Global Environmental Change 24 (2014) 123–131

Contents lists available at ScienceDirect

Global Environmental Change

journal homepage: www.elsevier.com/locate/gloenvcha

Risk mitigation and the social cost of carbon

Richard B. Howarth [a,*], Michael D. Gerst [b,1], Mark E. Borsuk [b,2]

[a] Environmental Studies Program, Dartmouth College, Hanover, NH 03755, USA
[b] Thayer School of Engineering, Dartmouth College, Hanover, NH 03755, USA

ARTICLE INFO

Article history:
Received 14 February 2013
Received in revised form 8 November 2013
Accepted 11 November 2013

Keywords:
Climate stabilization
Integrated assessment
Risk and uncertainty
Social cost of carbon

ABSTRACT

The social cost of carbon – i.e., the marginal present-value cost imposed by greenhouse gas emissions – is determined by a complex interaction between factual assumptions, modeling methods, and value judgments. Among the most crucial factors is society's willingness to tolerate potentially catastrophic environmental risks. To explore this issue, the present analysis employs a stochastic climate–economy model that accounts for uncertainties in baseline economic growth, baseline emissions, greenhouse gas mitigation costs, carbon cycling, climate sensitivity, and climate change damages. In this model, preferences are specified to reflect the high degree of risk aversion revealed by private investment decisions, signaled by the large observed gap between the average rates of return paid by safe and risky financial instruments. In contrast, most climate–economy models assume much lower risk aversion. Given high risk aversion, the analysis finds that investment in climate stabilization yields especially large net benefits by forestalling low-probability threats to long-run human well-being. Accordingly, the social cost of carbon attains the markedly high value of $25,700 per metric ton of carbon dioxide in a baseline scenario in which emissions are unregulated. This value falls to just $4 per ton as the stringency of control measures is successively increased. These results cast doubt on the idea that the social cost of carbon takes on a uniquely defined, objective value that is independent of policy decisions. This does not, however, rule out the use of carbon prices to achieve the benefits of climate stabilization using least-cost mitigation measures.

1. Introduction

The social cost of carbon – defined as the marginal present-value cost imposed by greenhouse gas emissions – has emerged as a central concept in the economics of climate change (Tol, 2011). In 2002, for example, the United Kingdom adopted an official social cost of £19 per metric ton of carbon dioxide (or $29 per ton) for use in policy evaluation (see Department for Environment Food and Rural Affairs, 2002; Pearce, 2003). In the United States, carbon pricing is now required under the procedures of the Office of Management and Budget, which mandate the use of cost-benefit analysis to review all significant new and revised federal regulations, even when statutory requirements explicitly rule out a balancing of costs and benefits in the promulgation of environmental standards (Clinton, 1993; Hahn and Sunstein, 2002). Applications of this approach have established that accounting for net carbon emissions can have non-trivial impacts on the desirability of policy options (Kopp and Mignone, 2012). This is true, for example, of the U.S. Corporate Average Fuel Economy standards, where the level of fuel economy that is judged to be economically efficient is sensitive to the monetary value assigned to reductions in greenhouse gas emissions (Masur and Posner, 2011).

In 2009, the Obama Administration convened an Interagency Working Group with representation from the Environmental Protection Agency and five cabinet level departments (Agriculture, Commerce, Energy, Transportation, and Treasury) to survey the literature and assign a range of quantitative values to the social cost of carbon for use in official policy analysis. In the pursuit of this task, the Working Group employed three major models of the interplay between climate change and the global economy: Nordhaus' (2008) "Dynamic Integrated Climate Economy" (DICE) model; Hope's (2008) "Policy Analysis of the Greenhouse Effect" (PAGE) model; and Anthoff and Tol's (2010) "Climate Framework for Uncertainty, Negotiation, and Distribution" (FUND) model (see also Tol, 1997). Although these models differ in various details, they adopt broadly similar assumptions regarding the costs of greenhouse gas emissions reductions, the economic impacts of climate change, and future trends in technology, population, and

* Corresponding author. Tel.: +1 603 646 2752; fax: +1 603 646 1682.
E-mail addresses: RBHowarth@Dartmouth.edu (R.B. Howarth), Michael.Gerst@Dartmouth.edu (M.D. Gerst), Mark.Borsuk@Dartmouth.edu (M.E. Borsuk).
[1] Tel.: +1 704 516 5067; fax: +1 603 646 2277.
[2] Tel.: +1 603 646 9944; fax: +1 603 646 2277.

