

Climate Economics: A Meta-Review and Some Suggestions for Future Research

Geoffrey Heal*

Introduction

About two years ago, the U.K. government published *The Economics of Climate Change: The Stern Review*, written by a team led by Nicholas Stern (Stern 2006). The publication of the Stern Review provoked an unprecedented outpouring of articles on the same topic: we have probably seen more economics articles on climate change in the last year and a half than in all preceding history, including an entire issue of *The Economists' Voice*¹ and large collections in this journal and in *World Economics*. Many of these articles have been written by very distinguished colleagues. So the Stern Review has triggered 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, 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 obvious 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.

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 in this article. I suggest that the recent debate has clarified many important issues, and that we are now in a position to identify those conditions that are sufficient to make a case for strong action on climate change. However, more work is needed before we can have a fully satisfactory account of the relevant economics. In particular, we need to better understand how climate change affects natural capital—the natural environment and the ecosystems comprising it—and how this in turn affects human welfare. This article takes some first steps in this direction.

*Columbia Business School, New York, NY, 10027; E-mail: gmh1@columbia.edu.

I am grateful to Scott Barrett, Tony Fisher, Claude Henry, Michael Hoel, Bengt Kriström, Karl-Goran Maler, Bill Nordhaus, Bob Pindyck, Nick Stern, and a referee for valuable comments on earlier versions of this article. This does not imply that they agree with it! Errors are of course my own responsibility.

¹Arrow 2007; Barrett 2007; Gans 2007; Goulder 2007; Olmsted and Stavins 2007; Schelling 2007; and Stiglitz 2007.

Review of Environmental Economics and Policy, volume 3, issue 1, winter 2009, pp. 4–21

doi:10.1093/reep/ren014

Advance Access publication on September 24, 2008

© The Author 2008. Published by Oxford University Press on behalf of the Association of Environmental and Resource Economists. All rights reserved. For permissions, please email: journals.permissions@oxfordjournals.org

Welfare Economics and Climate Change

Let me begin with the basic economic theory of climate change. I will focus on the discount rate, but first I want to discuss a simple, important, and interesting point that has recently been emphasized by Foley (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 be possible in principle to make a Pareto-improving change to the world economy. If we do this then there is in the aggregate no net cost to correcting climate change: that is, 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 actually do is a matter of institutional design. This is something 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.

Discount Rates and the Environment

Now to discount rates. As anyone who has spent even a short time studying climate change must be aware, one of the most controversial issues is the choice of a discount rate. By this we mean the pure rate of time preference (PRTP), as distinguished clearly from the consumption discount rate (CDR). 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 in general 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.² 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, for example, one valued the utility of people in Asia differently from the utility of people in Africa, except that we are using differences in date rather than place as the basis for differentiation. In contrast, the CDR ρ is the rate of change of the present value of the marginal utility of consumption.

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 (2006) uses this argument to justify a very low PRTP. Civilization might, for example, be destroyed by a meteor hitting the earth. My own judgment is that over the period we are considering—say the next two centuries—the chance of this is small enough to be neglected. However, the point seems to be valid conceptually.

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

²We could also value them differently for all manner of other reasons—differences in nationality, ethnicity, and proximity either physically or genetically. In general we do not do these things, at least explicitly, which to me makes it strange that we do explicitly discriminate by proximity in time.

taxation (see Sen 1973), and is very widely 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 $\eta(c_t)$, the elasticity of the marginal utility of consumption. Pulling 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 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 others would continue to rise—recognizing this requires that we treat consumption as a vector of different goods that can be affected differently by climate change.³

Let us follow this line of thought and disaggregate consumption at date t into a vector $c_t = (c_{1,t}, c_{2,t}, \dots, c_{n,t})$ of n different goods. Utility is increasing at a diminishing rate in all of these and is a concave function overall. In this case, there is a CDR for each type of consumption and we have n equations, 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 $\eta_{ij}(c_t)$ of good i with respect to good j times the growth rate of consumption of good j .⁴ Here the own elasticities such as $\eta_{ii}(c_t)$ are positive numbers, but the cross-elasticities $\eta_{ij}(c_t)$, $j \neq i$, are zero if the utility function is additively separable and can otherwise have either sign.

