review article

climate and development

Inside the integrated assessment models: Four issues in climate economics

ELIZABETH A. STANTON*, FRANK ACKERMAN and SIVAN KARTHA

Stockholm Environment Institute, US Center, Tufts University, 11 Curtis Avenue, Somerville, MA 02144-1224, USA

Good climate policy requires the best possible understanding of how climatic change will impact on human lives and livelihoods in both industrialized and developing counties. Our review of recent contributions to the climate-economics literature assesses 30 existing integrated assessment models in four key areas: the connection between model structure and the type of results produced; uncertainty in climate outcomes and projection of future damages; equity across time and space; and abatement costs and the endogeneity of technological change. Differences in treatment of these issues are substantial and directly affect model results and their implied policy prescriptions. Much can be learned about climate economics and modelling technique from the best practices in these areas; there is unfortunately no existing model that incorporates the best practices on all or most of the questions we examine.

Keywords: climate economics; integrated assessment models

1. Introduction

There is no shortage of models that join climate to economy with the goal of predicting the impacts of greenhouse gas emissions in the decades to come and offering policy advice on when, where, and by how much to abate emissions. Some models are designed to offer a detailed portrayal of the climate, or the process of economic growth, or the feedback between these two systems; others focus on the long-run or the short-run, economic damages or environmental damages, carbon-based energy sectors or abatement technology. The best models produce results that inform and lend clarity to the climate policy debate. Some models surprisingly conclude – in direct contradiction of the urgency expressed in the scientific literature – that rapid, comprehensive emissions abatement is both economically unsound and unnecessary. Some models seem to ignore (and implicitly endorse the continuation of) gross regional imbalances of both emissions and income.

Good climate policy requires the best possible understanding of how climatic change will impact on human lives and livelihoods, in industrialized countries and in developing countries. No model gets it all right, but the current body of climate-economics models and theories contains most of the ingredients for a credible model of climate and development in an unequal world.

Unfortunately, many climate-economics models suffer from a lack of transparency, which affects both their policy relevance and their credibility. Building a model of the climate and the economy inevitably involves numerous judgement calls; debatable judgements and untestable hypotheses turn out to be of great importance in determining the policy recommendations of climate-economics models and should be visible for debate. A good climate-economics model

■ *Corresponding author. ***E-mail:*** **liz.stanton@sei-us.org**

doi:10.3763/cdev.2009.0015 ISSN: 1756-5529 (print), 1756-5537 (online) www.earthscanjournals.com earthscan

would be transparent enough for policy relevance, but still sophisticated enough to get the most important characteristics of the climate and the economy right.

Our review of recent contributions to the climate-economics literature assesses 30 existing integrated assessment models (IAMs) in four key areas:

1. Choice of model structure and the type of results produced
2. Uncertainty in climate outcomes and the projection of future damages
3. Equity across time and space
4. Abatement costs and the endogeneity of technological change

These models were chosen based on their prominence in the climate-economics literature over the last 10 years. Most of them have both climate and economic modules, and report results as damages in money values or as a share of GDP. A few models discussed here are better classified as physical impact IAMs, which report results in terms of physical damages.[1] In addition, a few models treat emissions as exogenous to the model structure.[2]

The next four sections of this review evaluate the body of existing climate economics models in terms of these key model characteristics, with illustrative examples of both problems and solutions taken from the literature. The concluding section summarizes our findings and their implications for the construction of climate-economics models.

2. Choice of model structure

This review examines 30 climate-economics models, all of which have been utilized to make contributions to the IAM literature within the last 10 years.[3] These models fall into five broad categories, with some overlap: welfare optimization, general equilibrium, partial equilibrium, simulation, and cost minimization (see Table 1).[4] Each of these structures has its own strengths and weaknesses and each provides a different perspective on the decisions that are necessary for setting climate and development policy. In essence, each model structure asks a different question and that question sets the context for the results it produces.

2.1. Differences in model structures

2.1.1. Welfare optimization models

Welfare optimization models tend to be fairly simple, which adds to their transparency. The production of goods and services causes both emissions and economic output, which can be used either for consumption or investment. Greenhouse gas emissions affect the climate, causing damages that reduce production. Abatement reduces emissions but causes costs that

TABLE 1 Climate-economics models reviewed in this study

Model category	Global	Regionally disaggregated
Welfare maximization	DICE-2008; ENTICE-BR; DEMETER-1CCS; *MIND*	RICE-2004; FEEM-RICE; FUND; MERGE; CETA-M; GRAPE; AIM/Dynamic Global
General equilibrium	JAM; IGEM	IGSM/EPPA; SMG; WORLDSCAN; ABARE-GTEM; G-CUBED/MSG3; MS-MRT; AIM; IMACLIM-R; WIAGEM
Partial equilibrium		MiniCAM; *GIM*
Simulation		PAGE-2002; ICAM-3; E3MG; *GIM*
Cost minimization	GET-LFL; *MIND*	DNE21+; MESSAGE-MACRO

Note: Italics indicate that a model falls under more than one category.

reduce economic output. The models maximize the discounted present value of welfare (which grows with consumption, although at an ever-diminishing rate)[5] across all time periods by choosing how much emissions to abate in each time period, where abatement costs reduce economic output (see Figure 1). The process of discounting welfare (or 'utility', which is treated as a synonym for welfare here and in many models) requires imputing speculative values to non-market 'goods' like ecosystems or human lives, as well as assigning a current value to future costs and benefits. Dynamic optimization models – including all of the welfare optimization and cost minimization models reviewed here – solve for all time periods simultaneously, as if decisions could be made with perfect foresight.[6]

Our review of climate-economics models includes four global welfare optimization models: DICE-2007 (Nordhaus, 2008), ENTICE-BR (Popp, 2006), DEMETER-1CCS (Gerlagh, 2006) and MIND (Edenhofer et al., 2006b), and seven regionally disaggregated welfare maximization models: RICE-2004 (Yang and Nordhaus, 2006), FEEM-RICE (Bosetti et al., 2006), FUND (Tol, 1999), MERGE (Manne and Richels, 2004), CETA-M (Peck and Teisberg, 1999), GRAPE (Kurosawa, 2004) and AIM/Dynamic Global (Masui et al., 2006).

2.1.2. General equilibrium models

General equilibrium models represent the economy as a set of linked economic sectors (markets for labour, capital, energy etc.). These models are solved by finding a set of prices that have the effect of 'clearing' all markets simultaneously (i.e. a set of prices that simultaneously equate demand and supply in every sector). General equilibrium models tend to use 'recursive dynamics' – setting prices in each time period and then using this solution as the beginning point for the next period (thus assuming no foresight at all). Eleven general equilibrium models are reviewed in this study: JAM (Gerlagh, 2008), IGEM (Jorgenson et al., 2004), IGSM/EPPA (Babiker et al., 2008), SMG (Edmonds et al., 2004), WORLDSCAN (Lejour et al., 2004), ABARE-GTEM (Pant, 2007), G-CUBED/MSG3 (McKibbin and Wilcoxen, 1999), MS-MRT (Bernstein et al., 1999), AIM (Kainuma et al., 1999), IMACLIM-R (Crassous et al., 2006) and WIAGEM (Kemfert, 2001).

