Modeling Climate Policies:
A Critical Look at Integrated Assessment Models

1. Introduction

This paper addresses a topic at the intersection of the philosophy of science and technology, on the one hand, and practical ethics, on the other. It investigates the link between the predictions of climate models concerning the effects of anthropogenic climate change and policy recommendations concerning appropriate abatement measures. Climate change presents us with a problem of intergenerational justice. While any costs associated with climate change mitigation measures will have to be borne by the world's present generation, the main beneficiaries of mitigation measures will be future generations. This raises the question whether present generations have a responsibility to shoulder these costs: Do we have a duty toward future generations and if yes, to what extent would any such duty require us to reduce our carbon emissions today?

One approach for addressing this question, influential in particular in philosophy and public policy discussions, is to appeal to a neo-classical economic cost-benefit analysis, which weighs the present (and future) costs associated with mitigation measures against their future benefits, to determine what course of action a principle of intergenerational welfare maximization would require of us. According to this approach we have the obligation to implement mitigation measures to the extent that their marginal costs (to us) do not exceed their marginal benefits (to future generations).

While climate scientists are nearly unanimous in arguing that climate change poses grave risks to the biosphere and that dramatic steps ought to be taken soon to drastically reduce global emissions, the perhaps surprising conclusion of many economic cost-benefit analyses is that only modest controls of greenhouse gas emissions would be optimal (see, e.g., Nordhaus and Boyer 2003; Nordhaus 2008; Tol 2002a, b).[1] The core conclusion of Nordhaus and others is that rather than diverting resources to reduce emissions now, it would be more equitable and efficient to build up the productive base of economies and postpone emissions reductions.

Economic cost-benefit analyses and economy-climate integrated assessment models (IAMs) have had a strong influence on public policy debates concerning climate change, especially in the United States. It appears that the United States response to the Kyoto negotiations and the treaty

---

[1] See (Kelly and Kolstad 2000) for a survey of different models. (Tol) gives a more recent summary of existing IAMs and offers a proposal on how to aggregate their predictions.

both during the Clinton and Bush (II) administrations was influenced by IAMs (deCanio 2003, 4; Spash 2007) and they are used by the United States Department of Energy to determine the 'social cost of carbon.'[2]

I will argue here that cost-benefit analyses and attempts to determine the optimal climate policy by maximizing intergenerational expected utility in economy-climate integrated assessment models are deeply problematic. First, IAMs face a problem of underdetermination and induction: They are very sensitive to a number of highly conjectural assumptions about economic responses to a temperature and climate regime, for which we have no empirical evidence. Second, they involve several extremely simplifying assumption which cannot be justified empirically. And, third, some of the assumptions underlying the construction of economic models are normative assumptions that reflect value judgments of the modeler.

In criticizing IAMs for being 'value-laden' I do not wish to invoke an ideal of science as value-free. Climate models are value-laden, too, if Eric Winsberg is correct. Normative assumptions, Winsberg argues, enter into climate modeling in two ways in particular: model choices reflect attitudes toward 'inductive risks' and models are constructed with particular purposes and metrics of success in mind (Winsberg forthcoming). As Winsberg argues, however, due to the models' complexity and complicated histories, the normative assumptions are buried in the "nooks and crannies" and can no longer be recovered or made explicit. Thus, normative assumptions cannot be shown to have introduced any systematic biases. This contrasts sharply with the value-ladenness of IAMs, or so I will argue. IAMs are value-laden in ways that systematically affect their predictions and the normative assumptions *can* and therefore *should* be made explicit. My objection to optimization IAMs is that they purport to provide us with precise numbers as policy guidance, yet these numbers are deeply uncertain, ultimately highly conjectural and buried beneath them lie important normative assumptions that need to be made reflected upon explicitly.[3]

---

[2] (Interagency Working Group on Social Cost of Carbon 2010). See (Dietz 2012) for a critical discussion. Economic cost-benefit analysis also greatly influences certain philosophical discussions of public policy. See (Posner and Sunstein 2008) and (Posner and Weisbach 2010). Nordhaus DICE model also influenced (Lomborg 2005), which, Sir Partha Dasgupta says, "although it's understandably tempting to think otherwise, is not by itself a reason for not taking DICE seriously." (Dasgupta 2007)

[3] The economist Stephen deCanio puts this concern this way: "the application of general equilibrium analysis to climate policy has produced a kind of specious precision, in which the assumptions of the analysis masquerade as results that are solidly grounded in theory and the data." (deCanio 2003, 7)

While most IAMs are rather sanguine about prospective damages from climate change, there are exceptions, in particular the model developed by William Cline (1992) and the PAGE model developed by the researchers of the Stern Review (Stern 2007). Even these models, however, assume that global GDP per capita will continue to grow under business as usual and no mitigation measures (Dasgupta 2007) and that environmental and produced goods are perfectly substitutable. Especially the disagreement between the recommendations of the Stern Review and those based on William Nordhaus's DICE model has received a significant amount of attention in the literature. While Nordhaus's model predicts "that economic damages from climate change with no interventions will be in the order of 2.5 percent of world output per year by the end of the 21st century" (P. W. D. Nordhaus 2008, 14), the *Stern Review*'s PAGE model predicts that business as usual will lead to a minimum damage of 5% GDP per year, reaching as high as 20% (Sterner and Persson 2008, 63). Corresponding to these widely diverging predictions concerning expected economic damages, the two models arrive at very different recommendations on optimal abatement measures, Which of these two models, if any, should we accept? Is the disagreement between Nordhaus and Stern purely one between two rival scientific hypotheses that can be settled by an appeal to empirical evidence? I want to argue here that the answer to the second question is 'no' and that the strong divergence in views is further evidence for the fundamental problems faced by any economic cost-benefit approach to climate change.

I will proceed as follows. In the next section I will make some background remarks concerning the role of simplifications, idealizations, and abstractions in scientific modeling. In section three I will examine how both the climate system and economic systems are modeled in integrative assessment models. Due to its prominence in the literature, I will throughout focus on Nordhaus's DICE model as concrete example. Section four critically discusses how economic damages as a result of climate change are modeled with the help of a so-called 'damage function'. In section five I examine how IAMs aggregate utilities across time and the role of future discounting. Section six discusses how modeling uncertainties can significantly affect the predictions of integrated assessment models. Section six offers a summary of my criticisms of IAMs and briefly suggests an alternative role for IAMs as toy models.[4]

Before we begin, it is worthwhile to recall some of the most recent projections from climate models. CO_2 emissions continue to increase and are estimated to have set a new annual record in the year 2011, just as they did in the year 2010. CO_2 concentrations are approaching 400 ppm,

---

[4] Critical examinations of IAMs in the economics literature, from which I have benefited, include (DeCanio 2003; Spash 2005; Ackerman et al. 2009). (Gardiner 2011) contains a philosophical critique of cost-benefit analyses.

which is substantially higher than the pre-industrial level of 280 ppm. According to the 2012 outlook published by the Massachusetts Institute of Technology Joint Program on the Science and Policy (Prinn and Reilly 2012), CO_2 concentrations are projected to reach 816 ppm in 2100, almost tripling preindustrial levels, and CO_2-equivalent concentrations (which also include other greenhouse gases) will reach 1226 ppm CO_2-equivalent in 2100. For global mean surface temperatures the projected 90% likelihood range of warming by 2100 is 3.5° to 6.7°C with a median of about 4.3°C. This range is derived from a joint probability distribution of climate sensitivity and the strength of aerosol forcing based on observed 20th century climate (Prinn and Reilly 2012). In light of these sobering predictions, it is striking that existing IAMs only consider temperature increases between 1° and 3°C—a temperature range that lies fully outside the 90% likelihood range predicted by the MIT model for 2100.