Let us test our intuition and take the case where an environmental good and produced consumption are highly complementary, so that indifference curves are near to right angled and the elasticity of substitution 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 CDR on the regular good. Conversely, if the availability of the environmental good is falling then this raises the CDR 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 CDR, 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 CDR negative.

³For an early recognition of this point see Fisher and Krutilla (1975), who note that increasing scarcity of in situ resources of the natural environment will drive up our valuation of them relative to some extractive resources and produced goods generally.

⁴That is,

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

where ρ_{it} is the consumption discount rate on good i at date t , $R(c_{i,t})$ is the rate of change of consumption of good i at date t , and $\eta_{ij}(c_t)$ 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 paper (1953).




Figure 1. Preferences for produced goods and ecosystem services.

Whether produced goods and environmental services are substitutes or complements in consumption is not an issue that has been discussed in the literature, as we almost always work with one-good models. There do however seem to be reasons to suppose that complementarity is the better assumption. Dasgupta and Heal (1979), following Berry, Heal, and Salamon (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 that if the utility function has constant elasticity of substitution (CES) then the elasticity is less than one.

Figure 1 illustrates this idea: it shows indifference curves for a two-argument utility function, consumption of produced goods and ecosystem services. 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.

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. In the case of a one-good model with

just aggregate consumption in the utility function, we can easily show that under certain assumptions the CDR equals the rate of return on capital. What assumptions are needed to reach 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 return on capital becomes more complex if there are several goods in the utility function. In the case of just two goods, the CDR on the consumption good 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 CDRs 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 CDR upwards.⁵ In particular, Nordhaus's (Nordhaus 1993; Nordhaus and Boyer 2000) estimates of the exogenous rate of growth of total factor productivity in his DICE (dynamic integrated climate and economy) model will be too high.

Uncertain Discount Rates

Weitzman (1998) has argued 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 do not know which is the right one, then in the long run we should work with the lowest of all the possible rates. My own judgment is that the right PRTP is zero. 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 Weitzman's result is undoubtedly technically correct, I am not totally certain of its philosophical foundations and implications. If we disagree over the discount rate, 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 article and they are not obvious to me. It seems to me that in this case 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 you a nonconstant discount rate, one that changes over time and converges asymptotically on the lowest rate suggested. Any intertemporal plan constructed using such a discount rate will be dynamically inconsistent, in the sense that if we follow it for a period of time and then stop and ask what is the best continuation from where we are, it will not be the plan that we originally adopted. I am not sure if this matters, but it certainly makes me pause and think.

⁵Personal communication, Karl-Goran Mäler, 2008.

Choosing a Discount Rate

Having set out some of the relevant theory, let us return to the issue of the choice of a discount rate, and see whether there are any noncontroversial 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 δ and the CDR ρ . 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 certainly depends on δ , which is exogenous, but also on $R(c)$, the consumption growth rate, which is clearly endogenous and is part of 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.⁶ 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 than to a poor person? The answer to this question is η , which is clearly normative.

The bottom line, then, is that we cannot select either discount rate without making ethical judgments. Ramsey (1928) famously commented that discounting future utilities is ethically indefensible and arises purely from a weakness of the imagination, Harrod (1948) spoke in similar terms, and Sidgwick (1890), the philosopher, commented: “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” (412). This is an ethical judgment not a theorem, so you do not 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 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 think there is often a lot of confusion between δ and ρ : 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.

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 confusion between 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, and indeed generally will be, even when the PRTP is zero.

Operationally when do we use δ and when do we use ρ ? To pick a CDR you need a consumption growth rate. So you can not 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

⁶These statements are obviously true whether c is a scalar or a vector.