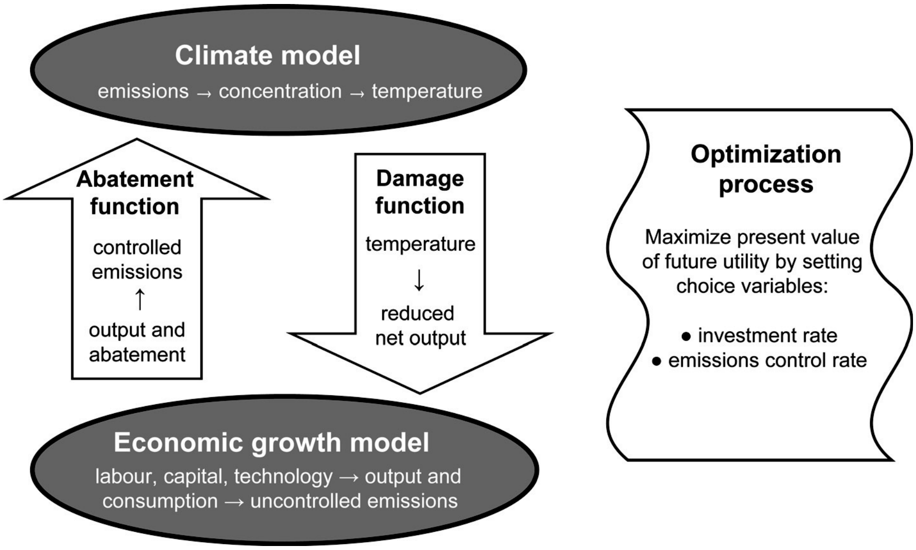

FIGURE 1 Schematic representation of a welfare optimizing IAM

In dynamic versions of general equilibrium theory, multiple equilibria cannot always be ruled out (Ackerman, 2002).[7] When multiple equilibria are present, general equilibrium models yield indeterminate results that may depend on details of the estimation procedure. For this reason, an assumption of constant or decreasing returns is often added to their production functions, an arbitrary theoretical restriction which is known to assure a single optimal result (Köhler et al., 2006). Because increasing returns to scale are important to accurate modelling of endogenous technological change, general equilibrium modellers must skirt between oversimplifying their representation of the energy sector and allowing unstable model results.

2.1.3. Partial equilibrium models

Partial equilibrium models – for example, MiniCAM (Clarke et al., 2007) and GIM (Mendelsohn and Williams, 2004) – make use of a subset of the general equilibrium apparatus, focusing on a smaller number of economic sectors by holding prices in other sectors constant; this procedure also can help to avoid problems with increasing returns to scale.

2.1.4. Simulation models

Simulation models are based on off-line predictions about future emissions and climate conditions; climate outcomes are determined by an economic model of production, damages, consumption, investment and abatement costs. A predetermined set of emissions by time period dictates the amount of carbon that can be used in production and model output includes the cost of abatement and cost of damages. Simulation models cannot, in and of themselves, answer questions of what policy makers *should* do to maximize social welfare or minimize social costs. Instead, the simulation models reviewed in this study – PAGE2002 (Hope, 2006), ICAM-3 (Dowlatabadi, 1998), E3MG (Barker et al., 2006) and GIM (Mendelsohn and Williams, 2004) – estimate the costs of various likely future emission paths.

2.1.5. Cost minimization models

Cost minimization models are designed to identify the most cost-effective solution compatible with a particular objective. Some cost minimization models explicitly include a climate module, while others abstract from climate by representing only emissions, and not climatic change and damages. The four cost minimization models included in this review – GET-LFL (Hedenus et al., 2006), MIND (Edenhofer et al., 2006b), DNE21+ (Sano et al., 2006) and MESSAGE-MACRO (Rao et al., 2006) – have very complex 'bottom-up' energy supply sectors, modelling technological choices based on detailed data about specific industries. Three of these models, excluding GET-LFL, combine a bottom-up energy supply sector with a top-down energy end-use sector, modelling technology from the vantage point of the macroeconomy.

2.2. Evaluation of model structures

The different types of model structures provide results that inform climate and development policy in very different ways. All five categories have strengths and weaknesses. Many of the best-known IAMs attempt to find the 'optimal' climate policy, one that maximizes long-term human welfare. This calculation depends on several unknowable or controversial quantities, including the numerical measurement of human welfare, the physical magnitude and monetary value of all current and anticipated climate damages, and the relative worth of future versus present benefits.

General equilibrium models can be extremely complex, combining very detailed climate models with intricate models of the economy; yet despite their detail, general equilibrium models' reliance on decreasing returns is a serious limitation to their usefulness in modelling endogenous technological change. When models are overly complex, both transparency and the plausibility of final results are compromised (this latter point is discussed in more detail in section 5 of this article).[8] Partial equilibrium models circumvent the problem of increasing returns, at the cost of a loss of generality. In some cases, there appears to be a problem of spurious precision in overly elaborated models of the economy, with, for example, projections of long-term growth paths for dozens of economic subsectors.

Simulation models are well suited for representing uncertain parameters and for developing IAM results based on well-known scenarios of future emissions, but their policy usefulness is limited by a lack of feedback from their economic damages and abatement modules to emissions. Finally, cost minimization models address policy issues without requiring calculations of human welfare in money terms, but existing cost minimization models may suffer from the same tendency towards spurious precision exhibited in some general and partial equilibrium models.[9]

3. Uncertain outcomes and projections of future damages

IAMs inevitably rely on forecasts of future climate outcomes and the resulting economic damages, under conditions that are outside the range of human experience.[10] This aspect of the modeling effort raises two related issues: the treatment of scientific uncertainty about climate change and the functional relationships used to project future damages.

3.1. Scientific uncertainty in climate outcomes

There are inescapable scientific uncertainties surrounding climate science, for instance, in the climate sensitivity parameter (the temperature increase resulting from a doubling of CO_2 concentrations). As a result, low-probability, enormous-cost climate outcomes cannot be ruled out; the response to these extreme risks is often central to policy debate and would ideally be incorporated in economic models of climate change. Yet we found that most IAMs use central or average estimates to set parameter values, typically addressing uncertainty through a few sensitivity analyses of responses to selected changes in parameter values.[11] Those few models that express parameter values as distributions often use truncated distributions that inappropriately exclude or de-emphasize low-probability, high-cost catastrophes.

Uncertainty is inescapable despite the ever-expanding body of climate research, because there are only a limited number of empirical observations relevant to questions such as estimation of the climate sensitivity parameter. As a result, the best estimates of the relevant probability distributions inevitably exhibit 'fat tails', meaning that extreme outcomes are much more likely than a normal distribution would imply (Weitzman, 2008). According to Weitzman, an economist who has raised this problem in recent debates, IPCC (2007) data imply that an atmospheric concentration of 550 ppm of CO_2-equivalent would lead to a 98th percentile chance of 6 °C increase in temperature, a point at which we 'are located in the terra incognita of … a planet Earth reconfigured as science fiction… [where] mass species extinctions, radical alterations of natural environments, and other extreme outdoor consequences will have been triggered by a geologically-instantaneous temperature change that is significantly larger than what separates us now from past ice ages.' (Weitzman, 2007: 716).[12]

In the face of such worst case risks, it is misleading to look only at the most likely range of conditions. The future will happen only once. Suppose we knew that there were one hundred equally likely future scenarios, of which only one or a few would lead to truly catastrophic climate change. If we plan well for the most likely outcomes but instead one that we consider unlikely comes to pass, will we be comforted by our parsimonious rationality?

A thorough treatment of uncertainty through Monte Carlo analysis that varies multiple unknown parameters is seen in just a few IAMs. Even then it is difficult to fully explore the parameter space, especially given the fat-tailed distributions that characterize many key climate parameters and their poorly understood interactions.

One of the best-known models that incorporates Monte Carlo analysis is Hope's PAGE2002 (Hope, 2006), the model used in the Stern Review (Stern, 2006). PAGE2002 includes triangular distributions for 31 uncertain parameters; Hope's standard analysis is based on 1000 iterations of the model; as in other multivariate Monte Carlo analyses, he uses Latin Hypercube sampling[13] to select the uncertain parameters. This level of sensitivity analyses has a major impact on results. For the Stern Review, replacing the Monte Carlo analysis with a deterministic analysis using the modal parameter values decreases annual climate damages by an average of 7.6% of world output (Dietz et al., 2007).