0959-3780/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.gloenvcha.2013.11.012

economic growth. (FUND, however, is relatively optimistic about climate impacts and adaptation, especially for small changes in mean global temperature.)

One point of difficulty for the Interagency Working Group was choosing the rate at which to discount future costs and benefits. Nordhaus and others have asserted that investments in greenhouse gas mitigation are warranted if and only if they provide returns at least as high as those available on financial markets, or approximately 6% per year (Nordhaus, 2008). In contrast, authors including Cline (1992) and Stern (2007) have argued that annual discount rates on the order of 1–2% are justified if one accepts the moral premise that equal weight should be attached to the welfare of present and future generations. Based on its review of the literature, the Working Group decided to consider discount rates of 2.5%, 3%, and 5% for all three models. With a discount rate of 5% per year, the Working Group concluded that the social cost of carbon attains a value in the year 2010 of $4.7 per metric ton of carbon dioxide, or just 4 cents per gallon of gasoline equivalent. With a 2.5% discount rate, the Working Group estimated a social cost of carbon of $35.1 per ton for the year 2010. Thus the use of low discount rates favors more aggressive steps to stabilize climate (Stern, 2007).

The reception of the Interagency Working Group report on the social cost of carbon has been mixed. On the one hand, it is clearly very important to assign an accounting price to changes in greenhouse gas emissions for use in regulatory impact analysis (Rose, 2010). An appropriately chosen carbon price can guide decision-makers to the adoption of behaviors and technologies that achieve society's environmental goals at the least economic cost. Pragmatically, the Working Group report provides a framework that federal agencies can utilize to pursue this objective.

On the other hand, critics have argued that the three models considered by the Interagency Working Group are based in part on optimistic assumptions concerning the projected economic impacts of climate change coupled with an incomplete analysis of risk (Ackerman et al., 2009; Pindyck, 2013; Stern, 2013). The DICE model, for example, assumes that a 3 °C increase in mean global temperature would lead to a 2.5% reduction in economic output. PAGE and FUND assume even lower damages. This contrasts with Hansen et al.'s (2008) warning that increases in mean global temperature exceeding 1–2 °C could potentially trigger positive feedback processes related to ice sheet collapses and the destabilization of global ecosystems that would impose truly catastrophic costs (see also Lenton et al., 2008).

To address uncertainty, the Interagency Working Group adopted Roe and Baker's (2007) fat-tailed distribution on climate sensitivity – i.e., the change in mean global temperature caused by a doubling of greenhouse gas concentrations – which implies a 20% chance of exceeding 5.0 °C. Monte Carlo simulations were then used to estimate a 95th percentile estimate for the social cost of carbon given a 3% discount rate. The resulting estimate of $64.9 per metric ton of carbon dioxide for the year 2010 is in one sense surprising. Ackerman and Stanton (2012), for example, found that assigning plausible values to uncertain parameters can result in a carbon price as high as $900 per ton of carbon dioxide. Anthoff et al. (2009) found that even higher values can arise in a sensitivity analysis involving low time preference and high risk aversion, though their central estimate was $16 per ton based on their interpretation of decision-makers' revealed preferences in the absence of equity weighting. These points are linked to Neumayer's (2007) concern that the current generation of integrated assessment models does not fully account for the potentially "irreversible and non-substitutable damage" that climate change will inflict on the stability and functioning of ecosystems and the role of natural capital in supporting human activity.

In the present paper, we develop Kousky et al.'s (2011) argument that appropriately accounting for the role of risk mitigation might substantially alter the numerical value assigned to the social cost of carbon. Following Weitzman (2009), we work with a formal model of decision-making under uncertainty that allows for major risks of the type described by Hansen et al. and Roe and Baker. Using a theoretical model, Weitzman concluded that aggressive climate change policies might generate highly valuable (at face value potentially infinite) net benefits by reducing the statistically low probability that unmitigated climate change would lead to future economic collapse. In previous work (Gerst et al., 2013), we confirmed this finding in a plausibly specified numerical model in which preferences regarding time and risk were inferred from market data on consumption growth and the rates of return paid by safe and risky investments using methods from the macrofinance literature (Lucas, 1978; Mehra and Prescott, 1985; Barro, 2006). Iverson and Perrings (2011) provide a related analysis based on an application of the asymmetric minimax regret criterion as a framework for characterizing rational decisions under strong uncertainty. In a similar vein, McInerney et al. (2012) describe how an array of decision-theoretic approaches can be applied to evaluate climate change policies in a modified version of DICE.