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 a partial equilibrium your analysis will make only marginal changes about the preexisting 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 CDR as part of the solution. So the DICE model developed by Nordhaus (1993), probably the most widely used IAM, needs to be given a PRTP 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.

Let us return to the relationship between the CDR and the return on capital. As we have seen above, for the aggregate one-good model this is equal to the return on capital under certain assumptions. These assumptions are very strong indeed, and it is not clear to me that they are relevant in the context of climate change. One assumption that is necessary 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 CDR for a model of climate change? An assumption of perfect foresight in all capital markets indefinitely far into the future is also required. 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.⁷ There are lots of questionable things here. One is that, as noted we need very strong assumptions, assumptions that seem to be particularly out of place in a discussion of climate change. There is also the issue of deducing an ought from an is, recognized as a philosophically dubious step since the time of Hume. But leaving these aside, you can argue that it is using the relationship between the CDR and the return on investment (ROI) in the 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 do not 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 the equation for the CDR and the ROI to make statements about this PRTP, and another reason why its relevance is limited. 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

⁷They take the equation relating the CDR and the return on investment (ROI), 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 $\delta + \eta(c_i) R(c_i) = f'(k)$.

example they will behave differently in response to climate change. To develop this point we need to talk about natural capital and ecosystem services. As the World Bank (2006) recently emphasized, natural capital is an important part of the wealth of nations. Ecosystem services are some of the returns that come from that capital, and are of great 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 (IPCC 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 it is important to distinguish at least two types of consumption: conventional consumption c and 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 Krautkramer (1985), was developed by Barrett (1994), and Heal (1998), and has been emphasized recently by Guesnerie (2004) and Sterner and Persson (2008). On the assumption that s is falling over time because of climate change, we see that the CDR on the consumption good is

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

Here $R(s) < 0$ but we do not 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. By arguments given earlier, it seems likely to be negative. This means that even in a 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.

A “Sterner” Perspective

It is worth looking in more detail at the Sterner and Persson (2008) 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, 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. Consequently, the present value of an increment of environmental services may be rising over time, and the CDR on environmental services may thus be negative, precisely the point that we were making above. This could be the case even with a positive PRTP, which is the main point of the Sterner and Persson article. They also present an interesting modification of Nordhaus’s DICE model to incorporate this point. They replace the standard utility function with a CES function that is modified to reflect a constant relative risk aversion and is a function of two variables, consumption of the produced good and consumption of ecosystem services. As suggested above, they assume that the supply of environmental services is negatively affected by temperature according to the square of temperature, and that the share of environmental goods in consumption is about 20 percent. They use these assumptions to calibrate the modified DICE model and then run the model with the discount rate used by Nordhaus. Their runs show that even with such a high discount rate, the presence of an environmental stock that is damaged by higher temperatures radically transforms the optimal emissions path of CO₂. This leads to a vastly more conservative policy toward climate change, with emissions both

staying lower and falling faster. In fact it leads to an even more aggressive reduction in greenhouse gases than what is recommended by the Stern Review.

The Role of Uncertainty

Uncertainty is one of the dominant facts in any analysis of climate change. It is not that the underlying science is uncertain: the mechanisms through which the accumulation of greenhouse gases in the atmosphere warm the earth are simple and well understood. The point is that there is some debate about whether this will lead to a temperature rise of 2 degrees or 6 degrees Celsius, 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, as noted by 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 do not 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 probability distribution. In the terminology of Frank Knight (1921), 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.

Another point about uncertainty that has been made by a number of authors is that climate change and many of its consequences are irreversible.⁸ Melting of ice sheets and glaciers, extinction of species, and destruction of coral 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.

