is given by

$$E(M) = \beta e^{-\eta Y}$$

where $\beta \in (0, 1]$ is the one-period discount factor and M is the amount of consumption we would give up today to get one unit next period with certainty. Weitzman assumes that the growth rate Y is normally distributed as $N(\mu, \sigma^2)$ and that we know μ but not σ, which we have to estimate from data on the history of the climate system. In this case

$$E(M) = \exp\left(-\delta - \eta(c_t) R(c_t) + \frac{1}{2}\mu^2\sigma^2\right) \tag{10}$$

by the same arguments as gave Eq. 7 above. So far this is all straightforward. Next we look at the consequences of not knowing σ: to develop these Weitzman adopts a Bayesian approach, and assumes a prior distribution of σ. This he takes to be a uniform distribution of $\ln \sigma$ on $(0, \infty)$. With these assumptions he shows that $E(M) = +\infty$, so that we are willing to give up an unlimited amount of consumption today to assure one unit in the future. We are, in other words, extraordinarily risk-averse about the future under these assumptions.

This is a somewhat surprising result. Technically it comes from the fact that a range of well-behaved distributions such as power law distributions have infinite means. Weitzman assumes that Y is distributed normally, but when he adds in that the variance of the normal is unknown with a noninformative prior then the result is a distribution of Y that has some of the properties of power law distributions, namely it has a "fat tail," meaning that there is much more weight in the tails of the distribution than there is in the normal and in most of the distributions we usually work with. This in turn means that the risks of extreme outcomes—very large temperature changes and consequent large economic losses—are significantly greater than one would expect from working with a normal distribution of risks.

This result raises a number of questions—about its robustness and its economic implications. Let's look at the robustness issue first. Several assumptions are made in deriving this outcome. One is a noninformative prior distribution over the variance of the growth rate. Another is that we know the mean of the distribution of growth rates but not the variance. Yet another is that the utility function is unbounded. It is not clear how well the result would survive changes in these assumptions (it would not survive a change in the third). However these are technical points, and it the result would probably survive in some form for other reasonable assumptions. A more fundamental issue is a more abstract one. Weitzman considers a decision problem where the probabilities over outcomes are unknown, and rather than adopting an axiomatic framework similar to but more general than that underlying the classical von Neumann–Morgenstern framework—as used by the Henrys and those on whose work they draw—he assumes we can work within the standard expected utility framework which assumes known probabilities. There is a difficult problem here: what is the right way of analyzing decision problems where probabilities are not known? Should we work within the expected utility framework, even though this has known probabilities built in as fundamental assumptions, take a Bayesian approach and model learning about the unknown parameters, or should we deal directly with the ambiguity associated with unknown probabilities? What has become the standard approach to such problems in the finance literature is to work with multiple priors. This involves recognizing that the information available is insufficient to determine a single probability distribution and that there are many distributions compatible with what we know. These are all possible priors: this multiple-prior-based approach is now widely used in finance as a way modelling risk or ambiguity.

5 Equity and climate change

There are two dimensions of equity that are important in the context of climate change. We have already discussed equity between present and future generations. The other dimension concerns equity between rich and poor countries both now and in the future. This second dimension is invisible in aggregative one-good models, and we have already noted that we need a many-good model to talk seriously about discount rates and climate change. The discussions below will reinforce the need for some measure of disaggregation in the analysis of the economics of climate change.

As noted, the parameter η, the elasticity of the marginal utility of consumption, summarizes our preference for equality: it determines how fast marginal utility falls as income rises. There are two ways in which this affects the case for action on climate change.

As η rises, the marginal utility of consumption falls more rapidly. If consumption is growing over time, then this means that the marginal utility of future generations falls more rapidly with larger values of η and therefore we are less concerned about benefits or costs to future generations. We are less future-oriented—the consumption discount rate ρ is higher—and so place less value on stopping climate change. So via this mechanism, a stronger preference for equality leads to a less aggressive position on the need for action on climate change. Preferences for equality and action on climate change are negatively linked here.

There is another offsetting effect, not visible in an aggregative model. Climate change is an external effect imposed to a significant degree by rich countries on poor countries. The great majority of the greenhouse gases currently in the atmosphere were put there by the rich countries, and the biggest losers will be the poor countries—though the rich will certainly lose as well. Because of this, a stronger preference for equality will make us more concerned to take action to reduce climate change.