Here, we employ the Gerst et al. model to produce a seemingly paradoxical result. On the one hand, we find that deep cuts in greenhouse gas emissions can produce very high net social benefits. On the other hand, once an aggressive control path is initiated, the *marginal* benefit of further emissions reductions is quite low. We see this result as consistent with the well-known "diamond–water" paradox, in which actions that are essential to sustaining human welfare have high total net benefits yet low marginal benefits given appropriate levels of provisioning (see Farber et al., 2002). In our model, this occurs when emissions cuts are sufficient to reduce the relatively low probability of catastrophic climate impacts to essentially zero.

These results contrast strongly with the Interagency Working Group's (2010) finding that the social cost of carbon is relatively independent of the stringency of emissions abatement. Such independence can occur in a deterministic model like DICE (Nordhaus, 2008), in which equilibrium temperature is logarithmic with respect to greenhouse gas concentrations and climate change damages are quadratic with respect to temperature. This implies a nearly linear relationship between temperature and damages and, hence, a marginal cost of greenhouse gas emissions that is independent of the state of the environment (Hope, 2006). This, however, appears to be an idiosyncratic condition rather than a general phenomenon, especially when the complex dynamics of risk mitigation are considered.

Our analysis suggests that risk-averse decision-makers attach especially high value to the early elimination of catastrophic risks and that the level of total net benefits provided becomes nearly invariant to specific policy scenarios under emissions control rates of 40% or more by the year 2050. This implies that the main risks to welfare are from failing to stabilize climate, not from cutting emissions by too much, too soon. Thus, balancing the marginal costs and benefits of emissions controls may be less important than attaining the overall benefits of climate stabilization. Pragmatically, this may favor an approach to carbon pricing aimed at the cost-effective achievement of policy-specified emissions targets, rather than a focus on the 'correct' social cost of carbon, especially given the uncertainties associated with integrated assessment models and the strong role of moral values in climate governance (see Howarth, 2011; Dietz, 2012).

2. The model

For the purposes of analysis, we work with Gerst et al.'s (2013) stochastic integrated assessment model of climate–economy interactions, which adopts the mitigation cost, damage cost, carbon cycle, and climate system modules of DICE (Nordhaus, 2008). To allow for tractable analysis of low probability risks, Gerst et al. base the economic module on the well-known Lucas–Mehra–Prescott model (Lucas, 1978; Mehra and Prescott, 1985), which plays a central role in understanding the coupled dynamics of asset valuation and economic growth under conditions of uncertainty. Although we do not endogenously model decisions concerning capital accumulation, the Lucas–Mehra–Prescott model may be understood as a reduced-form representation of a suitably specified stochastic growth model in which a fixed savings rate emerges as socially optimal (see Barro, 2006). This type of analysis abstracts away from the feedbacks that arise when climate damages lead to reductions in economic output that in turn influence capital investment, mirroring the methods employed by the PAGE and FUND models, which treat baseline economic growth as exogenously specified.

Potential economic output in the model is allocated to four components: capital costs, greenhouse gas emissions abatement costs, climate-change damages, and final consumption. For the sake of tractability, we assume that a fixed share of potential output ($\beta = 0.22$) is allocated to the cost of maintaining the productivity of a fixed stock of non-reproducible capital assets. After adjusting for mitigation costs (m_t) and damage costs (d_t) as shares of economic output, this results in the per capita consumption level:

$$c_t = (1 - \beta - m_t - d_t)\hat{y}_t. \tag{1}$$

Here, $\hat{y}_t$ is the level of potential economic output per capita, which grows at a rate g_t that follows a random-walk statistical process based on detailed historical data (Barro, 2006). The distribution of g_t has non-trivial skewness (Fig. 1), an empirical characteristic important in establishing our preference calibration with respect to market returns (Ding et al., 2012).

Human population grows according to a probabilistic specification based on the long-run scenarios developed by Lutz et al. (2008). The main uncertainty in this specification is the year in which population peaks, represented by a beta distribution with a mean of 2067 and a standard deviation of 18 (Gerst et al., 2010).