2. Integrated Assessment Models and Idealizations in Science

Integrated assessment models represent extraordinarily complex geophysical and economic processes in an astonishingly simplified manner. This fact is in itself not yet problematic, since it is a general feature of scientific models that they involve idealizations and abstractions. In fact, the search for simplifying assumptions in the form of idealizations and abstractions, which allow us to successfully represent complex phenomena, seems to be at the very heart of the scientific enterprise. Scientific models idealize by representing certain complex real world phenomena in a mathematically simplified and more easily tractable manner; and they abstract by ignoring factors that are taken to be unimportant for certain representational purposes. Thus, in mathematically representing the motion of a billiard ball on a smooth surface we might *idealize* the shape of the billiard ball by treating it as a perfect sphere and we might *abstract* away from any frictional forces between the ball and the table. Moreover, models of complex systems can also contain dependencies that are 'put in by hand' without being fully motivated in terms of the underlying mechanisms or processes that might give rise to them.

The very fact that, then, that IAMs contain gross oversimplifications is not yet an argument against them. Indeed, William Nordhaus takes the simplicity of his DICE model as one of its main virtues and for his most recent book chose the following epigraph, which he attributes to Leonardo da Vinci: "Simplicity is the highest form of sophistication" (Nordhaus 2008) Yet while scientists frequently take simplicity to be a virtue of theories or models, introducing idealizations and abstractions also requires an empirical or theoretical justification. If we want to base predictions or policy recommendations on a model, we need to have good reasons to believe that the simplifications of our model do not introduce any errors too large to significantly affect the model's output and that the model's predictions are largely insensitive to possibly arbitrary or unprincipled choices of model parameters. Simplicity is only a virtue as long as the simplifying assumptions are not in conflict with our representational aims.

Sometimes scientists investigate models which are known to contain empirically ill-founded and unrealistic abstractions and idealizations. Examining how a such 'toy models' behave under a range of values for their parameters arguably can teach us something about the processes modeled. For example, we can use the law of free fall to model the motion of a falling object without air resistance. Yet to what extent this model can be used to model real-world systems and what inferences the model allows us to draw depends on a careful examination of the idealizing assumptions. We can use the model to describe an object falling in an evacuated tube in a physics lecture hall. We can also use the model to describe the motion of an iron ball dropped from a tower. But it would be a mistake to appeal to the model uncritically and unreflectively to describe the motion of a piece of paper falling from the tower. It is a controversial question in the philosophy of science, under what conditions we are licensed to think of the laws governing an idealized model as also applying to de-idealized real-world situation (see Cartwright 1999). But no matter what view one takes on this issue, an idealized free-fall model does not allow us to predict the motion of the piece of paper. Thus, we need to carefully keep in mind limitations on a model's domain of use when reasoning with idealized models.

I will suggest below that IAMs can be useful as 'toy models' that allow us to investigate how, in extremely simplified scenarios, various outcomes may qualitatively and broadly depend on variations in different policy-relevant assumptions. For example, toy models may give us a feel for the risks associated with humanity's continuing business-as-usual approach to climate change. But trying to read off directly from their predictions what climate policy would be optimal to adopt is as legitimate as using Galileo's law of free fall to model the motion of a piece of paper blowing in the wind.

3. Constructing a climate-policy model

3.1 *Climate Sensitivity*

Integrated assessment models couple an economic general equilibrium model to an extremely simplified climate model with the aim of representing the impacts of climate change on human welfare, the impact of changes in economic activity on GHG emissions, and the effect of mitigation measures on economic growth. Thus, the model has two core components, a climate model and an economy model, which are coupled through two different channels: economic activity is assumed to affect climate change through the emission of greenhouse gases (GHGs); and economic activity is modeled as being affected by climate change through a so-called 'damage function'. Optimization IAMs, which will be the focus of my investigation here, are used to determine what the optimal emission abatement strategy would be by maximizing the present value of overall utility, which consists in an aggregate of utilities across time.[5]

Sophisticated full-fledged climate models, so-called 'general circulation models' (GCMs), attempt to answer how average global temperatures respond to an increase in GHG concentrations in the atmosphere. This response is usually expressed in terms of a probability density function (PDF) for the 'climate sensitivity,' which is defined as the equilibrium mean surface temperature response to a doubling in atmospheric CO_2. There are a number of different GCMs—the fourth assessment report of the IPCC cites twenty-two peer-reviewed studies—all with their own PDF for climate sensitivity (See IPCC-4AR, Fig. 9.20 showing the PDF for nine of these studies). The question is how to use the predictions of GCMs as input into an integrated economy-climate model.

One strategy would be to aggregate the PDFs of different GCMs into a single overall PDF, which could then be fed into a probabilistic economy-climate model to derive the expected present (discounted) utility of various mitigation policies. A problem for this approach, however, is that it is unclear that we know how to perform the aggregation.[6] In particular, we do not know how fat the bad upper tail of the aggregated PDF should be (Weitzman 2009). If the individual GCMs can be thought of as independently drawn from a class of models, then the aggregated PDF should have a thinner tail than the PDFs derived from the individual models. If, by contrast the models have not been constructed independently from one another, the tail should be fattened. Moreover, GCMs do not represent possible large fast feedback effects, such as melting of the Greenland and Antarctic ice sheets or large-scale release of methane from Siberian methane clathrates. Trying to include such feedbacks would presumably increase the median temperature rise upward, except we do not know by how much. Also, due to the non-linearity of the climate system, very small uncertainties concerning feedback effects can result in very large uncertainties for the overall climate sensitivity

---

[5] A less ambitious use of IAMs is as modeling costs of climate policies, without including a component in the model representing future damages of climate change on the economy. This second type of model is less problematic—a point to which I will briefly return in the concluding section.

[6] See (Knutti et al. 2008) for a review of uncertainties in climate models. They conclude that "aggregating temperature ranges across scenarios should be avoided" (2661)

and in very fat upper tails for the overall PDF. The problem is that even though it has so far not been possible to include large feedback effects in climate models and the 'tipping point' behavior that may result from these effects, experts do not think that these are dangers that can be ignored and assign it subjectively non-negligible probabilities (Kriegler et al. 2009).

An alternative strategy would be that we treat this as a situation in which we do not have a precise belief about the PDF for temperature increases. According to one view defended in the literature on imprecise belief (see, e.g., Joyce 2010; Gilboa and Schmeidler 1989), the best representation of our uncertainty in light of the evidence in this case would consist in the set of *all* the PDFs associated with the different climate models, in addition, presumably, to a set of probabilities for the risks of large feedback effects. A possible justification for retaining a range of PDFs rather than aggregating them into one is that the aggregation would result in a loss of information: given the available evidence, we know neither which of the range of GCMs best predicts climate sensitivity nor how to aggregate the models, and thus it is best to keep all models in play. Applied to our problem of finding the optimal mitigation policy, this would mean that we should construct multiple probabilistic IAMs, one corresponding to each PDF associated with a GCM, together with PDFs that take expert judgments on the probabilities of large feedback events into account (Kriegler 2009). These IAMs would allow us to calculate a set of expected present discounted utilities. Using the results of our IAMs as a guide in policy making would require a further step: we would have to choose a decision procedure for situations in which we have imprecise beliefs about what the optimal policy would be. One reasonable choice in the present context might be to choose that mitigation policy that has the highest minimal expected utility.