The 31 uncertain parameters in PAGE2002 include two sets of seven regional parameters, but there are still 19 orthogonal (i.e., presumed unrelated or independent) parameters with independent distributions to be sampled for each iteration. This makes it essentially impossible for a Monte Carlo analysis to explore simultaneous worst cases in many or most of the parameters. To have, on average, at least one iteration with values from the worst quintile for all 19 parameters, it would be necessary to run the model an unimaginable 20 trillion times – a result of the so-called 'curse of dimensionality' (Peck and Teisberg, 1995).[14] Of course, parameters that are treated as orthogonal in the model could be interdependent in the real world. Greater interdependency among parameters would make seemingly rare extreme events (based on multiple worst case parameter values) more likely. But as long as these parameters are represented as orthogonal in probabilistic IAMs, a high number of iterations will be necessary to assure even a single run with extreme values for multiple parameters. In PAGE2002, with 1000 iterations, it is highly unlikely that there are any results for which more than a few parameters are assigned 95th percentile or worse values.

Only one other model among those reviewed has a built-in method of randomizing parameter values. Carnegie Mellon's ICAM is a stochastic simulation model that samples parameter values from probability distributions for 2000 parameters for an unspecified number of iterations (Dowlatabadi, 1998). An enormous number of iterations would be necessary to assure even one result with low-probability values for any large subset of these parameters. With any plausible number of iterations, the 'curse of dimensionality' means that the primary choice being made by the Monte Carlo sampling is the selection of which parameters happen to have their worst cases influence the results of the analysis.

Several studies have added a Monte Carlo analysis onto other IAMs reviewed here. Nordhaus and Popp (1997) ran a Monte Carlo analysis on a modification of an earlier version of the DICE model – called PRICE – using eight uncertain parameters and 625 iterations, with five possible values for each of three parameters and a variation on Latin Hypercube sampling for the rest. Nordhaus also has run a Monte Carlo simulation using DICE-2007 (Nordhaus, 2008) with eight parameters and 100 iterations. Kypreos (2008) added five stochastic parameters to MERGE and runs 2500 iterations; Peck and Teisberg (1995) added one stochastic parameter to CETA-R with an unreported number of iterations; and Scott and co-authors (1999) added 15 stochastic parameters to MiniCAM with an unreported number of iterations. Webster et al. (1996) take a different approach to modelling uncertainty in ISGM/EPPA by using a collocation method that approximates the model's response as a polynomial function of the uncertain parameters.

None of the models reviewed here assumes fat-tailed distributions and reliably samples the low-probability tails. Therefore, none of the models provides adequate representation of worst case extreme outcomes – which are unfortunately not unlikely enough to ignore.

3.2. Projecting future damages

Most IAMs have two avenues of communication between their climate model and their economic model: a damage function and an abatement function (see Figure 1). The damage function translates the climate model's output of temperature – and sometimes other climate characteristics, such as sea-level rise – into changes to the economy, positive or negative.

Many models assume a simple form for this relationship between temperature and economic damage, such that damages rise in proportion to a power of temperature change:

$$D = aT^b \tag{1}$$

where D is the value of damages (in dollars or as a percent of output), T is the difference in temperature from that of an earlier period and the exponent b determines the shape or steepness of the curve. Damages are calculated for multiple time periods, often at intervals of 5 or 10 years, over the course of as long as 600 years; annual damages for any given year are calculated by interpolation between adjacent estimates.[15] Implicitly, the steepness of the damage function at higher temperatures reflects the probability of catastrophe – a characteristic that can have a far more profound impact on model results than small income losses at low temperatures.

Our literature review revealed three concerns with damage functions in existing IAMs: the choice of exponents and other parameters for many damage functions are either arbitrary or under-explained; the form of the damage function constrains models' ability to portray discontinuities; and damages are commonly represented in terms of losses to income, not capital.[16]

3.2.1. Arbitrary exponent

DICE, like a number of other models, assumes that the exponent in the damage function is 2 – that is, damages are a quadratic function of temperature change.[17] The DICE-2007 damage function was assumed to be a quadratic function of temperature change with no damages at 0 °C temperature increase, and damages equal to 1.8% of gross world output at 2.5 °C; this implies, for example, that only 10.2% of world output is lost to climate damages at 6 °C. (Nordhaus, 2007a).[18] Numerous subjective judgements, based on fragmentary evidence at best, are incorporated in the point estimate of 1.8% damages at 2.5 °C (much of the calculation is unchanged from Nordhaus and Boyer (2000), which provides a detailed description). The assumption of a quadratic dependence of damage on temperature rise is even less grounded in any empirical evidence.

Our review of the literature uncovered no rationale, whether empirical or theoretical, for adopting a quadratic form for the damage function – although the practice is endemic in IAMs, especially in those that optimize welfare.[19] PAGE2002 (Hope, 2006) uses a damage function calibrated to match DICE, but makes the exponent an uncertain (Monte Carlo) parameter, with minimum, most likely and maximum values of 1.0, 1.3 and 3.0, respectively. Sensitivity analyses of the Stern Review (Stern, 2006) results, which were based on PAGE2002, show that fixing the exponent at 3 – assuming damages are a cubic function of temperature – increases average annual damages across the 200 year forecast horizon (above the Stern Review's business-as-usual baseline) by a remarkable 23% of world output (Dietz et al., 2007). Thus the equally arbitrary assumption that damages are a cubic, rather than quadratic, function of temperature would have a large effect on IAM results, and consequently on their policy implications.

3.2.2. Continuity

Damage functions are often defined to be continuous across the entire range of temperature rise, even though it is far from certain that climate change will in fact be gradual and continuous. Several climate feedback processes point to the possibility of an abrupt discontinuity at some uncertain temperature threshold or thresholds. However, only a few IAMs instead model damages as discontinuous, with temperature thresholds at which damages jump to much worse, catastrophic outcomes.

Two leading models incorporate some treatment of catastrophic change, while maintaining their continuous, deterministic damage functions. MERGE (Manne and Richels, 2004) assumes all incomes fall to zero when the change in temperature reaches 17.7 °C – which is the implication of the quadratic damage function in MERGE, fit to its assumption that rich countries would be willing to give up 2% of output to avoid 2.5 °C of temperature rise. This formulation deduces an implicit level of catastrophic temperature increase, but maintains the damage function's continuity. DICE-2007 (Nordhaus, 2007b) models catastrophe in the form of a specified (moderately large) loss of income, which is multiplied by a probability of occurrence (an increasing function of temperature), to produce an expected value of catastrophic losses. This expected value is combined with estimates of non-catastrophic losses, to create the DICE damage function; that is, it is included in the quadratic damage function discussed above.

In the PAGE2002 model (Hope, 2006), the probability of a catastrophe increases as temperature rises above some specified temperature threshold. The threshold at which catastrophe first becomes possible, the rate at which the probability increases as temperature rises above the threshold, and the magnitude of the catastrophe when it occurs, are all Monte Carlo parameters with ranges of possible values.

3.2.3. Income damages

Damages are commonly modelled in IAMs as losses to economic output, or gross domestic product (GDP), and therefore losses to income (GDP per capita) or consumption, leaving the productive capacity of the economy (the capital stock) and the level of productivity undiminished for future use. For example, non-catastrophic damages in the DICE-2007 model (Nordhaus, 2007a) include impacts to agriculture, 'other vulnerable markets', coastal property from sea-level rise, health, time-use and 'human settlements and natural ecosystems', all of which are subtracted directly from total economic output. In reality, many of these categories are reductions to the capital stock and not directly to income, especially coastal property and human settlements damages. Others have multi-period effects on the marginal productivity of capital or labour, that is, the ability of technology to transform capital and labour into income; damages to agricultural resources and health are good examples of longer-term changes to productivity.