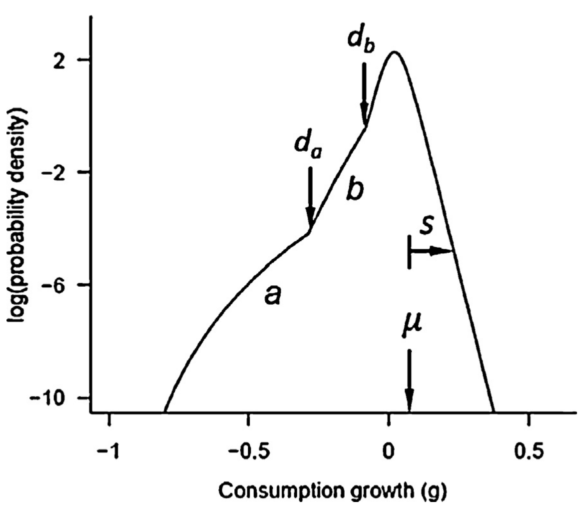

Fig. 1. Distribution of the potential growth rate for per capita economic output. The central portion of the probability density function is described by a logistic distribution while the lower tails are represented by a double power law distribution. The transitions among functions, which occur at $d_b = -0.0805$ and $d_a = -0.228$, are smooth, and all parameters are estimated simultaneously so that the function integrates to one. The shape of the inner and outer power law distributions are described by exponents $b = 15.14$ and $a = 4.96$, respectively. The logistic distribution is described by the mean 0.0212 and the standard deviation 0.0261. The data and estimation procedures are described in Ding et al. (2012).

The costs of climate change mitigation from reducing baseline emissions by the fraction μ_t are modeled as a percentage reduction in gross economic output:

$$m_t = B_t\sigma_t\left(\frac{1}{\omega}\right)\mu_t^{\omega} \tag{2}$$

Following DICE, the coefficient $\omega = 2.8$ represents the curvature of the abatement cost function. Uncertainty in mitigation costs is introduced through parameters that represent the cost of a carbon-free backstop technology (B_t; \$ per metric ton carbon) and changes in baseline emissions intensity (σ_t). Backstop costs decline according to an exogenous, deterministic trajectory in which the initial value (B_0) is normally distributed with a mean of 1170 and a standard deviation of 468 (Nordhaus, 2008). Baseline emissions per unit of output fall due to technological change and a transition toward higher energy efficiency and low-carbon energy carriers. Uncertainty in the decline of emissions intensity is introduced by the change in the rate of decarbonization per decade, represented as a normal distribution with a mean of −0.031 and a standard deviation 0.122 (Gerst et al., 2010).

The carbon cycle is represented by a three-box model representing carbon exchange among the atmosphere, the upper ocean, and the lower ocean. Uncertainty concerning the exchange coefficients, which regulate the movement of carbon among reservoirs, is modeled by assuming a normal distribution for atmospheric-upper ocean exchange with a mean of 0.189 and a standard deviation of 0.017 (Nordhaus, 2008). Other exchange coefficients are adjusted algebraically to preserve mass balance.

The climate system is modeled by two heat reservoirs, the atmosphere and the lower ocean. A key uncertainty in this representation is equilibrium climate sensitivity (S), which defines the equilibrium temperature change resulting from a doubling of atmospheric greenhouse gas concentrations. In the base version of the DICE model, climate sensitivity is normally distributed with a mean of 3 °C and a standard deviation of 1.1 °C. Importantly, this symmetric representation is strongly inconsistent with the recent scientific literature (Tomassini et al., 2007).

To address this issue, the present analysis focuses on two asymmetric distributions for S. One follows Roe and Baker's (2007) assumption that climate sensitivity is characterized by a fat-tailed probability distribution with a 20% chance of exceeding 5.0 °C and a 3.5% probability of exceeding 10 °C. The other utilizes a thin-tailed gamma distribution, which is calibrated to have a median value of 3 °C to match DICE's point estimate and a 23% probability of exceeding 4.5 °C (Zickfeld et al., 2010). While it is clear that other climate system uncertainties, such as ocean heat uptake, are worthy of exploration and quantification (Baker and Roe, 2009), here we focus on climate sensitivity alone to allow for easy comparison with other integrated assessment models.