So we have an ambiguous impact of a stronger preference for equity on our attitude towards climate change. Via the mechanism captured in the formula for the consumption discount rate, it makes us less future oriented—provided consumption is growing. (If consumption were to fall, it would make us more future oriented, and if consumption of some goods were to rise and that of others to fall, the effect would be a priori unclear.) And via our concern for the poor countries in the world today it makes us more future-oriented.

Unfortunately, without exception analytical models capture only the first of these effects. They are aggregative one-sector models or models with no distributive weights and so their operation does not reflect the second mechanism mentioned above. This explains the really puzzling and counter-intuitive result that a greater preference for equality in Nordhaus's DICE model leads to less concern about climate change.

To capture fully the contradictory impacts of preferences for equality on climate change policy, we need a model that is disaggregated both by consumption goods and by consumers, allowing us to study the consumption of environmental as well as non-environmental goods and also the differential impacts of climate change on rich and poor nations.

6 Costs of climate change

Having reviewed some of the theoretical issues raised by the recent literature on climate change, it is time to look briefly at the estimates that are available of the costs and benefits of action to reduce emissions of greenhouse gases.

6.1 The costs of action

I look first at the costs of action to prevent climate change, as this seems to be the less controversial of the two areas. The latest IPCC report estimates the cost of keeping CO_2 equivalent concentrations below about 450 parts per million (ppm) as less than 3% of world GDP by 2030 and less than 5.5% by 2050. The Stern review estimates the costs of keeping these concentrations at less than 500–550 ppm as within the range −1% to +3%, with a best estimate of 1%. This is a continuing cost. These statements already illustrate a point that concerns me about the analysis of the costs and benefits of climate change, which is this habit of expressing things as a percentage of national income. Clearly there are many reasons for wanting to do this, not the least of which is the desire to give a sense of scale. But there is an implication here that as GDP rises then the costs of stopping climate change will rise in proportion. Is this true? If we double our income, do we double the cost of stopping climate change? This is not obvious to me. The same issue arises with respect to the damages from climate change: these are typically expressed by a multiplicative factor on national income, which again implies that if income doubles then the damages will double. This is once again not a self-evident point. If incomes double, does the value of land lost due to sea level rise double? Does the cost of extinction of species double? We can probably write out models in which this is true, but it is not clear that they will be convincing.

We can get a rough sense of whether these numbers are reasonable from some back-of-the-envelope calculations. Suppose we need to reduce CO_2 emissions by about 30 gigatons annually: currently that would take emissions down close to zero, but as they are growing quite fast it would leave them positive but small for the next few decades. Assume that the average cost per ton of reducing emissions is $40, which is within the range considered by the IPCC and also by several other studies. (One of the most interesting and detailed is a recent study by McKinsey's, giving a detailed supply curve for reductions in carbon emissions (McKinsey & Company 2007)). Then reducing emissions by $30x10^9$ tons at $40 per ton leads to a total cost of 2.6% of world GDP. The cost of reducing emissions could on average be less than $40 per ton, as there are plenty of opportunities to reduce emissions at less than this cost, and this assumes a reduction to close to zero, much more than would be required to stabilize atmospheric concentrations of greenhouse gases. Furthermore, the deployment of greenhouse gas reduction technologies on so large a scale would almost certainly lead to large reductions in costs. So the numbers that Stern and the IPCC quote seem reasonable.

6.2 The costs of inaction

Here there is a lot of disagreement. Most of the integrated assessment models suggest that the costs of climate change would be of the order of 1% or 2% of national income. Stern suggests a much larger number, at least 5% and possibly as much as 20%. These numbers are the annuitized costs of climate change—the annuities with the same present value as the damages from climate change. I am inclined to think that Stern is much nearer the mark: it is impossible to read the IPCC reports and believe that the consequences of climate change along the business as usual (BAU) path are only 1% or 2% of national income. 1% is almost within the margin of accounting error, and the IPCC certainly gives the impression that climate change will have a far-reaching impact on many human activities, which is not consistent with so small a value. Recent work by Schlenker et al. (2006) suggests climate change on the BAU scenario will have a dire impact on U.S. agriculture, reducing the value of output by as much as 70% by the end of the century. William Cline's recent book (2007) also suggests that climate change on the BAU scenario will have a severe harmful impact on agricultural output in many countries, including many developing countries where agriculture is a large fraction of total output.[12] And of course while agricultural output accounts for only a small fraction of GDP in the U.S., if food were to become scarce it is clear that prices would rise to the point where this could change drastically. Our current spending on food greatly understates our willingness-to-pay for food.