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. Heal and Kriström (2002) investigate this, and present some simple yet suggestive calculations indicating the importance of risk aversion. These calculations were made in response to the question: what cost is it worth incurring to avoid the risk of climate change? Within a simple framework, Heal and Kriström (2002) carry out calculations that illustrate the issues involved, and how discount rates, risk aversion, and probabilities interact. They show that the amount it is worth paying to avoid the risk of climate change is very sensitive to

⁸See Fisher and Narain (2002), and for a survey see Heal and Kriström (2002).

the PRTP and the index of relative risk aversion, which is also the elasticity of marginal utility. In some respects, the amount it is worth paying is more sensitive to changes in these two economic parameters than to variation in the scientific forecasts concerning climate change.

Unknown Probabilities

Some interesting articles have recently tried to tackle directly the point that we do not know the probability distribution governing the consequences of climate change. Henry and Henry (2002) have looked at this issue from the perspective of ambiguity theory, the theory of choice in the face of incomplete probability distributions (see also Henry 2007), and Weitzman (2007) has looked at it from the perspective of learning about unknown parameters, including those of the probability distribution. In each case the assumption is that we do not know the distribution of damages from climate change. In the Henry article, the lack of knowledge is of a very general kind, with some information about the relative likelihoods of different outcomes but not enough to form a probability distribution. Ghiradato and Marinacci (2001) analyzed decision problems in this framework in an axiomatic framework similar to the one introduced by von Neumann and Morgenstern (1944) for expected utility theory. The central result of the Henry article is that 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 (2007b) assumes that we know the type of distribution that governs the damages from climate change but not its parameters, and that we are also uncertain about one of the key parameters driving climate change, the climate sensitivity parameter relating the change in atmospheric greenhouse gas concentrations to the change in global mean temperature. He models the process of learning about these parameters. In his analysis of uncertainty about the climate sensitivity parameter, he finds that we are willing to give up an unlimited amount of consumption today to assure one unit in the future. In other words, under Weitzman's assumptions, we are extraordinarily risk-averse about the future.

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 the growth rate is distributed normally. However, he adds that the variance of the normal distribution is unknown and can assume any non-negative value, which means that the result is a distribution of consumption, which has some of the properties of power law distributions. That is, it has a fat tail, meaning that there is much more weight in the tails of the distribution than there is in the normal distribution 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.

Weitzman considers a decision problem where the probabilities over outcomes are unknown. Rather than adopting an axiomatic framework similar to the one underlying the classical von Neumann–Morgenstern framework, as did Henry and Henry, 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

or subjective approach and model learning about the unknown parameters, or deal directly with the ambiguity associated with unknown probabilities? To take the Bayesian approach appears to contradict one of the Henrys' results on the suboptimality of using the subjective expected utility framework in the context of ambiguity. This is a fundamental question to which we do not know the answer.

Let us look at the implications of this result. It is saying that if we think that the risks we face from climate change are normally distributed but we do not know the variance of the distribution and have no prior information about the distribution, then we are actually facing risks governed by a much more threatening distribution than the normal, one that has far more weight out in the tails and therefore a far greater chance than we might have thought of a really negative outcome. Overall, I find this result to be worrying and thought-provoking, though as I indicated above I am not sure what is the best approach to decision-making with unknown risks. We need a lot more thinking about this issue.

Equity and Climate Change

There are two dimensions of equity that are important in the context of climate change. One 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 indeed 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.

We have already spent some time on equity between present and future and on the role of the elasticity of the marginal utility of consumption. As noted, this indicates, among other things, our preference for equality as it determines how fast marginal utility falls as income rises. There are two ways in which our preference for equality affects the case for action on climate change.

The first has already been discussed: as the elasticity rises, the marginal utility of consumption falls more rapidly. If consumption is growing over time, this means that the marginal utility of future generations falls more rapidly with larger values and therefore we are less concerned about benefits or costs to future generations. We are less future-oriented—the CDR is higher—and so we 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. In other words, preferences for equality and action on climate change are negatively linked here.

There is a second, offsetting effect, which as noted above is 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 about taking action to reduce climate change.