As in DICE, climate damage costs are modeled using the equation:

$$d_t = \left(\frac{\pi T_t^2}{1 + \pi T_t^2}\right). \tag{3}$$

Here, d_t represents the percentage reduction in economic output caused by increases in mean global temperature. In this equation, T_t is the temperature change above pre-industrial levels (°C), while π is a normally distributed coefficient with a mean of 0.00285 and a standard deviation of 0.0013 (Nordhaus, 2008). Note that while the exponent in Eq. (3) is also considerably uncertain (see Gerst et al., 2010), here we treat it as fixed to aid in comparison with standard DICE model runs. This suggests that our model is likely to understate true climatic risks; as noted in the introduction, authors such as Hansen et al. (2008) believe that climate damages may be highly nonlinear given temperature increases in excess of 2 °C. A further analysis of this issue is provided by Kopp et al. (2012), who find that uncertainty about the damage function given large temperature increases can have significant impacts on the social cost of carbon, especially when decision-makers are assumed to have high risk aversion.

A crucial aspect of our analysis concerns the representation of social preferences concerning time and risk. To address this issue, we assume that social welfare (W) may be represented by the function:

$$W = E\left[\sum_{t=0}^{\infty}(1+\rho)^{-t}N_t\left(\frac{c_t^{1-\gamma}}{1-\gamma}-k\right)\right]. \quad (4)$$

In this expression, $E[\cdot]$ is the expectations operator, which aggregates over uncertain future outcomes based on their relative probabilities; ρ is the pure rate of time preference, which measures the relative weight attached to present and future well-being; N_t is the population at date t; c_t is per capita consumption; γ is the coefficient of relative risk aversion; and k is a scaling constant.

For our main calibration of the coefficient of relative risk aversion, we follow the results established by Ding et al.'s (2012) multi-country revealed preference analysis, which concludes a value of $\gamma = 5.6$ is most consistent with people's observed behavior on financial markets. Specifically, this study builds on Barro's (2006) analysis of the equity premium paradox, which explains the high average returns on risky assets (r_r, approximately 6% per year) and the low observed returns on safe assets (r_s, approximately 1% per year) based on a model that accounts for people's attempts to insure themselves against the effects of low-probability, severe economic downturns. In terms of the data, the consumption growth rate (g) exhibits a fat-tailed statistical distribution with major downside risks (Fig. 1). Following Barro, Ding et al. account for this fact in a model that assumes sequentially rational investment behavior under uncertainty. This gives rise to a revised, expectational version of Ramsey's Rule known as the Consumption Capital Asset Pricing Model (see Romer, 2006, pp. 366–370):

$$(1+\rho) = E[(1+g)^{-\gamma}(1+r_r)] = E[(1+g)^{-\gamma}(1+r_s)] \quad (5)$$

The right-hand part of this equation may be understood as an optimal portfolio condition that uniquely determines the coefficient of relative risk aversion in a manner that is independent from the pure rate of time preference.

For the sake of easy comparison with other major integrated assessment models (e.g., Nordhaus, 2008; Anthoff et al., 2009), we choose a rate of pure time preference of 1.5% per year. Our coefficient of relative risk aversion, however, is significantly higher than the values between 1 and 2 that are typically applied in the integrated assessment literature. One key point is that assuming a value on the order of 1–2 implies that safe and risky assets should pay nearly identical average returns – a prediction that is strongly inconsistent with observed historical data. In addition, it is important to note that our estimate of this parameter is generally consistent with the results of macroeconometric studies of intertemporal substitution (Hall, 1988) and with stated preference studies concerning people's attitudes toward long-run health and environmental risks (Barsky et al., 1997; Atkinson et al., 2009). These independent literatures jointly support values of γ falling between 4 and 8. The point is not that using this estimate in a deterministic model would yield a realistic rate of economic growth. Instead, risk and uncertainty appear to play central roles in decisions concerning investment and portfolio allocation that are not captured by deterministic models. Given the importance of γ in the context of our model, we conduct sensitivity analyses that reduce the level of risk aversion to values of 2.0 (as in DICE) and 3.5. As we shall see, this change has a substantial bearing on our results.

An anonymous referee noted that paradoxes can arise in a social welfare function of the form stated by Eq. (4) when the level of population is stochastic. This general issue corresponds to the distinction between average and total utilitarianism that has been widely studied in consequentialist ethics (see Sinnott-Armstrong, 2012). In Gerst et al. (2013), we found that switching off uncertainty concerning future population growth does not significantly affect the net welfare gains generated by climate stabilization in the model under analysis. This leaves open the philosophical question of whether it is better to have a small population that enjoys a high quality of life or a large population of people who are less happy.