The Stern estimate of 5% of GDP is presented as the lower bound for the cost of climate change on the BAU scenario. The Review comments that "*Modeling work undertaken by the Review suggests that the risks and costs of climate change over the next two centuries could be equivalent to an average reduction in global per capita consumption of at least 5%, now and for ever. The estimated damages would be much higher if non-market impacts, the possibility of greater climate sensitivity, and distributional issues were taken into account*." They are leaving out any impact not reflected in market transactions, assuming a rather conservative value for the key climate sensitivity parameter—the parameter S in Eq. 9, about which Weitzman models uncertainty—and not taking in to account the fact that many of the costs of climate change will fall most heavily on the poor. By adopting a conservative value for the sensitivity parameter they are ruling out the extreme outcomes in the tail of the distribution of outcomes that are the focus of Weitzman's study.

It is easy to see the kinds of issues omitted by not considering non-market effects of climate change. The IPCC estimates that of the order of one third of all species could be driven to extinction along a BAU scenario. This would be a radical transformation and impoverishment of our biological environment, with far-reaching implications for the flow of ecosystem services to human societies (discussed in Section 3, where I suggested that there is a minimum level of ecosystem services required for any welfare level) and also with major ethical implications. Do we have the right to condemn to extinction many of the species with which we share the planet? Opinions vary, and interestingly this is an issue on which evangelical Christians are increasingly taking a position. For many people it is one of the most important issues associated with climate change.

The Stern review follows the mainstream of economic analysis in mentioning and then neglecting the distributional impacts of climate change. Yet as Sterner and Persson point out, it is not so long since economists were attaching distributional weights to the costs and benefits in project evaluation. The Stern Review mentions this possibility, but does not develop it at all. The report accepts the principle that the value of an increment of consumption decreases with the recipient's income level in using an elasticity of marginal utility of consumption of unity, so it would be quite consistent with its underlying ethical assumptions to apply the distributional weights implied by this choice. This is a point that Dasgupta (2007) emphasizes.

[12] A recent paper by Guiteras (2007) looks at the impact of climate change on Indian agriculture and predicts significant loss of output.

7 The case for action

Stern and his team argue that there is a strong case for immediate and effective action. Nordhaus, on the other hand, argues for a far more circumspect approach, with a gradual ramp up from small beginnings over several decades. Most IAMs give similar conclusions to Nordhaus. Who is right—or rather, under what assumptions is each side right? And which assumptions are better?

There are five key issues that between them determine whether the analysis suggests a case for strong and immediate action or not. One, of course, is the value assigned to the costs of climate change. (I am not mentioning the cost of preventing climate change as there is general agreement on this.) The second, inevitably, is the choice of a pure rate of time preference. The third is the choice of an elasticity of the marginal utility of consumption. Fourthly there is the issue of whether we break out and model explicitly the consumption of ecosystem services yielded by our stock of natural capital and the impact that climate change will have on this. And finally there is the issue of uncertainty and the possibility of really severe impacts associated with a climate path that the IPCC regards as possible but unlikely.

Different selections on these five issues get different conclusions about the need for action. Stern, for example, chooses $\delta = 0$ and $\eta = 1$ and estimates the damages from climate change to be high. This combination of assumptions, together with historically normal rates of consumption growth $R(c)$ and a single aggregate consumption good, justifies strong and immediate action. Nordhaus sets $\delta = 3$, $\eta = 1$ and makes conservative assumptions about damages, concluding that there is no case for strong immediate action. If we replace Stern's $\eta = 1$ by Nordhaus' $\eta = 3$, Stern's numbers no longer justify immediate action, even with $\delta = 0$. This reflects the fact that in an aggregative model, a stronger preference for equality reduces our concern from the (richer) future and therefore our interest in preventing climate change. It is hard to know what η should be: Weitzman argues for $\eta = 2$, Heal and Kristrom suggest that the index of relative risk aversion (which is η) is between 2 and 6. Dasgupta comments on $\eta = 1$ as follows: "$\eta = 1$ is to insist that any proportionate increase in someone's consumption is of equal social worth to that same proportional increase in the consumption of any other contemporary no matter how rich or poor. With $\delta = 0$ it implies that any proportionate increase in consumption today is of the same social worth as the same proportional increase at any other date no matter how rich or poor the people then." Another perspective is that $\eta = 1$ implies that taking \$1 from a person earning \$1,000 can be compensated by giving \$1,000,000 to Bill Gates.[13] Dasgupta and Weitzman argue for a higher value of η, which would imply an even more striking trade-off: you would have to give even more that \$1,000,000 to Bill Gates to compensate for taking \$1 from the person earning \$1,000. Weitzman suggests that the risk of an outlying outcome is sufficient to justify strong action—he thinks that Stern is right in his recommendations but for the wrong reason. Sterner and Persson argue, as have Heal and Guesnerie, that disaggregating consumption and modelling the flow of services from ecosystems, which are likely to be seriously damaged by changing climates, will justify action on climate change even with the discount rates and elasticities favored by Nordhaus.