When damages are subtracted from output, the implication is that these are one time costs that are taken from current consumption and investment, with no effects on capital, production or consumption in the next period – an unrealistic assumption even for the richest countries, as attested by the ongoing struggle to rebuild New Orleans infrastructure, still incomplete three years after Hurricane Katrina. FUND (Tol, 1999) is unusual among welfare optimizing IAMs in that it models damages as one-time reductions to both consumption and investment, where damages have lingering 'memory' effects determined by the rate of change of temperature increase.

4. Equity across time and space

Most climate economic models implicitly assume that little attention is needed to the problems of equity across time and space. In the area of intertemporal choice, most models have high discount rates that inflate the importance of the short-term costs of abatement relative to the long-term benefits of averted climate damage. Together with the common assumption that the world will grow richer over time, discounting gives greater weight to earlier, poorer generations relative to later, wealthier generations. Equity between regions of the world, in the present or at any moment in time, is intentionally excluded from most IAMs, even those that explicitly treat the regional distribution of impacts.

4.1. Equity across time

The impacts of climate change, and of greenhouse gas mitigation, will stretch centuries or even millennia into our future. Models that estimate welfare, income or costs over many years must somehow value gains and losses from different time periods. There are two leading approaches.

The early work of Ramsey (Ramsey, 1928) provides the basis for the 'prescriptive' approach, in which there are two components of the discount rate: the rate of pure time preference, or how human society feels about costs and benefits to future generations, regardless of the resources and opportunities that may exist in the future; and a wealth-based component – an elasticity applied to the rate of growth of real consumption – that reflects the diminishing marginal utility of income[20] over time as society becomes richer.

Algebraically, the discount rate, r(t), combines these two elements: it is the rate of pure time preference, ρ, plus the product of the elasticity of marginal utility with respect to consumption per capita, η, and the growth rate of income or consumption per capita, g(t).

$$r(t) = \rho + \eta g(t) \quad (2)$$

Some models use the alternative, 'descriptive' approach to discounting, where the market rate of interest or capital growth is taken to represent the discount rate.[21] These analyses typically either set the discount rate at 5%, or at an unspecified market rate of interest (e.g. Charles River Associates' MS-MRT (Bernstein et al., 1999), a general equilibrium model).

Because climate change is a long-term problem involving long time lags, climate-economics models are extremely sensitive to relatively small changes in the assumed discount rate. There are long-standing debates on the subject which are summarized well in the Stern Review (Stern, 2006). Remarkably, the model descriptions for many IAMs do not state the discount rate or methodology they use, even when discussing discounting.

Choices about the discount rate inevitably reflect value judgements made by modellers. The selection of a value for the pure rate of time preference is a problem of ethics, not economic theory or scientific fact. Pure time preference of zero would imply that (holding real incomes constant) benefits and costs to future generations are just as important as the gains and losses that we experience today. The higher the rate of pure time preference, the less we value harm to future generations from climate change and the less we value the future benefits of current actions to avert climate change. Pure rates of time preference found in this literature review range from 0.1% in the Stern Review's PAGE2002 analysis (Hope, 2006) to 3% in RICE-2004 (Yang and Nordhaus, 2006).

Only a few model descriptions directly state their elasticity of marginal utility of consumption, although the use of this elasticity, implying that marginal utility declines as consumption grows, is common to many IAMs. In DICE-2007 (Nordhaus, 2008), the elasticity of the marginal utility of consumption is set at 2, and the discount rate declines from 4.7% in 2005 down to 3.5% in 2395. In the Stern Review's version of PAGE2002 (Hope, 2006), the elasticity of the marginal utility of consumption is set at 1, and the discount rate averages 1.4%.

A higher elasticity of marginal utility of consumption reflects a greater emphasis on equity: the larger the elasticity, the greater the value to social welfare of an increase in consumption for a poorer person, versus a richer one.[22] However, in a global model – lacking regional disaggregation – there is only one utility function for the world as a whole. The practical upshot of this is that the diminishing marginal utility of income is applicable only in comparisons across time (e.g. the present generation versus the future) and not in comparisons across different regions or socio-economic characteristics (e.g. Africa versus north America today, or at any given point in time).

The four cost minimization models included in this literature review – GET-LFL (Hedenus et al., 2006), MIND (Edenhofer et al., 2006b), DNE21+ (Sano et al., 2006) and MESSAGE-MACRO (Rao et al., 2006) – all report a 5% discount rate.[23] The ethical issues involved in discounting abatement costs are somewhat more straightforward than those involved in discounting welfare. Abatement technologies have well-defined monetary prices, and thus are more firmly situated within the theoretical framework for which discounting was developed. Many abatement costs would occur in the next few decades – over spans of time which could fit within the lifetime and personal decisions of a single individual. To pay for \$1000 worth of abatement 50 years from now, for example, one can invest \$230 today in a low-risk bond with 3% annual interest. On the other hand, welfare optimization models must inevitably assign subjective, contestable values to the losses and gains to future generations that are difficult to monetize, such as the loss of human life or the destruction of ecosystems. No investment today can adequately compensate for a loss of life or irreversible environmental damage; and even if an agreeable valuation were established, there is no existing or easily imagined mechanism for compensating victims of climate change several hundred years in the future.

4.2. Equity across space

IAMs that optimize welfare for the world as a whole – modelled as one aggregate region – maximize the result of a single utility function by making abatement and investment choices that determine the emissions of greenhouse gases; emissions then determine climate outcomes and damages, one of the inputs into utility. This utility function is a diminishing function of per capita consumption. The IAM chooses emission levels for all time periods simultaneously – when more emissions are allowed, future periods lose consumption to climate damages; when emissions are lowered, abatement costs decrease current consumption.

The model's optimizing protocol (or more picturesquely, the putative social planner) balances damages against abatement costs with the goal of maximizing utility – not income or consumption. Because utility is modelled with diminishing returns to consumption, the value to society of a given cost or benefit depends on the per capita income level at the time when it occurs. A change to income in a rich time period is given a lower weight than an identical change to income in a poor time period (even if the rate of pure time preference is zero). If, as usual, per capita income and consumption are projected to keep growing, the future will be richer than the present.[24] Under that assumption, the richer future matters less, in comparison to the relatively poorer present.

Regional welfare optimizing IAMs apply the same logic, but with separate utility functions for each region. The model is solved by choosing abatement levels that maximize the sum of utility in all regions. Seemingly innocuous, the disaggregation of global IAMs into component regions raises a gnarly problem for modellers: with identical, diminishing marginal returns to income in every region, the model could increase utility by moving income towards the poorest regions. This could be done by reallocating responsibility for regional damage and abatement costs, or inducing transfers between regions for the purpose of fostering technical change, or funding adaptation, or purchasing emission allowances or any other channel available in the model for interregional transfers.

Modellers have typically taken this tendency toward equalization of income as evidence of the need for a technical fix. In order to model climate economics without any distracting rush toward global equality, many models apply the little-known technique of 'Negishi weights' (Negishi, 1972). Stripped of its complexity, the Negishi procedure assigns larger weight to the welfare of richer regions, thereby eliminating the global welfare gain from income redistribution.[25]

In more detail, the technical fix involves establishing a set of weights for the regional utility functions. The model is run first with no trade or financial transfers between regions; the regional pattern of per capita income and marginal product of capital from that autarkic (no-trade) run is then used to set the so-called Negishi weights, for each time period, that equalize the marginal product of capital[26] across all regions. Since the marginal product of capital is higher in lower-income regions, the Negishi weights give greater importance to utility in higher-income areas. In a second iteration, the normal climate-economics model, with transfers possible between regions, is restored, and the Negishi weights are hard-wired into the model's utility function. The result, according to the model descriptions, is that the models act as if the marginal product of capital were equal in all regions and, therefore, no transfers are necessary to assuage the redistributive imperative of diminishing marginal returns.[27] The (usually) unspoken implication is that the models are acting as if human welfare is more valuable in the richer parts of the world.