Many prominent IAMs, such as Nordhaus's DICE model, adopt neither of these two strategies, however. Instead, the models assume an average value for the climate sensitivity predicted by a GCM as input from which an optimal policy is derived deterministically. For example, Nordhaus chooses as value for climate sensitivity 3°C, which, he says, is "the center of the IPCC range of for an equilibrium CO2 doubling."[7] As uncertainty analysis Nordhaus then picks eight of his model's parameters, assumes that their values are normally distributed[8] and makes 100 runs of his model using random draws of these parameters. By assumption this procedure reduces uncertainty in two different ways. Not only does the use of a single probability distribution for each

---

[7] Even though the IPCC states the mean value as 3.2°C.

[8] Nordhaus emphasizes that the probability distributions represent his own subjective degrees of belief: "It should be emphasized that these distributions are indeed judgmental and have been estimated by the author." (Nordhaus 2008, 106) Note that the eight parameters that Nordhaus subjects to an uncertainty analysis do not include the parameters I will be concerned with below, representing the exponent of the damage function and the future discount rate.

of the parameters, including the climate sensitivity, presuppose that we know how to aggregate the predictions of different GCMs in a single PDF, but in using a normal distribution without fat tails it is also assumed that low-probability catastrophic outcomes can be safely ignored. Nordhaus justifies his choice of a normal distribution by saying that introducing a fat tailed "is highly speculative" (Nordhaus 2008, 106). But his own choice of distribution is of course no less speculative and suggests the rather optimistic decision rule that in situations of deep uncertainty we are entitled to limit the choice of models to those that are mathematically tractable and conservative in their damage estimates.[9]

3.2. *Modeling welfare*

If we wanted to construct the overall welfare function at the heart of an IAM from the 'ground up,' we would have to sum up utility functions representing billions of individual agents across different times with different preferences and endowments, who have access to a wide variety of different kinds of goods. The aim is then to find the market equilibrium at which the excess demand for each good is equal to zero. Not only would this problem be prohibitively complex to solve, it is also not obvious that it has a unique solution. That no unique equilibrium may exist can be made plausible, for example, be considering an extremely simple model consisting of just two agents existing at different times and two goods (see DeCanio 2003). As Stephen DeCanio shows, whether there are multiple equilibria for this model depends, among other things, on the value of the intertemporal elasticity of substitution, which indicates how easily the two goods can be substituted for each another. If the elasticity of substitution is small—that is, if the two goods are not easily substituted for each other—and if we assume symmetric utility functions for agents at different times, such that each agent exhibits a preference for her own time, then there are multiple equilibria. Only if the value of the elasticity of substitution is large, there is a unique equilibrium (DeCanio 2003, 89-90). This foundational problem is avoided in IAMs through specifying a 'representative consumer' (and, hence, by ignoring any differences in endowments and preferences between agents) and by choosing a utility function that implies a high degree of substitutability between different goods. The general strategy is, of course, well familiar from the natural sciences, including physics, where it is common to model complex situations in simplified ways, with an eye on mathematical tractability. In each case the strategy has to be justified and the assumption of substitutability is far from inocuous, however, as we will see below.

---

[9] (Spash 2007, 709) makes a related criticism of the *Stern Review*: "Thus Stern manage to convert unknown and unknowable futures into events with known probabilities, and miraculously strong uncertainty becomes weak uncertainty."

The present value of overall utility is, thus, represented as a sum over the welfare equivalent aggregate consumption of agents who all are characterized by the same utility function. In principle the concept of welfare equivalent consumption is meant to be very broad and include consumption and enjoyment of any goods that are of value to people, including goods that are not marketable and do not have a market price, including environmental goods and services.[10] This raises the question how environmental goods, such as climate stability or biodiversity are to be valued, for which there are no markets and which are not exchanged. How are environmental and produced goods to be aggregated to arrive at an overall welfare equivalent consumption? It is unclear that there is a procedure for doing this that does not depend on the values of the modeler. In practice, however, the answer often appears to be that, explicit claims to the contrary, non-marketable goods are simply ignored and welfare equivalent consumption is measured in terms of GDP. Clive Spash criticizes IAMs precisely for that reason: GDP cannot be used as a measure of welfare, he argues, "because it ignores non-monetary welfare (e.g. ecosystems functions, biodiversity, aesthetics) and informal economic activity (e.g. housework, or the 'black' economy), is boosted by disasters (e.g. cleanup of oil spills), and is generally concerned with material throughput rather than quality of life." (Spash 2007) Independently of whether one agrees with Spash's criticism, the decision to abstract from non-monetary aspects of quality of life for the purposes of climate policy recommendations and, hence, to model welfare by GDP alone cannot be given an empirical justification but reflects a value judgment by the modeler about the relative importance of non-marketable goods.

To the extent that consumption of environmental services is included at all, the use of a single aggregate notion of consumption as a measure of welfare implies that produced goods and environmental goods are treated as perfectly substitutable for each other (Sterner and Persson 67-68). That is, it is assumed that the rate of exchange between environmental and produced goods is not affected by the relative scarcity of one good with respect to the other and environmental damages can always be substituted for one-to-one by increases in material consumption. As Stephen Gardiner points out (Gardiner 2011), this assumption has the disturbing consequence that overall welfare could continue to increase, even if humans were eventually forced to live under artificial domes due to the negative consequences of climate change, as long as increases in the consumption of produced goods are large enough to make up for the loss in environmental goods. Gardiner takes this to be an objection to the assumption of perfect substitutability, but at the very

---

[10] "Economic welfare should include everything that is of value to people, even if those things are not included in the market place." (Nordhaus 2008, 13)

least Gardiner's observation, like Spash' criticism, reveals the value-ladenness of the economic model.

It is important to distinguish two distinct criticisms in this context. The first, raised by Spash and Gardiner, is the worry that enjoyment of the environment cannot be monetized at all. No matter how high the growth in GDP is in the dome world, the loss in non-monetary welfare cannot be made up for by consumption of produced goods, or so one might argue. The second criticism is that it is unrealistic to assume, *even to the extent that environmental goods can be monetized*, that they will be perfectly substitutable with produced goods. More plausibly we should expect that as environmental goods become scarcer their relative price will go up, which will result in an increased importance of these goods in the overall economy.