[13] On the assumption that his income is \$1 billion.

The following table summarizes these conclusions: it shows the combinations of assumptions used by each author and the conclusions reached with respect to the need for strong and immediate action. ESS denote ecosystem services—whether the author models these explicitly. Risk likewise indicates whether the author takes the risk of extreme outcomes into account. I have inserted a question mark under costs for the work of Sterner and Persson (S&P) because they work with Nordhaus's DICE model, which has low estimates of the costs of climate change, but by modeling climate damage to natural capital and the ecosystem services flowing from it, they are inserting a new element into the costs of climate change, an element not present in the original DICE model.

Author	Stern	Nordhaus	S&P	Weitzman
d	0	3	3	2
h	1	1	1	2
ESS	No	No	Yes	No
Risk	No	No	No	Yes
Costs	High	Low	?	High
Strong action	Yes	No	Yes	Yes

We can see from this that there are several ways of concluding that we need to take action. We can follow the route of the Stern review and use a low discount rate and set $\eta = 1$, or we can allow for climate impacts on ecosystem services, or we can be explicitly concerned about the risk of an outcome in the tail of the distribution of possible outcomes. Any of these seems sufficient to justify immediate action. And several of them seem plausible. My own personal judgment would be to set $\delta = 0$, to disaggregate consumption and model the effect of climate on ecosystem services, and to worry about the risk out extreme outcomes, a blend of Stern, Sterner and Persson and Weitzman. That would certainly secure a case for very strong and immediate action, more so than the Stern Review. But if you disagree with me on the choice of a pure rate of time preference, but agree with me on the other matters, or even on one of them, you would still have to agree that there is a case for strong action now. The table misses out one framework that it would be interesting to explore, because no one has looked it them till now. This is a model disaggregated by both consumers and consumption goods with an explicit use of distributional weights not only over time, as is customary in utilitarian models, but also across countries. My guess is that this would increase the case for action on climate change, but in the absence of a model incorporating this framework this remains a guess.

8 Conclusions

To return to the questions with which I opened this article: what have we learned from the outpouring of literature as a result of the Stern Review? A lot. We have explored the model space and the parameter space much more thoroughly, though there are still unexplored regions. I think this should change the presumption that economists hold about the need for strong action on climate change, which prior to Stern was largely negative, to positive. We can see many ways in which we can make a case for strong action now, and few in which we can deny it. While there are aspects of the Stern Review's analysis with which we can disagree, it seems fair to say that it has catalyzed a fundamental rethinking of the economic case for action on climate change. Recent developments both allow us to see more clearly conditions under which there is a case for acting quickly and strongly on climate change, and the conditions under which such action is not justified. There are several combinations of assumptions that justify strong action, depending on choices of the pure rate of time preference, the elasticity of marginal utility, the costs of climate change and the nature of uncertainty and the way in which we react to this. The analysis also reveals that in spite of the extensive literature, there are issues that remain to be explored, many of which are related to the fact that most modeling to date has been in the context of one-good one-country models. We have really not spent enough time on the impact of climate change on our natural capital and the ways in which it may compromise the flow of essential ecosystem services from this. Nor do we have much by the way of modelling of our preferences for such services and of the degree of substitutability between produced goods and services and ecosystem services. These parameters affect consumption discount rates and their relationships with market data such as interest rates. Understanding these issues will require models that are more disaggregated than those that have been used to date, as will representing more satisfactorily the issues raised by a concern for equity at the international level. It is very clear that most of the models analyzed to date are so aggregated as to miss many important issues.

Acknowledgements I am grateful to Scott Barrett, Tony Fisher, Claude Henry, Bengt Kristrom, Karl-Goran Maler, Bill Nordhaus, Bob Pindyck and Nick Stern for valuable comments on earlier versions of this paper. This does not imply that they agree with it! Errors are of course my own responsibility. A shortened and non-technical version of this paper appeared as "Climate Economics—A Meta-Review and Some Suggestions" in the Review of Environmental Economics and Policy. I am grateful to three referees for some very constructive comments.