Describing the use of Negishi weights as a mere technical fix obscures a fundamental assumption about equity. Negishi weights cause models to maximize welfare as if every region already had the same income per capita – suppressing the obvious reality of vastly different regional levels of welfare, which the models would otherwise highlight and seek to alleviate (Keller et al., 2003; Manne, 1999; Nordhaus and Yang, 1996).

In IAMs that do not optimize welfare, Negishi weights are not used and interregional effects can, therefore, remain more transparent. For example, in PAGE2002 (Hope, 2006) – a simulation model that reports regional estimates – no radical equalization of per capita income across regions occurs, but utility is not being maximized, and the simulations do not claim to represent optimal policy outcomes.[28]

By including discounting over time as well as Negishi weights, welfare optimizing IAMs accept the diminishing marginal utility of income for intergenerational choices, but reject the same principle in the contemporary, interregional context. Some justification is required if different rules are to be applied in optimizing welfare across space than those used when optimizing welfare across time. At the very least, a climate-economics model's ethical implications should be transparent to the end users of its analyses. While ethical concerns surrounding discounting have achieved some attention in policy circles, the highly technical but ethically crucial Negishi weights are virtually unknown outside the rarified habitat of modelers and welfare economists. The Negishi procedure conceals one strong, controversial assumption about welfare maximization, namely that existing regional inequalities are not legitimate grounds for shifting costs to wealthier regions, but inequalities across time are legitimate grounds for shifting costs to wealthier generations. Other assumptions, needless to say, could be considered.

5. Abatement costs and the endogeneity of technological change

The analysis of abatement costs and technological change is crucial to any projection of future climate policies. An unrealistic picture of fixed, predictable technological change, independent of public policy, is often assumed in IAMs – as is the treatment of investment in abatement as a pure loss. These choices are mathematically convenient, but prevent analysis of policies to promote and accelerate the creation of new, low-carbon technologies. This oversimplification supports the questionable conclusion that the best policy is to avoid immediate, proactive abatement, and wait for automatic technical progress to reduce future abatement costs.

5.1. Choices in modelling abatement technology

There have been rapid advances in recent years in the area of modelling endogenous technological change. A review by the Innovation Modeling Comparison Project (Edenhofer et al., 2006a; Grubb et al., 2006; Köhler et al., 2006) offers a thorough description of the most recent attempts to model endogeneity and induced technological innovation – an effort that we will not attempt to reproduce here. Instead, this section briefly discusses three choices that all IAM modellers must make with regard to their representation of abatement technology: how to model increasing returns; how much technological detail to model; and how to model macroeconomic feedback.

Many models, especially general equilibrium models, assume technologies are characterized by decreasing returns to scale (meaning that doubling all inputs yields less than twice as much output), a provision which ensures that there is only one, unique equilibrium result. The assumption of decreasing returns may be realistic for resource-based industries such as agriculture or mining, but it is clearly inappropriate to many new, knowledge-based technologies – and indeed, it is inappropriate to many branches of old as well as new manufacturing, where bigger is better for efficiency, up to a point. Some industries exhibit not only increasing returns in production, but also 'network economies' in consumption – the more people that are using a communications network or a computer operating system, the more valuable that network or operating system is to the next user.

The problem for modelling is that increasing returns and network economies introduce path dependence and multiple equilibria into the set of possible solutions. Small events and early policy choices may decide which of the possible paths or output mixes the model will identify as 'the solution'. An inferior computer operating system, energy technology or other choice may become 'locked in' – the established standard is so widely used and so low-priced because it is produced on such a large scale, that there is no way for individual market choices to lead to a switch to a technologically superior alternative. Modelling increasing returns, path dependence and multiple equilibria can bring IAMs closer to a realistic portrayal of the structure and nature of emissions abatement and economic development options, but at the expense of making models more difficult to construct and model results more difficult to interpret.

Knowledge spillovers are also related to increasing returns. Some of the returns to research and development are externalities, that is, they impact on third parties – other companies, industries or countries. Because of the public goods character of knowledge, its returns cannot be completely appropriated by private investors. Without public incentives for research and development, private firms will tend to under-invest in knowledge, with the result that the total amount of research and development that occurs is less than would be socially optimal. Increasing returns are modelled either as a stock of knowledge capital that becomes an argument in the production function, or as learning curves that lower technological costs as cumulative investments in physical capital or research and development grow.

A second choice that IAM modellers must make is how much technological detail to include. This encompasses not only whether to model increasing returns but also how many regions, industries, fuels, abatement technologies or end uses to include in a model. A more detailed technology sector can improve model accuracy but there are limits to the returns from adding detail – at some point, data requirements, spurious precision and loss of transparency begin to detract from a model's usefulness. On the other hand, a failure to model sufficient technological diversity can skew model results. Abatement options such as renewable energy resources, energy efficiency technologies and behavioural shifts serve to limit abatement costs; models without adequate range of abatement options can exaggerate the cost of abatement and therefore recommend less abatement effort than a more complete model would.

The final modelling choice is how to portray macroeconomic feedback from abatement to economic productivity. A common approach is to treat abatement costs as a pure loss of income, a practice that is challenged by new models of endogenous technological change, but still employed in a number of IAMs, such as DICE-2007 (Nordhaus, 2008). Two concerns seem of particular importance. Modelling abatement costs as a dead-weight loss implies that there are no 'good costs' – that all money spent on abatement is giving up something valuable and thereby diminishing human welfare. But many costs do not fit this pattern: money spent wisely can provide jobs or otherwise raise income, and can build newer, more efficient capital. A related issue is the decision to model abatement costs as losses to income. Abatement costs more closely resemble additions to capital, rather than subtractions from income. (A similar argument can be made regarding many kinds of damage costs: see the earlier section on projecting future damages.)

5.2. Cost minimization models

Many of the IAMs making the most successful inroads into modelling endogenous technological change are cost minimization models. All four of the cost minimization models reviewed in this study – GET-FL (Hedenus et al., 2006), DNE21+ (Sano et al., 2006), MIND (Edenhofer et al., 2006b) and MESSAGE-MACRO (Rao et al., 2006) – include learning curves for specific technologies and a detailed rendering of alternative abatement technologies.

GET-FL, DNE21+, MIND and MESSAGE-MACRO are all energy systems models that include greenhouse gas emissions but not climate change damages. These models include various carbon-free abatement technologies, carbon capture and storage and spillovers within clusters of technologies. GET-FL has learning curves for energy conversion and investment costs. DNE21+ has learning curves for several kinds of renewable energy sources and a capital structure for renewables that is organized in vintages. Both MIND and MESSAGE-MACRO combine an energy system model with a macroeconomic model. MIND has learning curves for renewable energy and resource extraction research; development investments in labour productivity; trade-offs between different types of research and development investment; and a vintaged capital structure for renewables and carbon capture and storage technologies. MESSAGE-MACRO models interdependencies from resource extraction, imports and exports, conversion, transport and distribution to end-use services; declining costs in extraction and production; and learning curves for several energy technologies (Edenhofer et al., 2006a; Köhler et al., 2006).