Whether or not we treat environmental and produced goods as perfectly substitutable can have significant consequences for the model's predictions. Nordhaus posits an isoelastic utility function $c(t)^{1-\eta}/(1-\eta)$ with diminishing marginal utility η, for a consumer whose consumption at time t is c. Sterner and Persson (2008) examine how Nordhaus's predictions are affected if instead one assumes an economy with just two representative goods that are not perfectly substitutable for each other—a produced good C and an environmental good E—represented by a constant elasticity of substitution (CES) utility function, $U(C, E) = [(1 - \gamma)C^{1-1/\sigma} + \gamma E^{1-1/\sigma}]^{(1-\eta)\sigma/(\sigma-1)}/(1-\eta)$ with elasticity of substitution σ. They show that, as the environmental good becomes scarcer, relative price changes between the two goods will lead to an increased importance of the environmental good. As a result the two-goods model recommends much more drastic abatement measures, even under the assumption of continued growth in the consumption of the produced good, than Nordhaus's original model does.[11]

Sterner and Persson are explicit that their model positing just two goods is merely a toy model used to illustrate the importance of the elasticity of substitution. A more realistic model would have to disaggregate different environmental goods, such as, for example, clean water and pollination services. Their own highly idealized model cannot be used as direct policy advice, they stress, since "it is hard to provide a good empirical estimate for the elasticity of substitution at this level of aggregation (with only two, representative goods) and particularly hard to say how it would evolve over time." (Sterner and Persson 2008, 71) Thus, while we can learn from their highly idealized two-goods model that assumptions concerning the substitutability of different goods have significant consequences for a model's predictions, it is extremely difficult, if not impossible, to know what empirically reasonable assumptions about the elasticity of substitution might be.

---

[11] See also (Neumayer 2007).

The only way environmental goods and damages enter into IAMs that model welfare equivalent consumption purely in terms of GDP is through the so-called 'damage' function that models damages to the economy as a function of environmental factors. This is the issue to which I will turn next.

4. The damage function

As we have seen, IAMs couple a climate model to a welfare function in two ways: first, consumption is taken to affect GHG emissions and, second, climate change is taken to affect consumption through a damage function. Standard approaches to modeling costs associated with abatement measures have been critically examined in (Spash 2007). Here I want to focus on the damage function.

It is common in IAMs to assume that the effects of climate change on the economy can be represented in terms of a single temperature-dependent damage function: economic damages due to climate change are represented as depending on average global temperatures alone. As Spash has argued, this extreme idealization is likely to ignore a range of important factors. First, agriculture and other climate sensitive sectors of the economy depend on other climate parameters, such as precipitation, as well. Second, damages to ecosystems and agriculture are likely to depend not only on what the new equilibrium temperatures will be but also on the *rate* of temperature increases.[12] If average temperatures rise increasingly quickly, as they are projected to do in the second half of this century unless strong abatement measures are undertaken, then ecosystems will have less time to adapt to climate change. Third, the variability and variances in temperatures and precipitation are important as well. With rising temperatures extreme weather conditions are projected to become increasingly common. An increase in the frequency of droughts and years with floods would likely have serious effects on agriculture, even if long-range precipitation averages did not increase appreciably. Thus, there are serious doubts that *any* damage function that represented damages purely as a function of global average temperatures could be legitimately used to predict economic damages associated with climate change.

The choice of particular function is severely underdetermined by the available evidence, as is the structural question as to how one ought to couple the damage function to the welfare function. Nordhaus's choice of damage function is

$$C(T)=1/[1+(T/a)^b].$$

---

[12] The damage function for the FUND model depends on temperature changes as well (Tol 2002a; Tol 2002b).

C is the welfare equivalent consumption as a fraction of what consumption would be without climate change and the values of the parameters a and b are chosen as a=20.46 and b=2. Choosing a quadratic damage function appears to be quite common. But what justifies this choice?

One motivation for this choice is simply the conviction that damages will increase faster than temperatures and that the damage function should be as simple as possible consistent with this constraint. But a quadratic damage function is still relatively benign—a feature the damage function of Nordhaus's DICE model shares with those of the Stern Review's PAGE and Tol's FUND models: For a difficult-to-fathom 8°C increase in global temperatures, the DICE and PAGE models predict 'only' damages amounting to roughly 15% global GDP while the FUND model predicts damages of roughly 6% GDP.[13] Nordhaus offers two additional motivations for his choice. First, he appeals to studies that aim to estimate the effects of climate change on specific sectors of the economy in specific geographic regions; and second, he supports his choice by appealing to 'expert judgment'. A third, and perhaps dominant justification is an appeal to simplicity: a quadratic function is the simplest function that grows faster than linearly.

Nordhaus's sector- and region-specific estimates are partly based on impact-studies and are partly his own estimates or his own estimated extrapolations from existing studies. The most extensive impact studies exist for agricultural damages and here Nordhaus's discussion comes closest to resembling a case of empirical confirmation of the damage function. Nordhaus divides the world into thirteen geographic regions and presents a graph that plots the loss in GDP per doubling of CO_2 concentrations against average temperatures for each region (see fig. 4.1 in W. D. Nordhaus and Boyer 2003). But even here some of the data points are Nordhaus's own guesses or estimates, and the shape of the best fitting function is highly sensitive to adjustments to the small set of data points.[14]

There is an even more serious problem. Even careful impact studies have to rely on evidence concerning a domain far different from the one for which we want to make predictions. Our data concern (at best) relatively modest regional variations in temperature around current

---

[13] From (Interagency Working Group on Social Cost of Carbon 2010, 10).

[14] Two of Nordhaus's estimated data points seem especially problematic. One estimated data point represents "rich OPEC countries" while another represents Africa. One might question the rationale for including rich OPEC countries (with presumably comparatively little agricultural activity) with equal weight in an assessment of the damages of climate change on global agricultural output. And for different reasons one might question Nordhaus's very low estimate of 0.06% loss in GDP for expected agricultural damages in Africa, especially given that in many African countries up to 50% or even 60% of total GDP is from agricultural activity. Contrary to Nordhaus's optimistic estimate, other studies suggest that food security will be severely threatened in many African countries as a result of climate change (Thornton et al. 2011).

averages. Yet we are interested in predictions for global average temperature increases between 3.5° and 6.7°C.[15] Moreover, we cannot expect to get good estimates for prospective damages to agriculture in a given climate zone by comparing data from that region with data from regions that are presently relatively hotter. One problem is that it is currently ill understood how various possible feedback effects will affect regional temperatures and precipitation patterns. Another problem is, as I already noted, that temperature increases will occur very quickly and it is unclear how easily ecological systems will be able to adapt to these changes.

Thus, IAMs face a problem of induction, as the Stern Review emphasizes: "Indeed, the knowledge base on which the cost of climate change is calibrated – specialized studies of impacts on agriculture, ecosystems and so on – is particularly patchy at high temperatures. In principle, the gaps that remain may lead to underestimates or overestimates of global impacts. In practice, however, most of the unresolved issues will increase damage estimates." (Stern 149) That is, the higher the temperature, the less reliable our models are likely going to be and, moreover, the more likely, according to Stern, it is that they will underestimate the projected costs.

The inductive problem can also be illustrated as follows. For an extremely simple toy-model Martin Weitzman compares Nordhaus's damage function with a function that contains a T^6 term in addition to a quadratic term (Weitzman 2012). Given Weitzman's particular choice of parameters, the two functions are empirically indistinguishable for small increases in temperature (given the limitations of the available data) but differ dramatically for larger increases. Weitzman purposefully chooses his damage function to exhibit a 'tipping-point' behavior: damages are relatively small for modest increases in temperature but grow dramatically for increases around 5°C - 6°C. Yet any empirical evidence we have clusters around the very low end of the domain of the damage function and is, thus, unable to distinguish between Nordhaus's quadratic and Weitzman's T^6 -function.