References

Barbier E, Heal G (2006) Valuing ecosystem services. The Economists' Voice, Berkeley Press, January 2006

Barrett S (1994) Economic growth and environmental preservation. J Environ Econ Manage 23:289–300

Berry S, Heal G, Salamon P (1978) On a relationship between economic and thermodynamic optima. Resources & Energy 1:125–137

Cline W (2007) Global warming and agriculture: impact estimates by country. The Peterson Institute, July 2007

Dasgupta P (2007) Discounting climate change. September 2007, memo, Department of Economics, University of Cambridge

Dasgupta P, Heal G (1979) Economic theory and exhaustible resources. Cambridge University Press

Fisher A, Krutilla J (1975) Resource conservation, environmental preservation and the rate of discount. Q J Econ 89(3):358–370

Fisher A, Narain U (2002) Global warming, endogenous risk and irreversibility. Working Paper, Department of Agricultural and Resource Economics, U.C. Berkeley

Foley D (2007) The economic fundamentals of global warming, memo, The New School, October 2007. Presented at the Workshop on the Economic Analysis of Global Warming at the New School on October 12 2007

Gerlagh R, van der Zwaan R (2002) Long-term substitutability between the environment and man-made goods. J Environ Econ Manage 44:329–345

Ghiradato P, Marinacci M (2001) Risk, ambiguity and the separation of utility and beliefs. ICER Working Paper 21

Guesnerie R (2004) Calcul économique et development durable. La Revue Economique

Guiteras R (2007) The impact of climate change on Indian Agriculture. Working paper. Department of Economics, MIT
Harrod SR (1948) Towards a dynamic economics. Macmillan, London
Heal G (1998) Valuing the future: economic theory and sustainability. Columbia University Press
Heal G (2005) Intertemporal welfare economics and the environment. Handbook of Environmental Economics, Volume 3.Edited by K-G Mäler and J.R. Vincent, Elsevier, Chapter 21, 1105–1145
Heal G, Kriström B (2002) Uncertainty and climate change. Environ Resour Econ 22:3–39
Henry C (2007) Somewhat different ways, but the same destination. On some controversial points in the Stern report. Working Paper, Laboratoire d'Économetrie de l'École Polytechnique
Henry C, Henry M (2002) Formalization and applications of the precautionary principle. Working Paper, Laboratoire d'Économetrie de l'École Polytechnique
Intergovernmental Panel on Climate Change (2007) Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Climate Change 2007: The Physical Science Basis. Available at www.ipcc.ch
Kant I (1950) Prolegomena to any future metaphysics, Carus translation, Bobbs-Merrill
Krautkramer J (1985) Optimal growth, resource amenities and the preservation of natural environments. Rev Econ Stud 52:153–170
Malinvaud E (1953) Capital accumulation and the efficient allocation of resources. Econometrica 21(2):233–268
McKinsey & Company. (2007) Reducing U.S. Greenhouse Gase Emissions: How Much at What Cost?
Nordhaus W (1993) Rolling the DICE: an optimal transition for controlling the emission of greenhouse gases. Resourc Energy Econ 15:27–50
Nordhaus W (2007) The stern review of the economics of climate change. J Econ Lit 45(3):686–702
Ramsey F (1928) A mathematical theory of saving. Econ J 38:543–559
Roughgarden T, Schneider S (1999) Climate change policy: quantifying uncertainties for damages and optimal carbon taxes. Energy Policy 27:415–429
Schlenker W, Fisher AC, Hanemann WM (2006) The impact of global warming on U.S. agriculture: an econometric analysis of optimal growing conditions. Rev Econ Stat 88(1):113–125
Sen A (1973) On economic inequality. Clarendon, Oxford
Sidgwick H (1890) The methods of ethics. Macmillan, London
Stern N (2006) The economics of climate change: the Stern Review. H.M. Treasury, UK
Sterner T, Persson M (2007) An even sterner review: introducing relative prices into the discounting debate. Discussion Paper, Resources for the Future, July 2007, RFF DP 07-37. Available at http://www.rff.org/rff/Documents/RFF-DP-07-37.pdf
Weitzman M (1998) Why the far distant future should be discounted at the lowest rate. J Environ Econ Manage 36:201–208
Weitzman M (2007a) The stern review of the economics of climate change. J Econ Lit 45(2):703–724
Weitzman M (2007b) Structural uncertainty and the value of statistical life in the economics of catastrophic climate change. Working Paper, July 2007
World Bank (2006) Where is the Wealth of Nations? Measuring capital in the 21st century. The World Bank, Washington