So the impact of a stronger preference for equity on our attitude toward climate change is ambiguous. Via the mechanism captured in the formula for the CDR, it makes us less future

oriented—provided consumption is growing.⁹ 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 counterintuitive result in Nordhaus's DICE model that a greater preference for equality 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 nonenvironmental goods and also the differential impacts of climate change on rich and poor nations.

Costs of Climate Change

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

The Costs of Action

I look first at the costs of action to prevent climate change, as this seems to be the less controversial issue. The latest IPCC report (2007) estimates the cost of keeping CO₂ equivalent concentrations below about 450 parts per million (ppm) as less than 3 percent of world GDP by 2030 and less than 5.5 percent by 2050. The Stern Review (2006) estimates the costs of keeping these concentrations at less than 500–550 ppm as being within the range –1 percent to +3 percent of world GDP, with a best estimate of 1 percent. 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, 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? The answer 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. Once again this is 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 CO₂ 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

⁹If 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.

average cost per ton of reducing emissions is \$40, which is within the range considered by the IPCC and also by several other studies.¹⁰ Then reducing emissions by 30 gigatons at \$40 per ton leads to a total cost of 2.6 percent 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. Moreover, this assumes emissions are reduced 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.

The Costs of Inaction

Here there is a lot of disagreement. Most of the IAMs suggest that the costs of climate change would be on the order of 1 or 2 percent of national income. Stern suggests a much larger number, at least 5 percent and possibly as much as 20 percent. 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 percent of national income. 1 percent 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, Hanemann, and Fisher (2006) suggests that climate change under the BAU scenario will have a dire impact on US agriculture, reducing the value of output by as much as 70 percent by the end of the century. Cline's recent book (2007) also suggests that climate change under 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.¹¹ And of course while agricultural output accounts for only a small fraction of GDP in the United States, 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 Review presents 5 percent of GDP as the lower bound for the cost of climate change under the BAU scenario, noting 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 percent, 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." The Review leaves out any impact not reflected in market transactions, assumes a rather conservative value for the key climate sensitivity parameter, and does not take into 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 that are the focus of Weitzman's study.

¹⁰One of the most interesting and detailed is a recent study by McKinsey (2007), which gives a detailed supply curve for reductions in carbon emissions.

¹¹A recent paper by Guiteras (2007) looks at the impact of climate change on Indian agriculture and predicts significant loss of output.

It is easy to see the kinds of issues omitted by not considering non-market effects of climate change. The IPCC estimates that about 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 as well as 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 ignoring the distributional impacts of climate change. Yet as Sterner and Persson (2008) point out, it is not so long since economists were attaching distributional weights to the costs and benefits in project evaluation. In using an elasticity of marginal utility of consumption of unity, the Stern Review accepts the principle that the value of an increment of consumption decreases with the recipient's income level, 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.

The Case for Action

Stern and his team argue that there is a strong case for immediate and effective action. Nordhaus (2007), on the other hand, argues for a far more circumspect approach, with a gradual ramp up from small beginnings over several decades. Most IAMs offer conclusions similar to those of Nordhaus. Who is right—or rather, under what assumptions is each side right? And which assumptions are better?

There are five key issues that together determine whether or not the analysis suggests a case for strong and immediate action. The first, 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 PRTP. The third is the choice of an elasticity of the marginal utility of consumption. Fourth, 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 choices concerning these five issues result in different conclusions about the need for action. Stern, for example, chooses PRTP $\delta = 0$ and elasticity $\eta = 1$ 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 for 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 Kriström (2002) cite the empirical literature on financial markets, which

suggests that the index of relative risk aversion (which is η) is between 2 and 6. Dasgupta (2007) 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 interpretation is that $\eta = 1$ implies that taking \$1 from a person earning \$1000 can be offset by giving \$1,000,000 to Bill Gates.¹² 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 than \$1,000,000 to Bill Gates to compensate for taking \$1 from the person earning \$1000. Weitzman suggests that the risk of an outlier outcome is sufficient to justify strong action—he thinks that Stern is right in his recommendations but for the wrong reason. Sterner and Persson (2008) argue, as have Heal (2005) and Guesnerie (2004), that disaggregating consumption and modeling 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.