These energy system models demonstrate the potential for representing induced innovation and endogeneity in technological change. Unfortunately, the very fact of their incredible detail of energy resources, technologies and end uses leads to a separate problem of unmanageably large and effectively opaque results in the most complex IAMs. (For example, the RITE Institute's DNE21+ models historical vintages, eight primary energy sources and four end-use energy sectors, along with five carbon capture and storage methods, several energy conversion technologies and separate learning curves for technologies like wind, photovoltaics and fuel cells.) A model is constructed at the level of detail achievable from present day energy sector data, providing accuracy in the base year calculations. Then the model is extended into the future based on unknowable and untestable projections, running the risk of turning historical accuracy into spurious precision in future forecasts. A high level of specificity about the future of the energy sector cannot be sustained over the number of years or decades necessary to analyse the slow, but inexorable, advance of climate change.

6. Conclusions

The best-known climate-economics models weigh the costs of allowing climate change to continue against the costs of stopping or slowing it, and thus recommend a 'best' course of action: one that, given the assumptions of the model, would cause the least harm. The results of such models are, of course, only as good as their underlying structures and parameter values.

Analysis of climate change, in economics as well as in science, inescapably involves extrapolation into the future. To understand and respond to the expected changes, it is essential to forecast what will happen at greenhouse gas concentrations and temperature levels that are outside the range of human experience, under regimes of technological progress and institutional evolution that have not yet even been envisioned. While some progress has been made toward a consensus about climate science modelling, there is much less agreement about the economic and societal laws and patterns that will govern future development.

IAMs seek to represent both the impacts of changing temperature, sea level and weather on human livelihoods, and the effects of public policy decisions and economic growth on greenhouse gas emissions. IAMs strive not only to predict future economic conditions but also to portray how we value the lives, livelihoods and natural ecosystems of future generations – how human society feels about those who will inherit that future. The results of economic models depend on theories about future economic growth and technological change, and on ethical and political judgments.

Model results are driven by conjectures and assumptions that do not rest on empirical data and often cannot be tested against data until after the fact. To the extent that climate policy relies on the recommendations of IAMs, it is built on what looks like a 'black box' to all but a handful of researchers. Better informed climate policy decisions might be possible if the effects of controversial economic assumptions and judgements were visible and were subjected to sensitivity analyses.

Our review of the literature has led to several concrete lessons for model development:

- Many value-laden technical assumptions are crucial to policy implications and should be visible for debate. Existing models often bury assumptions deep in computer code and parameter choices, discouraging discussion.
- Crucial scientific uncertainties – such as the value of the climate sensitivity parameter and the threshold for irreversible catastrophe – must be addressed in the model structure. Most IAMs use central or average estimates, and thereby ignore catastrophic risk.
- Modelling climate economics requires forecasts of damages at temperatures outside historical experience; there is no reason to assume a simple quadratic (or other low-order polynomial) damage function.
- Today's actions affect the climate and economy of future generations, thus linking current and future welfare. Many models effectively break this link by using high discount rates, inflating the importance of near-term abatement costs while trivializing long-term benefits of mitigation.
- Climate choices occur in an unequal world and inevitably affect opportunities for development. Most regionally disaggregated models use a technical device ('Negishi welfare weights') that freezes the current income distribution, constraining models to ignore questions of interregional equality.
- Measures to induce or accelerate technological change will be crucial for a successful climate policy; a realistic model must allow endogenous technical change and increasing returns. Many IAMs assume decreasing returns and/or exogenous technological progress and treat abatement costs as an unproductive loss of income, not an investment in energy-conserving capital.

Climate-economics models have improved over the years, including expanded treatment of externalities, technological innovation and regional disaggregation. But there is still tremendous scope for further improvement, including more extensive sensitivity analyses and more rigorous examination of risk and uncertainty. Fundamentally subjective judgements, especially those that embody deeply value laden assumptions, can be made more explicit.

What difference would it make to change these features of climate economics modelling? In the absence of a better model, we can only speculate about the results. Our guess is that the modifications we have proposed would make a climate economics model more consistent with the broad outlines of climate science models, portraying the growing seriousness of the problem, the ominous risks of catastrophe and the need for immediate action.

Notes

1. See the Goodess et al. (2003) model classification system, in which AIM and ISGM are both physical impact IAMs.
2. Examples include E3MG and several simulation models.
3. Two climate-economics modelling projects published as special issues of the *Energy Journal* were indispensible in preparing this review. The first was organized by the Stanford Energy Modeling Forum (Weyant and Hill 1999) and the second by the Innovation Modeling Comparison Project (Edenhofer et al., 2006a; Grubb et al., 2006; Köhler et al., 2006). For definitions of IAMs and accounts of their development over time see Goodess et al. (2003), Courtois (2004), Risbey et al. (1996), Rotmans and Dowlatabadi (1998).
4. A sixth category, macroeconomic models, could be added to this list, although the only example of a pure macroeconomic model being used for climate analysis may be the Oxford Global Macroeconomic and Energy Model (Cooper et al., 1999). Publically available documentation for this model is scarce and somewhat cryptic, perhaps because it was developed by a private consulting firm. Macroeconomic models analyse unemployment, financial markets, international capital flows, and monetary policy (or at least some subset of these) (Weyant and Hill, 1999). Three general equilibrium or cost minimization models with macroeconomic features are included in this literature review: G-CUBED/MSG3, MIND and MESSAGE-MACRO.
5. In these models, consumption's returns to welfare are always positive but diminish as we grow wealthier. Formally, the first derivative of welfare is always positive and the second is always negative. A popular, though not universal, choice defines individual welfare, arbitrarily, as the logarithm of per capita consumption or income.
6. For a critique of IAMs that focuses on the shortcomings of welfare-optimization models, see Courtois (2004).
7. See also DeCanio (2003).
8. On transparency of value-laden assumptions in IAMs see Schneider (1997), Morgan and Dowlatabadi (1996), Risbey et al. (1996), DeCanio (2003), Rotmans and Dowlatabadi (1998) and Parson (1996). On transparency of IAMs code and software see Ha-Duong (2001). For a discussion of how overly complex models can falsely convey model accuracy see Rotmans and Dowlatabadi (1998).
9. Several discussions of how best to assess IAMs exist in the literature, including Morgan and Dowlatabadi (1996), Risbey et al. (1996), and Rotmans and Dowlatabadi (1998).
10. Numerous reviews of IAMs critique their oversimplification of the physical climate model, the lack of clear standards in interdisciplinary work and the degree to which they lag behind current scientific findings. See Courtois (2004), Hall and Behl (2006), Parson (1996), Risbey et al. (1996) and Rotmans and Dowlatabadi (1998).
11. Morgan and Dowlatabadi (1996) stress the importance of portraying uncertainty in their 'hallmarks of good IAMs'. For other reviews and general discussions of uncertainty in IAMs see Scott et al. (1999) Morgan et al. (1999), Warren et al. (2006), Rotmans and Dowlatabadi (1998) and Heal and Kristrom (2002).
12. In more recent work, Weitzman has suggested that climate science implies even greater risks at the 95th–99th percentile (Weitzman, 2008). Of course, his argument does not depend on an exact estimate of these risks; the point is that accuracy is unattainable and the risks do not have an obvious upper bound, yet effective policy responses must be informed by those low-probability extreme events.
13. Latin Hypercube sampling, a technical procedure widely used in Monte Carlo analyses, ensures that the selected sets of parameters are equally likely to come from all regions of the relevant parameter space.
14. If the uncertain parameters were all truly independent of each other, such combinations of multiple worst case values would be extraordinarily unlikely. The danger is that the uncertain parameters, about which our knowledge is limited, may not be independent.
15. For discussions of the problems arising from long time scales in IAMs see Parson (1996) and Morgan et al. (1999).
16. For a review of damages functions in DICE, RICE, FUND, MERGE, and PAGE see Warren et al. (2006).
17. DICE-2007 actually uses a slightly more complicated equation which is equivalent to our equation (1), with the exponent $b = 2$, for small damages.
18. See Ackerman et al. (2008) for a more detailed critique of the DICE-2007 damage function.