Note that general circulation models (GCMs) considered by climate scientists do not face the inductive problem with the same force, even though they, too, contain large uncertainties. IAMs are constructed on the basis of often very limited empirical evidence or guesswork. In such a situation, the problem of induction looms large: There is very little reason to be confident that the highly idealized simple relationships postulated to hold for economic systems under current temperatures would continue to hold in a world that is 4°C or even 7°C warmer. By contrast, the amount of data

---

[15] This point is stressed in a study on the impact of climate change on United States crop yields: "Although the crop-simulation models were developed and tested over a range of climate conditions, they have not been tested under the temperature conditions suggested by the GCMs" (Adams et al. 220)

on which GCMs are based is vastly larger (including, for example ice core data reaching back 800,000 years) and span much wider temperature ranges than that on which IAMs can be based.

The inductive problem also applies to Nordhaus second strategy for justifying his choice of a quadratic damage function—expert judgment. DeCanio shows that economists do not in fact have a very good track record in forecasting economic activity, even for time spans much shorter as those to which IAMs are applied (DeCanio 2003, ch. 5). Yet let us grant that experts may possess tacit knowledge of economic systems that allows them to estimate the effects of small scale fluctuations in temperatures on economic activity even more accurately than any complicated impact model could. Yet the experience on which such tacit knowledge is based does not include any experience with rapid temperature rises and economic activity in a 4.3°C-world.

The fact that the choice of damage function is unprincipled and cannot be empirically justified would be less problematic if the models' results were largely insensitive to the choice of function. But this is not so, as a sensitivity analysis by some of the researchers involved in the Stern Review shows. The Stern Review treats the exponent as a Monte Carlo parameter ranging from one to three, with a modal estimate of 1.3. (Dietz et al. 2007) show that if the damages were a cubic function of temperature, then this would increase the estimated damages in their model by more than 20% of world output! Now, Nordhaus's DICE model is less sensitive to the choice of damage function (see Weitzman 2010). But the reason for this is that Nordhaus assumes a very high future discount rate—so high that his model appears to be largely insensitive even to severe damages under dramatic future temperature increases. I will examine the role of discount in the following section.

Even Nordhaus's model, however, is sensitive to the question as to how the damage function ought to be coupled structurally to the utility function. The standard choice is to take damages to affect utilities multiplicatively. An alternative option, investigated in (Weitzman 2009), is to couple the damage function additively. As Weitzman shows, this choice significantly affects the predictions of the model. In fact, assuming an additive damage function is formally equivalent to Sterner and Persson's introduction of a second environmental good through a CES-type utility function (Sterner and Persson 2008). As in the case of different exponents of the damage function our empirical evidence based on presently observed temperatures does not allow us to distinguish between the two structural approaches, since they tend to agree for small temperature rises and low damages. Again, the problem of underdetermination and induction looms large.

5. Future Discounting

Any cross-temporal aggregation of utilities faces the problem what relative weight to assign to utilities at different times. It is common practice to discount future utilities with respect to the present. But what is the correct discount rate to use? There is widespread disagreement in the literature not only on what choice of discount rate is appropriate but even on what the proper methodology for choosing a discount rate ought to be. Moreover, the predictions of IAM's are extremely sensitive to what discount rate is chosen. Thus, the stark disagreement between Nordhaus's DICE model and the Stern Review researchers' PAGE model on optimal abatement measures is to a significant extent due to the large difference in the presupposed discount rates between the two models: 5.5% in Nordhaus's case as opposed to 1.5% in the Stern Review. Indeed, it is often claimed that whether an IAM recommends modest or stringent abatement measures is largely determined by the choice of discount rate.

One can distinguish two basic approaches to picking a discount rate—an ostensibly purely empirical or positivist approach and an overtly normative or ethicists' approach. Both positivists and ethicists tend to agree that a starting point for the debate is the Ramsey rule, according to which the overall discount rate ρ consists of two factors δ and η satisfying the equation $\rho = \eta g + \delta$. The factor δ is the so-called 'rate of pure time difference' and captures how much future welfare equivalent consumption is discounted with respect to present consumption purely because it is occurring at a time later than the present. The factor η is the elasticity of the marginal utility of consumption and multiples the rate of growth g. A positive value of η implies a diminishing marginal utility—that is, that the added utility of an additional unit of consumption decreases as overall consumption increases. Thus, even though η is not inherently concerned with intertemporal comparisons, as long as the growth rate g is positive, η works together with δ to discount future increases in utility.

How should one go about choosing values for ρ, δ, and η? According to the positivist approach, advocated for example in (Nordhaus 2008), ρ and g are chosen to match the empirically observed rate of return on capital and the empirically expected growth rate, respectively, and δ and η are then picked to have 'reasonable' values satisfying the Ramsey equation. Any choice other than the empirically observed market rate of return, according to the positivists, would result in inefficient investments, since the discount rate reflects the opportunity costs of investments. An investment's opportunity costs are given by the alternative rate of return available on the investment by putting the money in the bank. If, for example, we assume an interest rate of 6%, then the present value of having \$100 in two years is less than \$100 and is equal to $\$100/1.06^2 = \$89$: Since investing \$89 now will provide us with \$100 in two years, \$100 two years from now should have the same worth to us as $89 today. Future discounting for the positivists is simply a consequence of the requirement that we ought to investment capital and resources as efficiently as possible. That this is also the appropriate approach to discounting in the case of climate change is justified by arguing that "market rates should be used to discount future benefits and costs, because it is the market interest rate that would govern the returns potentially set aside today to compensate future individuals for climate damages that they bear." (Interagency Working Group on Social Cost of Carbon 2010, 19) Note that this argument, too, presupposes perfect substitutability between environmental and produced goods.

By contrast, a normative approach, as adopted for example in (Stern 2007), relies partly on ethical arguments to first pick appropriate values for δ and η, which then, together with assumptions about future growth, determine the overall discount rate ρ. A core premise of the normative approach is that all generations ought to be treated equitably and that the welfare of future generations ought to count as much as that of present generations.

A comprehensive examination of the issue of future discounting would require far more space than I have here.[16] Rather, my aim is more limited. I want to argue that the positivists' view that the choice of discount rate is exclusively an empirical issue and that the correct discount rate can be determined by observing market rates of return on investments is deeply problematic. As in the case of the damage function we do not have the empirical data that would allow us to determine a long-term rate of return on investment of a kind that would be appropriate in the context of evaluating climate policies; and adopting a discount rate involves inherently normative considerations.

First, economists want to use IAMs to offer policy advice. That is the explicit role of the model is normative and part of this role involves the question as to how to value the interests of future generations in comparison with those of present generations.[17] Moreover, to the extent that the positivist approach simply projects certain historical market rates into the future, the approach amounts to the normative position that this the appropriate rate to use in discounting future goods and whatever ethical societal choices are reflected in the presently observed market rates of return are the ones we should continue to endorse.

---

[16] See (Broome 1993) and (Price 1993) for detailed critical examinations of future discounting.