To summarize, it appears that there are several ways to justify strong and immediate action on climate change. 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 conditions seems sufficient to justify immediate action, and several of them seem plausible. My own personal judgment would be to use a blend of Stern, Sterner and Persson, and Weitzman and set $\delta = 0$, disaggregate consumption, and model the effect of climate on ecosystem services, and worry about the risk of extreme outcomes. That would certainly secure a case for very strong and immediate action, even more so than the Stern Review. If you disagree with me on the choice of a PRTP, 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. One framework that no one has looked at until now but would be interesting to explore 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 intuition is that this would strengthen the case for action on climate change, but in the absence of a model incorporating such a framework, this remains only a guess.

Conclusions

Returning to the questions raised at the beginning of 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 from largely negative (prior to Stern) to positive. We can see many ways for making a case for strong action now, and few for denying it. While

¹²On the assumption that his income is \$1 billion.

there are aspects of the Stern Review's analysis with which we may disagree, it seems fair to say that it has catalyzed a fundamental rethinking of the economic case for action on climate change. Recent developments allow us to see more clearly both the 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 PRTP, the elasticity of marginal utility, the costs of climate change, the nature of uncertainty, and the way in which we react to this.

The analysis presented here also reveals that despite 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 this may compromise the flow of essential ecosystem services. Nor do we have much in the way of modeling our preferences for such services and the degree of substitutability between produced goods and services and ecosystem services. These parameters affect CDRs and their relationships with market data such as interest rates. Unfortunately, most of the models analyzed have been so aggregated as to miss these important issues. Understanding these issues, as well as representing more satisfactorily the issues raised by a concern for equity at the international level, will require models that are more disaggregated than those that have been used to date.

References

- Arrow, Kenneth J. 2007. Global climate change: A challenge to policy. *Economists' Voice* (July 2006). www.bepress.com/ev.
- Barbier, Edward, and Geoffrey Heal. 2006. Valuing ecosystem services. *Economists' Voice* (January 2006).
- Barrett, Scott. 1994. Economic growth and environmental preservation. *Journal of Environmental Economics and Management* 23:289–300.
- Barrett, Scott. 2007. Proposal for a new climate change treaty system. *Economists' Voice* (July 2006). www.bepress.com/ev.
- Berry, Stephen, Geoffrey Heal, and Peter Salamon. 1978. On a relationship between economic and thermodynamic optima. *Resources & Energy* 1:125–37.
- Cline, William. 2007. *Global Warming and Agriculture: Impact Estimates by Country*. Washington, DC: Peterson Institute for International Economics.
- Dasgupta, Partha. 2007. Discounting climate change. Memo, Department of Economics, University of Cambridge, September 2007.
- Dasgupta, Partha, and Geoffrey Heal. 1979. *Economic Theory and Exhaustible Resources*. Cambridge, UK: Cambridge University Press.
- Fisher, Anthony, and John Krutilla. 1975. Resource conservation, environmental preservation and the rate of discount. *Quarterly Journal of Economics* 89(3):358–70.
- Fisher, Anthony, and Urwishay Narain. 2002. Global warming, endogenous risk and irreversibility. Working Paper, Department of Agricultural and Resource Economics, University of California at Berkeley.
- Foley, Duncan. 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.
- Gans, Joshua. 2007. Do voluntary carbon offsets work? *Economists' Voice* (July 2006). www.bepress.com/ev.
- Ghiradato, P., and M. Marinacci. 2001. Risk, ambiguity and the separation of utility and beliefs. ICER Working Paper 21.