19. Risbey et al. (1996) refer to this practice as the 'wholesale uncritical adoption of archetype models'.
20. Diminishing marginal utility of income is the ubiquitous assumption in neo-classical economics that each new dollar of income brings a little less satisfaction than the last dollar.
21. The terminology of descriptive and prescriptive approaches was introduced and explained in Arrow et al. (1996).
22. If the elasticity of the marginal utility of consumption is a constant η, as in equation (2), and per capita consumption is c, then utility $= c^{(1-\eta)}/(1-\eta)$, except when $\eta = 1$, when utility $= \ln c$. See the Stern Review (Stern, 2006), technical annex to Chapter 2 on discounting, or other standard works on the subject.
23. The MIND model (Edenhofer et al., 2006b), which combines cost minimization with welfare maximization, uses a pure rate of time preference of 1% and a total discount rate of 5%.
24. Many models make the implicit assumption that resource availability is infinite and do not explicitly consider resource limitations or resource use efficiency.
25. For examples of how this procedure is discussed in the climate-economics literature see Kypreos (2005: 2723), Peck and Teisberg (1998: 3–4) and Yang and Nordhaus (2006: 731–738).
26. The marginal product of capital is the increase in output resulting from the last unit of capital added to the economy. It tends to be higher when capital is scarce, that is, in poorer regions.
27. For an example of the Negishi weights methodology see Yang and Nordhaus (2006) or Manne and Richels (2004).
28. Earlier versions of PAGE2002, in fact, applied equity weights that boost the relative importance of outcomes in developing countries; the Stern Review modeling effort dropped the equity weights in favour of a more explicit discussion of regional inequality (Chris Hope, personal communication, 2008).

References

Ackerman, F., 2002. Still dead after all these years: interpreting the failure of general equilibrium theory. *Journal of Economic Methodology* 9(2). 119–139.

Ackerman, F., Stanton, E. A. and Bueno, R., 2008. *Fat Tails, Exponents and Extreme Uncertainty: Simulating Catastrophe in DICE*. Working Paper WP-0901, Stockholm Environment Institute, US Center.

Arrow, K. J., Cline, W. R., Maler, K.-G., Munasinghe, M., Squitieri, R. and Stiglitz, J. E., 1996. Intertemporal equity, discounting, and economic efficiency. *Climate Change 1995 – Economic and Social Dimensions of Climate Change. Contribution of Working Group III to the Second Assessment Report of the IPCC*. H. L. James, P. Bruce and Erik F. Haites (eds). IPCC and Cambridge University Press, New York. 125–144.

Babiker, M., Gurgel, A., Paltsev, S. and Reilly, J., 2008. *A Forward Looking Version of the MIT Emissions Prediction and Policy Analysis (EPPA) Model*. MIT Joint Program on the Science and Policy of Global Change, Report No. 161.

Barker, T., Pan, H., Köhler, J., Warren, R. and Winne, S., 2006. Decarbonizing the global economy with induced technological change: scenarios to 2100 using E3MG. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 241–258.

Bernstein, P. M., Montgomery, W. D., Rutherford, T. F. and Yang, G.-F., 1999. Effects of restrictions on international permit trading: the MS-MRT model. *Energy Journal*, Special Issue on the Costs of the Kyoto Protocol: A Multi-Model Evalution. 221–256.

Bosetti, V., Carraro, C. and Galeotti, M., 2006. The dynamics of carbon and energy intensity in a model of endogenous technical change. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 191–206.

Clarke, L., Wise, M., Kim, S., Smith, S., Lurz, J., Edmonds, J. and Pitcher, H., 2007. *Model Documentation for the MiniCAM Climate Change Science Program Stabilization Scenarios: CCSP Product 2.1a*. Report for the US Department of Energy under Contract DE-AC05-76RL01830, Pacific Northwest National Laboratory, PNNL-16735.

Cooper, A., Livermore, S., Rossi, V., Wilson, A. and Walker, J., 1999. The economic implications of reducing carbon emissions: a cross-country quantitative investigation using the Oxford global macroeconomic and energy model. *Energy Journal*, Special Issue on The Costs of the Kyoto Protocal: A Multi-Model Evaluation. 335–365.

Courtois, P. 2004. The status of integrated assessment in climatic policy making: an overview of inconsistencies underlying response functions. *Environmental Science & Policy*, 7(1). 69–75.

Crassous, R., Hourcade, J.-C. and Sassi, O., 2006. Endogenous structural change and climate targets modeling experiments with Imaclim-R. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 259–276.

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

Dietz, S., Hope, C., Stern, N. and Zenghelis, D., 2007. Reflections on the Stern Review (1): a robust case for strong action to reduce the risks of climate change. *World Economics* 8(1). 121–168.

Dowlatabadi, H., 1998. Sensitivity of climate change mitigation estimates to assumptions about technical change. *Energy Economics* 20. 473–493.

Edenhofer, O., Lessmann, K., Kemfert, C., Grubb, M. and Köhler, J., 2006a. Induced technological change: exploring its implications for the economics of atmospheric stabilization: synthesis report from the innovation modeling comparison project. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 57–108.

Edenhofer, O., Lessmann, K. and Bauer, N., 2006. Mitigation strategies and costs of climate protection: the effects of ETC in the hybrid model MIND. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 207–222.

Edmonds, J., Pitcher, H. and Sands, R., 2004. Second Generation Model 2004: An Overview. *United States Environmental Protection Agency under contracts AGRDW89939464-01 and AGRDW89939645-01*. Joint Global Change Research Institute and Pacific Northwest National Laboratory.

Gerlagh, R., 2006. ITC in a global growth-climate model with CCS: the value of induced technical change for climate stabilization. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 223–240.

Gerlagh, R., 2008. A climate-change policy induced shift from innovations in carbon-energy production to carbon-energy savings. *Energy Economics*, 30. 425–448.

Goodess, C. M., Hanson, C., Hulme, M. and Osborn, T. J., 2003. Representing climate and extreme weather events in integrated assessment models: a review of existing methods and options for development. *Integrated Assessment Journal*, 4(3). 145–171.

Grubb, M., Carraro, C. and Schellnhuber, J., 2006. Technological change for atmospheric stabilization: introductory overview to the innovation modeling comparison project. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 1–16.

Ha-Duong, M., 2001. Transparency and control in engineering integrated assessment models. *Integrated Assessment*, 2(4). 209–218.

Hall, D. C. and Behl, R. J., 2006. Integrating economic analysis and the science of climate instability. *Ecological Economics*, 57(3). 442–465.

Heal, G. and Kristrom, B., 2002. Uncertainty and climate change. *Environmental and Resource Economics* 22. 3–39.

Hedenus, F., Azar, C. and Lindgren, K., 2006. Induced technological change in a limited foresight optimization model. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 109–122.

Hope, C., 2006. The marginal impact of CO2 from PAGE2002: an integrated assessment model incorporating the IPCC's five reasons for concern. *Integrated Assessment Journal*, 6(1). 19–56.

Intergovernmental Planet on Climate Change, 2007. *Climate Change 2007 – IPCC Fourth Assessment Report*. Cambridge, UK, Cambridge University Press.

Jorgenson, D. W., Goettle, R. J., Hurd, B. H., Smith, L. B., Chesnut, L. G. and Mills, D. M., 2004. US market consequences of global climate change. Pew Center on Global Climate Change, Arlington, VA.

Kainuma, M., Matsuoka, Y. and Morita, T., 1999. Analysis of post-kyoto scenarios: the Asian-Pacific integrated model. *Energy Journal*, Special Issue on the Costs of the Kyoto Protocol: A Multi-Model Evalution. 207–220.

Keller, K., Yang, Z., Hall, M. and Bradford, D. F., 2003. Carbon dioxide sequestration: when and how much? Working Paper 94 Center for Economic Policy Studies, Princeton University.