[17] "We are not simply observing the market as we do in positive or empirical studies; we are providing arguments for public action that involve the provision of very complex public goods." (Sterner and Persson 2008, 64-5) See also (Gardiner 2011).

Second, there are many different rates of return and it is unclear which of these might be the correct one to use in discounting future welfare equivalent consumption. For one, there is difference between the private rate of return seen by an investor and the social rate of return to an investment (see Broome 1993). Moreover, different types of investment offer different returns and one needs to decide whether one should adopt a rate equivalent to the rate of return on risky investments or rather a (relatively) risk free rate, such as the historical rate of return on treasury bills.

Third, in evaluating different climate policies with the help of an IAM, we need to make projections about expected rates of return and growth rates for the next fifty, one hundred, or even several hundreds of years. One might think that an appropriate procedure for choosing such long-term growth and interest rates, is simply to take average rates for the last century and use these as projections for the indefinite future. But what this approach ignores is that currently observed market rates of return do not adequately reflect environmental damages, in particular those associated with climate change, and that climate change itself will arguably significantly affect the rate of return on investment. Up until now the only constraints on the energy flow powering industrial growth have been on the production and transmission of power—that is on the source and transmission side—but not on the atmosphere as waste sink with a finite capacity. Moreover, there are huge inertias on the sink side that don't exist on the source side and make projecting potential future rates of return particularly difficult. Expected damages from climate change will occur decades or even centuries after the emissions have occurred and are not reflected in present market rates of return even for investment vehicles with the longest available time horizons.

One response to the uncertainty about future growth and interest rates is to suggest that we should calculate the discount rate as the (weighted) average of a range of possible future interest rates. However, such an averaging would probably result in significantly lower discount rates than Nordhaus, for example recommends, since low rates of return dominate in the proper averaging procedure. Rather than averaging interest rates, the appropriate procedure would be to average the present value of future goods, according to the different interest rates.[18] If we allowed for the possibility that catastrophic climate change would result in negative growth, then the expected discount rate is likely to be very low or even negative.

---

[18] For example, if we assumed as two possibilities an optimistic 6% and a pessimistic 0% annual interest rate for the next century, Then the present value of € 1 million are € 52 033 and € 1 million, respectively. The average of these two values us € 501 474, implying an interest rate of 0,7%, which is much lower than 3%, computed from averaging the interest rates. (See Price 1993, ch. 17)

At first sight this last point appears to be not an objection to a positivist approach to discounting in general, but only to the very high discount rates that some positivists posit. But the objection goes beyond that. Given the many unquantifiable uncertainties concerning both climate predictions and possible economic effects, it is doubtful that we have any clear idea on what the proper range of possible future interest rates should be and what probability distribution we should postulate over that range. Thus, it is unclear that we are in a position of being able to apply the empiricist's criterion of discounting in accord with the expected long-term market rate of return on investment.

Fourth, it is unclear whether the argument from opportunity costs applies to environmental goods. To the extent that the environment provides us also with non-monetary benefits (or costs), these cannot, by their very nature, be invested and reinvested. Thus, considerations of market rates of return fail to apply.

A fifth problem for a purely empirical approach is that the Ramsey equation leaves the values of δ and η underdetermined for given values of ρ and g, yet how the overall rate of return ρ is distributed among δ and η can also make a significant difference to the prediction of the model. δ, recall, is the rate of pure time difference and most normative approaches argue that it follows from the assumption of intergenerational neutrality that this value has to be equal to zero. As Frank Ramsey famously put it: "it is assumed that we do not discount later enjoyments in comparison with earlier ones, a practice which is ethically indefensible and arises merely from the weakness of the imagination." (quoted in Weitzman 2012)[19]

Choices different from zero are sometimes motivated by appealing to empirical evidence suggesting that people have a clear preference for present benefits over future benefits. But there are several problems with such an argument. First, note once more that the aim of IAMs is to issue policy recommendations and it does not follow from the fact that we humans do *as a matter of fact* frequently discount future benefits that this is what we *ought* to do. Individual humans might simply tend to act irrationally in their preference for the present moment, but government policy does not have to do the same. Second, there is evidence that suggests that the proper conclusion to draw from the empirical evidence is not that people asymmetrically discount benefits that occur at later times, but rather that people exhibit a preference for *present* consumption or benefits over *both earlier and later* benefits. (see Price 1993, ch. 8). That is, the observed discounting effect is one that privileges immediate consumption, rather than future discounting and shows the same kind of disregard for past as for future consumption. Such symmetric 'tyranny of the present' cannot be

---

[19] Weitzman, following Ramsey, adopts $\delta=0$. See also (Dasgupta 2007).

used to justify time-asymmetric future-discounting. Third, it is not clear what relevance a characteristic feature of our *personal* preference structure should have for the question of *intergenerational* discounting. It may well be that I value an apple today more highly than the promise of an apple next week. But this does not imply that the utility associated with my consuming the apple today is significantly higher than the utility is of consuming an apple for my grandchildren in 50 years. The former point is associated with the claim that the utility I derive from benefits in the future decreases with my lifetime. But the latter claim implies that the utility future generations derive from future benefits is less than that of benefits for the present generation simply because they were born later in time.

Given the underdetermination of δ and η in the Ramsey equation one could agree with the ethicist that the appropriate value of ρ is zero and then argue that the future ought to be discounted follows from the fact that the elasticity of the marginal utility of consumption η is positive. The discount factor, is thus given by the equation $\rho = \eta g$. This approach amounts to denying that there is *pure* time-discounting—future benefits are not discounted simply because they occur in the future—and the only reason for discounting future utility is due to a decrease in the marginal utility of consumption. η is a measure of the aversion to interpersonal inequality. Thus, Dasgupta has criticized the Stern Review for adopting a low value for η and setting it equal to one: "Curiously, the Review adopts a very *inegalitarian* attitude with regard to the distribution of well-being across people when futurity is not the issue - for example, when comparing the well-beings of the poor and rich in the contemporary world." (Dasgupta 2007) A low value for η, Dasgupta argues, would imply that we should save large amounts for later generations, even if they are richer than us.

Note that the equation $\rho = \eta g$ implies that future costs and benefits are discounted exactly if there is positive economic growth. Even though this assumption is standardly made in much of the relevant literature, it is far from trivial. First, as Price argues, using GDP-growth as a measure of increases in per capita consumption can be misleading, especially for developing countries, since it does not account for that portion of GDP that is used to pay foreign debt or for population growth. A better measure would be gross national income per head, which may stagnate or decline even when GDP increases (Price 1993, ch. 15). Second, GDP growth does not take into account depreciation of natural capital. Thus, even as GDP has increased in the past, the long-term income generating capacity may have declined. Third, it seems optimistic to assume that global GDP could continue to grow if temperatures increase 4.3 or even 6.7°C by the end of this century. If we assumed negative future growth, this would result in negative discounting, with the marginal utility of income tending toward infinity as income decreases and approaches survival income (Price 1993, ch. 15).