Finally, a model with high risk aversion and non-trivial catastrophic risks is prone to exhibit the behavior that lies behind Weitzman's (2009) "Dismal Theorem." In Weitzman's analysis, catastrophic risks impose infinitely large cost on social welfare in low-probability states where the consumption level is driven to zero. To address this anomaly, our analysis assigns a value to the parameter k that bounds the utility function from below at zero in the case where consumption falls to the subsistence level of c_m = $228 per person per year (2005 dollars; Ahmed et al., 2007). Below the subsistence level, the population is unable to sustain itself, and the economy comes to a finite, stochastic end. Such events, while extremely rare, do arise under the assumptions of this model.

3. Policy scenarios and welfare implications

The model summarized above defines an equilibrium path for the economy once policy-makers stipulate a set of greenhouse gas emissions abatement policies. As in Gerst et al. (2013), we focus on a family of alternative policy scenarios in which the greenhouse gas emissions control rate (μ_t) rises from a value of zero in 2010 to unity in the year 2270 following an S-shaped Weibull function

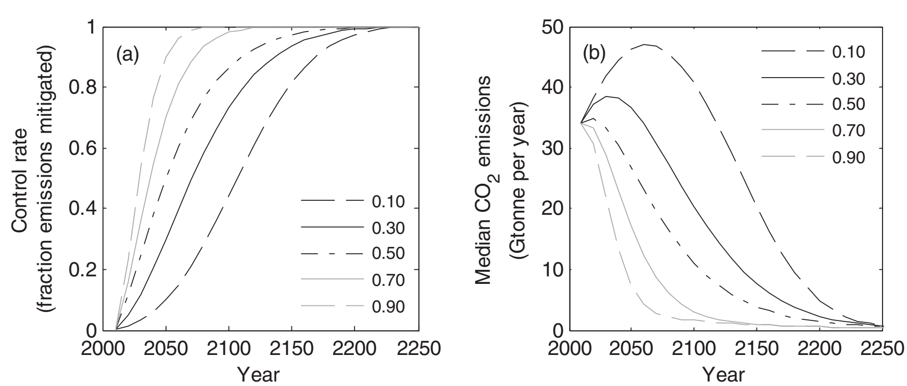

Fig. 2. Emissions control rates and median carbon dioxide emissions for selected policy scenarios. The legend denotes the emissions control rate in the year 2050.

(Fig. 2). The speed of the transition is then defined by the stringency of emissions reductions in the year 2050. In the baseline (or "no control") path, greenhouse gas emissions remain unregulated, although market forces lead to a transition to a low-carbon energy economy in the very long run. In the remaining cases, the timing of the move to a low-carbon future varies greatly according to the chosen scenario.

In the absence of mitigation measures, the model anticipates large increases in greenhouse gas emissions, which rise to 68 billion metric tons of carbon dioxide per year in 2100. This leads the atmospheric concentration of carbon dioxide to increase to a maximum median value of 740 parts per million. The most stringent scenario, in contrast, involves a rapid phase-out of carbon-based energy technologies and a move toward higher energy efficiency. In this case, the median atmospheric carbon dioxide concentration peaks at 460 parts per million around mid-century and declines from that point forward.

In the "no controls" scenario, high levels of emissions lead to a median temperature increase of 3.6 °C and 7.4 °C for the years 2100 and 2300, respectively (Fig. 3a). With fat-tailed uncertainty about climate sensitivity, this large median change is accompanied by a 5% chance that temperatures will exceed 5.7 °C and 16.5 °C in years 2100 and 2300 (Fig. 3c). This in turn leads to a 3.3-in-1000 chance of what we term a "climate catastrophe" – a case in which climate damages become so severe that the standard of living is driven all the way down to the subsistence level at some point during the next 400 years. This result gives weight to the factual assumptions behind Weitzman's (2009) "Dismal Theorem." By comparison, the assumption that climate sensitivity follows a thin-tailed, gamma distribution yields to somewhat lower median temperature increases and a reduced level of variability (Fig. 3b and d). This change in itself reduces the probability of a climate catastrophe to 2.5-in-10,000 chance in the case where greenhouse gas emissions remain unregulated.