Kemfert, C., 2001. Economy-energy-climate interaction: the model wiagem. NOTA DI LAVORO 71.2001 Fondazione Eni Enrico Mattei.

Köhler, J., Grubb, M., Popp, D. and Edenhofer, O., 2006. The transition to endogenous technical change in climate-economy models: a technical overview to the innovation modeling comparison project. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 17–56.

Kurosawa, A., 2004. Carbon concentration target and technological choice. *Energy Economics*, 26. 675–684.

Kurosawa, A., Yagita, H., Zhou, W., Tokimatsu, K. and Yanagisawa, Y., 1999. Analysis of carbon emissions stabilization targets and adaptation by integrated assessment model. *Energy Journal*, Kyoto Special Issue. 157–175.

Kypreos, S., 2005. Modeling experience curves in MERGE (model for evaluating regional and global effects). *Energy*, 30(14). 2721–2737.

Kypreos, S., 2008. Stabilizing global temperature change below thresholds: Monte Carlo analyses with MERGE. *Journal of Computational Management Science*, 5(1–2). 141–170.

Lejour, A., Veenendaal, P., Verweij, G. and Leeuwen, N. V., 2004. *WorldScan: A Model for International Economic Policy Analysis*. CPB Netherlands Bureau for Economic Policy Analysis, No. 111.

Manne, A. S., 1999. Greenhouse gas abatement: toward pareto-optimality in integrated assessment. *Education in a Research University*. K. J. Arrow, R. W. Cottle, B. C. EavesI. Olkin (eds). Springer, Netherlands.

Manne, A. S. and Richels, R. G., 2004. *MERGE: An Integrated Assessment Model for Global Climate Change*. http://www.standford.edu/group/MERGE/.

Masui, T., Hanaoka, T., Hikita, S. and Kainuma, M., 2006. Assessment of CO2 reductions and economic impacts considering energy-saving investments. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 175–190.

McKibbin, W. J. and Wilcoxen, P. J., 1999. The theoretical and empirical structure of the G-cubed model. *Economic Modelling*, 16. 123–148.

Mendelsohn, R. and Williams, L., 2004. Comparing forecasts of the global impacts of climate change. *Mitigation and Adaptation Strategies for Global Change*, 9(4). 315–333.

Morgan, M. G. and Dowlatabadi, H., 1996. Learning from integrated assessment of climate change. *Climatic Change*, 34. 337–368.

Morgan, M. G., Kandlikar, M., Risbey, J. and Dowlatabadi, H., 1999. Why conventional tools for policy analysis are often inadequate for problems of global change. *Climatic Change*, 41(3). 271–281.

Negishi, T., 1972. *General Equilibrium Theory and International Trade*. North-Holland Publishing Company, Amsterdam, London.

Nordhaus, W. D., 2007a. *Accompanying Notes and Documentation on Development of DICE-2007 Model: Notes on DICE-2007.v8 of September 28, 2007*. http://nordhaus.econ.yde.edu/Accom_Notes_100507.pdf.

Nordhaus, W. D., 2007b. *The Challenge of Global Warming: Economic Models and Environmental Policy*. http://nordhaus.econ.yale.edu/DICE2007.htm.

Nordhaus, W. D., 2008. *A Question of Balance: Economic Modeling of Global Warming*. Yale University Press, New Haven.

Nordhaus, W. D. and Boyer, J., 2000. *Warming the World: Economic Models of Global Warming*. MIT Press, Cambridge, Massachusetts.

Nordhaus, W. D. and Popp, D., 1997. What is the value of scientific knowledge? An application to global warming using the PRICE model. *Energy Journal* 18(1). 1–45.

Nordhaus, W. D. and Yang, Z., 1996. A regional dynamic general-equilibrium model of alternative climate-change strategies. *American Economic Review*, 86(4). 741–765.

Pant, H. M., 2007. *GTEM Draft: Global Trade and Environmental Model*. Canberra Australian Bureau of Agricultural and Resource Economics.

Parson, E. A., 1996. Three dilemmas in the integrated assessment of climate change. *Climatic Change*, 34(3). 315–326.

Peck, S. C. and Teisberg, T. J., 1995. Optimal CO2 control policy with stochastic losses from temperature rise. *Climatic Change*, 31. 19–34.

Peck, S. C. and Teisberg, T. J., 1998. *CO2 Concentration Limits, The Costs and Benefits of Control, and the Potential for International Agreement. Fondazione Eni Enrico Mattei Working Paper 6.98*. DOI:10.2139/ssrn.122029.

Peck, S. C. and Teisberg, T. J., 1999. CO2 emissions control agreements: incentives for regional participation. *Energy Journal*, Special Issue on the Costs of the Kyoto Protocol: A Multi-Model Evaluation. 367–390.

Popp, D., 2006. Comparison of climate policies in the ENTICE-BR model. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 163–174.

Ramsey, F. P., 1928. A mathematical theory of saving. *The Economic Journal*, 138(152). 543–559.

Rao, S., Keppo, I. and Riahi, K., 2006. Importance of technological change and spillovers in long-term climate policy. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 123–139.

Risbey, J., Kandlikar, M. and Patwardhan, A., 1996. Assessing integrated assessments. *Climatic Change*, 34(3). 369–395.

Rotmans, J. and Dowlatabadi, H., 1998. Integrated assessment modeling. *Tools for Policy Analysis*, S. Rayner and E. L. Malone (eds). Battelle Press, Columbus, OH. Vol. 3.

Sano, F., Akimoto, K., Homma, T. and Tomoda, T., 2006. Analysis of technological portfolios for CO2 stabilizations and effects of technological changes. *Energy Journal*, Special Issue on Endogenous Technological Change and the Economics of Atmospheric Stabilisation. 141–161.

Schneider, S. H., 1997. Integrated assessment modeling of global climate change: transparent rational tool for policy making or opaque screen hiding value-laden assumptions? *Environmental Modeling and Assessment*, 2(4). 229–249.

Scott, M. J., Sands, R. D., Edmonds, J., Liebetrau, A. M. and Engel, D. W., 1999. Uncertainty in integrated assessment models: modeling with MiniCAM 1.0. *Energy Policy*, 27. 855–879.

Stern, N., 2006. *The Stern Review: The Economics of Climate Change*. HM Treasury, London.

Tol, R. S. J., 1999. Kyoto, efficiency, and cost effectiveness: applications of FUND. *Energy Journal*, Special Issue on the Costs of the Kyoto Protocol: A Multi-Model Evaluation. 131–156.

Warren, R., Hope, C., Mastrandrea, M., Tol, R., Adger, N. and Lorenzoni, I., 2006. Spotlighting Impacts Functions in Integrated Assessment. Tyndall Centre, Norwich, UK.

Webster, M., Tatang, M. A. and McRae, G. J., 1996. *Report #4: Application of the Probabilistic Collocation Method for an Uncertainty Analysis of a Simple Ocean Model*. Massachusetts Institute of Technology, Joint Program on the Science and Policy of Global Change.

Weitzman, M. L., 2007. A review of the Stern Review on the economics of climate change. *Journal of Economic Literature* 45(3). 703–724.

Weitzman, M. L., 2008. On modeling and interpreting the economics of catastrophic climate change (June 6, 2008 version). http://www.economics.harvard.edu/faculty/weitzman/files/modeling.pdf.

Weyant, J. P. and Hill, J. N., 1999. Introduction and overview. *Energy Journal*, Special Issue on the Costs of the Kyoto Protocol: A Multi-Model Evaluation. vii–xliv.

Yang, Z. and Nordhaus, W. D., 2006. Magnitude and direction of technological transfers for mitigating GHG emissions. *Energy Economics* 28. 730–741.