Moreover, it is highly doubtful that the idealization of assuming a single discount rate for all consumption is legitimate. Realistically, the elasticity of marginal utility of consumption is expected to differ for different goods and for different income levels, something that is masked in the assumption of a universal discount rate. As Price puts it: "In a world where every person's income grew exponentially at a constant rate, where tastes were constant across time and where elasticity of marginal utility of income was constant across income, where the abundance of all goods and services could be confidently predicted, in every likely circumstance, to grow in proportion to the increasing demand for them, in such a world one could safely discount marginal units of consumption at a single discount rate which reflected the diminishing marginal utility of all consumption. The list of symptoms does not provide a recognizable diagnosis of the real world." (Price 1993, 155)

Finally, η tells us how much we ought to discount *marginal* increases in utility. We cannot conclude from the fact that the utility of an added unit of consumption decreases *at the margins*, that we can discount future consumption in general and that *intramarginal* consumption is discounted as well. More generally, it is questionable that the entire framework of marginal utility theory is appropriate to a discussion of climate policy. It appears widely optimistic that global economic, social, and political structures could remain stable in the light of temperature increase of 6.7°C and could respond through marginal changes and normal 'crisis management'. Far more plausible is that climatic changes will result in catastrophic disruptions and radically different socio-economic structures. The framework of marginal utility theory simply does not apply in this case (see deCanio 2003, 3). Conversely, the use of IAMs to model climate policy has as an implicit presupposition that damages from climate change to economic structures will only be marginal and that perhaps still low probability yet catastrophic events can be safely ignored in our analysis.

Overall, then, the positivist's claim that the parameters relevant for future discounting both can and ought to be determined empirically is deeply problematic. The many idealizing assumptions made in assuming a single universal discount factor are empirically highly questionable. The deep uncertainties involved in modeling the economic impacts of climate change teach us that we simply do not have the empirical data that would allow us reliably to predict interest and growth rates in a way that would allow a meaningful comparison of the costs and benefits associated with competing climate policies. And, finally, any choice of the pure time discount factor and the marginal utility of consumption reflects inherently normative assumptions.

6. Uncertainty and Risk Aversion

Standard IAMs, such as Nordhaus's DICE model, assume a utility function with negative curvature equal to -η, because it is assumed that the marginal utility of additional units of consumption decreases. There is a second reason for positing a positive η (and a utility function with negative curvature) and that is that one assumes that people are risk averse. Nordhaus's DICE model, like most other IAMs, takes a deterministic approach: the model assumes a certain value for the climate sensitivity—that is the temperature rise associated with a doubling in the concentration of atmospheric CO_2—and then calculates what path of abatement measure will maximize overall utility, assuming (i) a damage function and (ii) that abatement measures will reduce economic growth. Yet climate models do not give us a determinate value for the climate sensitivity but rather provide us with a PDF of possible temperature changes for a given rise in CO_2 concentration. Thus, arguably a better approach would be to treat the temperature T as a random variable with a probability density function, which, when fed into the damage function, provides us with a probability distribution over expected economic damages. In this framework η functions as the coefficient of relative risk aversion. That is to say, high values of η imply a larger willingness to pay now for a modest yet almost certain reduction in future growth to drastically reduce the probability of catastrophic future damages.

Interestingly, the effect of η in a probabilistic framework can be reversed from its role in a deterministic framework. In the latter case, as we have seen, positive values of η have the effect of discounting future consumption, under the assumption of positive growth g. Since abatement measures decrease present day economic growth (in addition to future growth), while climate damages only decrease future growth, larger values of η reduce our present-day willingness to pay for abatement measures. But in the probabilistic case increases in η can be associated with a larger willingness to pay to reduce the probability of future catastrophic climate change.[20]

(Weitzman 2010) shows by means of a simple probabilistic toy model how extremely sensitive the results of an economy-climate model can be not just to changes in the overall interest rate, but to choices of how the overall rate is distributed among the rate of pure time difference ρ and a factor of risk aversion η. Weitzman presupposes the same overall interest rate as Nordhaus

---

[20] It is unclear why our risk aversion should be so tightly wedded to the diminishing marginal utility as implied by the use of a single parameter η. In principle different combinations of attitude toward risk and of the utility of marginal increases in consumption seem possible.

but sets $\rho=0$ and therefore assumes a relatively high risk aversion of $\eta=3$. This results in a willingness to pay to mitigate future climate change which is much higher than that implied by Nordhaus choice of discount factors. Moreover, the results are sensitive to the shape of the damage function. Finally, the results are also sensitive to the choice of probability density function for the temperature increase and in particular to how 'fat' we assume the bad tail of the distribution to be. As in the case of the damage function, there are no good empirical arguments that could allow us to decide between different tail distributions.

Weitzman's sensitivity analysis suggests that Nordhaus's recommendation of only modest abatement measures is the result of assuming a relatively benign quadratic damage function coupled with a very high rate of pure time preference. Together these assumptions imply that future climate change cannot appreciably affect present day policy considerations almost no matter how high we assume future temperature increases to be.

7. Conclusion

The predictions of some influential integrated assessment models, such as Nordhaus's DICE model, are extremely sensitive to a number of empirically ill-supported and conjectural assumptions, including the structure of the model, the shape of the damage function, the choice of future discount rate and risk aversion factor, and the precise shape and the fatness of the tails of the probability density function for temperature increases. The gaps between the models' thin empirical basis and the models' predictions is bridged by normative assumptions reflecting value judgments of the modeler, such as the choice of discount factor ρ or the question as to how to treat environmental goods and whether to assume perfect substitutability between environmental and produced goods. Because of their value-ladenness, IAMs should not be used as a direct policy guide without additional arguments supporting the specific normative assumptions made. What is more, the intergenerational nature of the climate change problem threatens that we impose our own perhaps idiosyncratic ethical choices on future generations.

Still we need some criteria to guide policy decisions. What might take the place of using optimization IAMs to provide numerically precise policy recommendations? In this concluding section I briefly want to sketch two possible roles for economy-climate models that can avoid the criticisms raised in this paper. First, rather than feigning precision where none is to be had one could appeal to IAMs as toy models that allow us to explore plausible scenarios or 'stories' concerning possible risks we face from climate change. Instead of relying on IAMs to deliver precise policy recommendations concerning the optimal and most 'efficient' path to maximizing overall utility, we could use them as toy models that allow us examine a range of possible risks we face and the mechanisms by which GHG emissions could lead to economic damages.

IAMs and the sensitivity analyses performed by Weitzman or by Sterner and Persson, seem to be particularly helpful in exploring the grave risks associated with catastrophic climate change. Sometimes drastic abatement measures are urged as an insurance policy against low probability yet catastrophic climate change (see, e.g., Ackerman et al. 2009; Weitzman 2012). While this analogy is useful, it also might suggest unwanted implications. Insurance companies pay compensations in the case of a low probability loss. Yet there is no 'cosmic insurance company', to which we could turn in the event of catastrophic climate change and the compensation model (once more) seems to suggest perfect substitutability between environmental and produced goods. A better analogy, perhaps, is that of technological safety requirements. Political decision makers and regulatory bodies would not approve a permit for a type of aircraft that had a 5% probability of suffering a fatal crash. In fact, many technological systems are constructed to include significant margins of safety and redundancies. Yet we are currently on a GHG emissions path that, according to the MIT model, has a 5% probability of resulting in a temperature increase upwards of 6.7°C by the end of this century. Toy model IAMs, such as Weitzman's, can illustrate the grave risks associated with our current emissions path and help us find what may be appropriate abatement measures that would provide us with a sufficient margin of safety.