- Goulder, Lawrence. 2007. California's bold new climate policy. *Economists' Voice* (July 2006). www.bepress.com/ev.
- Guesnerie, Roger. 2004. Calcul économique et développement durable. *La Revue Economique*.
- Guiteras, Raymond. 2007. The impact of climate change on Indian agriculture. Working Paper, Department of Economics, MIT.
- Harrod, Sir Roy. 1948. *Towards a Dynamic Economics*. London: Macmillan.
- Heal, Geoffrey. 1998. *Valuing the Future: Economic Theory and Sustainability*. New York: Columbia University Press.
- Heal, Geoffrey. 2005. Intertemporal welfare economics and the environment. Vol. 3, *Handbook of Environmental Economics*, ed. K.-G. Mäler and J. R. Vincent, chap. 21, 1105–45. Oxford, UK: Elsevier.
- Heal, Geoffrey, and Bengt Kriström. 2002. Uncertainty and climate change. *Environmental and Resource Economics* 22:3–39.
- Henry, Claude. 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, Claude, and Marc Henry. 2002. Formalization and applications of the precautionary principle. Working Paper, Laboratoire d'Économetrie de l'École Polytechnique.
- Intergovernmental Panel on Climate Change (IPCC). 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. www.ipcc.ch.
- Knight, Frank H. 1921. *Risk, Uncertainty, and Profit*. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin.
- Krautkramer, Jeffrey. 1985. Optimal growth, resource amenities and the preservation of natural environments. *Review of Economic Studies* 52:153–70.
- Malinvaud, Edmond. 1953. Capital accumulation and the efficient allocation of resources. *Econometrica* 21(2):233–68.
- McKinsey & Company. December 2007. Reducing U.S. greenhouse gas emissions: How much at what cost? McKinsey & Company. http://www.mckinsey.com/client/service/ccsi/pdf/US_ghg_final_report.pdf.
- Nordhaus, William. 1993. Rolling the DICE: An optimal transition for controlling the emission of greenhouse gases. *Resource and Energy Economics* 15:27–50.
- . 2007. The Stern Review of the economics of climate change. *Journal of Economic Literature* 45(3):686–702.
- Nordhaus, William, and William Boyer. 2000. *Roll the DICE Again: Economic Models of Global Warming*. Cambridge, MA: MIT Press.
- Olmsted, Sheila, and Robert Stavins. 2007. A meaningful second commitment period for the Kyoto protocol. *Economists' Voice* (July 2006). www.bepress.com/ev.
- Ramsey, Frank. 1928. A mathematical theory of saving. *Economic Journal* 38:543–59.
- Schelling, Thomas. 2007. Climate change: The uncertainties, the certainties and what they imply about action. *Economists' Voice* (July 2006). www.bepress.com/ev.
- Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. 2006. The impact of global warming on U.S. agriculture: An econometric analysis of optimal growing conditions. *Review of Economics and Statistics* 88(1):113–25.
- Sen, Amartya. 1973. *On Economic Inequality*. Oxford: Clarendon Press.
- Sidgwick, Henry. 1890. *The Methods of Ethics*. London: Macmillan.
- Stern, Nicholas. 2006. *The Economics of Climate Change: The Stern Review*. London: H. M. Treasury.
- Stern, Thomas, and Martin Persson. 2008. An even Stern review: Introducing relative prices into the discounting debate. *Review of Environmental Economics and Policy* 2:77–93.
- Stiglitz, Joseph. 2007. A new agenda for global warming. *Economists' Voice* (July 2006). www.bepress.com/ev.
- Von Neumann, John, and Oskar Morgenstern. 1944. *The Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Weitzman, Martin. 1998. Why the far distant future should be discounted at the lowest rate. *Journal of Environmental Economics and Management* 36:201–8.

Weitzman, Martin. 2007a. The Stern Review of the economics of climate change. *Journal of Economic Literature* 45(3):703–24.

———. 2007b. Structural uncertainty and the value of statistical life in the economics of catastrophic

climate change. Department of Economics, Harvard University. Working Paper, July 2007.

World Bank. 2006. *Where is the Wealth of Nations? Measuring Capital in the 21st Century*. Washington, DC: The World Bank.