Alternatively, highly sophisticated economy-climate models, such as the "MIT Emissions Prediction and Policy Analysis (EPPA) Model" (Babiker et al. 2001) can be used to evaluate the costs associated with various abatement policies. The EPPA model avoids most of the criticisms raised here. Since the aim of the model is not to optimize costs of abatement measures with respect to future costs of climate change, it needs to include neither a damage function nor make normative choices concerning the discount factor. Moreover, the model is sophisticated enough to distinguish between different goods with different elasticities of substitution and different geographic regions between which goods can be exchanged with differing ease. And as long as we are modeling the costs of abatement measures that have a high chance of preventing catastrophic climate change, the framework of marginal utility analysis still applies.[21]

In light of these two kinds of polar opposite uses of economy-climate models—as toy models or as full-blown cost-assessment models—optimization IAMs, such as the DICE or PAGE model, seem to occupy an uneasy halfway house. They are too simplistic to compete empirically with sophisticated economic cost-models, but they are too complex to allow their normative

---

[21] See (Ackerman et al. 2009) for a similar suggestion.

presuppositions and empirically ill-supported assumptions to be perspicuous. Any policy decision concerning emission abatement measures involves normative choices concerning how we value the welfare of future generations and the environment and what attitude toward risk we find acceptable. An ethically responsible climate policy has to address these issues explicitly and should not defer to economic integrated assessment models to make these choices for us.

References

Ackerman, Frank, Stephen DeCanio, Richard Howarth, and Kristen Sheeran. 2009. "Limitations of Integrated Assessment Models of Climate Change." *Climatic Change* 95 (3): 297–315. doi:10.1007/s10584-009-9570-x.

Babiker, Mustafa H. M., John M. Reilly, Monika Mayer, Richard S. Eckaus, Ian Sue Wing, and Robert C. Hyman. 2001. "The MIT Emissions Prediction and Policy Analysis (EPPA) model : revisions, sensitivities, and comparisons of results" (February). http://dspace.mit.edu/handle/1721.1/3574.

Broome, John. 1993. *Counting the Cost of Global Warming*. The White Horse Press.

Cartwright, Nancy. 1999. *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press.

Cline, William R. 1992. *The Economics of Global Warming*. Peterson Institute.

Dasgupta, Partha. 2007. "Commentary: The Stern Review's Economics of Climate Change." *National Institute Economic Review* 199 (1): 4–7.

DeCanio, Stephen J. 2003. *Economic Models of Climate Change: A Critique*. Palgrave Macmillan.

Dietz, Simon. 2012. "The Treatment of Risk and Uncertainty in the US Social Cost of Carbon for Regulatory Impact Analysis." *Economics: The Open-Access, Open-Assessment E-Journal*. doi:10.5018/economics-ejournal.ja.2012-18. http://www.economics-ejournal.org/economics/journalarticles/2012-18.

Dietz, Simon, Chris Hope, Nicholas Stern, and Dimitri Zenghelis. 2007. "Reflections on the Stern Review (1): A Robust Case for Strong Action to Reduce the Risks of Climate Change." *World Economics* 8 (1).

Gardiner, Stephen M. 2011. *A Perfect Moral Storm: The Ethical Tragedy of Climate Change*. Oxford University Press, USA.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with Non-unique Prior." *Journal of Mathematical Economics* 18 (2): 141–153. doi:10.1016/0304-4068(89)90018-9.

Interagency Working Group on Social Cost of Carbon. 2010. "Appendix 15A. Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866. Final Rule Technical Support Document (TSD): Energy Efficiency Program for Commercial and Industrial Equipment: Small Electric Motors. U.S. Department of Energy. Washington, DC, U.S. Department of Energy." http://go.usa.gov/3fH.

Joyce, James M. 2010. "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24 (1): 281–323. doi:10.1111/j.1520-8583.2010.00194.x.

Kelly, David L., and Charles D. Kolstad. 2000. "Integrated Assessment Models for Climate Change Control." In *The International Yearbook of Environmental and Resource Economics 1999/2000: A Survey of Current Issues*, ed. Henk Folmer and Tom Tietenberg. Edward Elgar Publishing.

Knutti, R., M. R. Allen, P. Friedlingstein, J. M. Gregory, G. C. Hegerl, G. A. Meehl, M. Meinshausen, et al. 2008. "A Review of Uncertainties in Global Temperature Projections over the Twenty-First Century." *Journal of Climate* 21 (11) (June): 2651–2663. doi:10.1175/2007JCLI2119.1.

Kriegler, Elmar, Jim W. Hall, Hermann Held, Richard Dawson, and Hans Joachim Schellnhuber. 2009. "Imprecise Probability Assessment of Tipping Points in the Climate System." *Proceedings of the National Academy of Sciences* (March 16). doi:10.1073/pnas.0809117106. http://www.pnas.org/content/early/2009/03/13/0809117106.

Neumayer, Eric. 2007. "A Missed Opportunity: The Stern Review on Climate Change Fails to Tackle the Issue of Non-substitutable Loss of Natural Capital." *Global Environmental Change* 17 (3–4) (August): 297–301. doi:10.1016/j.gloenvcha.2007.04.001.

Nordhaus, Prof William D. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.

Nordhaus, William D., and Joseph Boyer. 2003. *Warming the World: Economic Models of Global Warming*. The MIT Press.

Posner, Eric A., and David Weisbach. 2010. *Climate Change Justice*. Princeton Univ Pr.

Price, Colin. 1993. *Time, discounting, and value*. Oxford; Cambridge, Mass.: Blackwell.

Prinn, Ronald G., and John M. Reilly. 2012. "Energy and Climate Outlook 2012." http://globalchange.mit.edu/research/publications/other/special/2012Outlook.

Spash, Clive. 2005. *Greenhouse Economics: Value and Ethics*. New ed. Routledge.

Stern, Nicholas. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge University Press.

Sterner, Thomas, and U. Martin Persson. 2008. "An Even Sterner Review: Introducing Relative Prices into the Discounting Debate." *Review of Environmental Economics and Policy* 2 (1) (January 1): 61–76. doi:10.1093/reep/rem024.

Tol, Richard. 2002a. "Estimates of the Damage Costs of Climate Change. Part 1: Benchmark Estimates." *Environmental and Resource Economics* 21 (1): 47–73.

———. 2002b. "Estimates of the Damage Costs of Climate Change, Part II. Dynamic Estimates." *Environmental and Resource Economics* 21 (2): 135–160.

———. "On the Uncertainty About the Total Economic Impact of Climate Change." *Environmental and Resource Economics*: 1–20. doi:10.1007/s10640-012-9549-3.

Weitzman, Martin L. 2009. "Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting." *Economics: The Open-Access, Open-Assessment E-Journal*. doi:10.5018/economics-ejournal.ja.2009-39. http://www.economics-ejournal.org/economics/journalarticles/2009-39.

———. 2012. "GHG Targets as Insurance Against Catastrophic Climate Damages." *Journal of Public Economic Theory* 14 (2): 221–244.

Winsberg, Eric. forthcoming. "Values and Uncertainties in the Predictions of Global Climate Models." http://usf.academia.edu/EricWinsberg/Papers/1285650/Values_and_uncertainties_in_the_predictions_of_global_climate_models.