As one would expect, the lower emissions pathways lead to lower changes in median global temperature and to substantially smaller climatic risks. The risk of a climate catastrophe under fat-tailed uncertainty, for example, is reduced to the order of 1-in-100-million when the emissions control rate in the year 2050 meets or exceeds 55%. This low probability points to an important methodological feature of the model. Given the emphasis we attach to low-probability, catastrophic events, to obtain statistically reliable results, it is necessary to conduct Monte Carlo simulations involving 100 million draws. Otherwise, rare events that are quite important to the results would be under-sampled, leading to biased welfare estimates. Given the "curse of dimensionality" (Bellman, 1957), this raises major issues of tractability in more complex models, especially those that aim to apply numerical algorithms to solve for optimal policies. See McInerney et al. (2012) for a discussion of computationally feasible approaches to evaluating climate change policy decisions under strong uncertainty.

As in Gerst et al. (2013), we measure the welfare effects of the various policy scenarios using a metric we call "stationary consumption." In conceptual terms, stationary consumption represents the level of consumption that – if maintained at a constant level at all points in time under all states of nature – would provide the same level of welfare associated with a given, uncertain consumption stream. In technical terms, stationary consumption (SC) may be derived by replacing c_t by SC in Eq. (4) and solving for SC as a function of the applicable level of social welfare (W):

$$SC = \left[\frac{W + (1-\gamma)^{-1} c_m^{1-\gamma} E\left[\sum_{t=0}^{\infty} (1+\rho)^{-t} N_t\right]}{(1-\gamma)^{-1} E\left[\sum_{t=0}^{\infty} (1+\rho)^{-t} N_t\right]}\right]^{1/(1-\gamma)}. \quad (6)$$

In Fig. 4, we report the level of stationary consumption as a percentage of the highest level that arises in our policy scenarios for each assumed level of risk aversion. This normalization allows for the easy comparison of results using just one set of axes.

As the figure shows, the level of risk aversion has a very important bearing on the welfare implications of different policy decisions. Notably, applying the standard DICE parameter value of $\gamma = 2.0$ implies that emissions controls provide almost negligible net social benefits. Given this preference parameter, the main risk is that policy makers will opt for excessively strict greenhouse gas

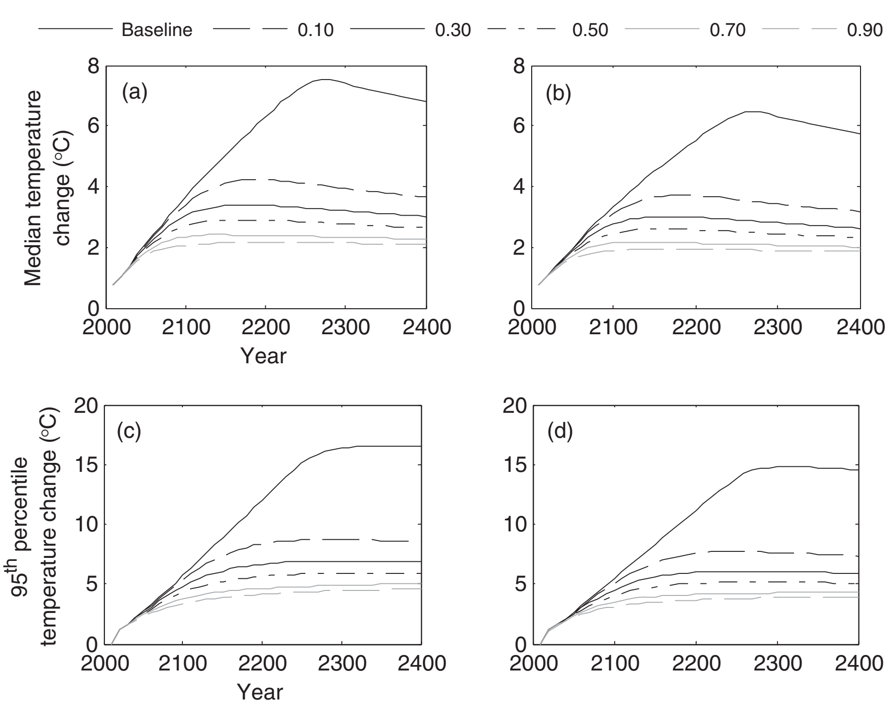

Fig. 3. Temperature increase relative to the pre-industrial norm for selected policy scenarios given fat-tailed (a and c) and thin-tailed (b and d) climate sensitivity. Thin-tailed climate sensitivity fat-tailed climate sensitivity

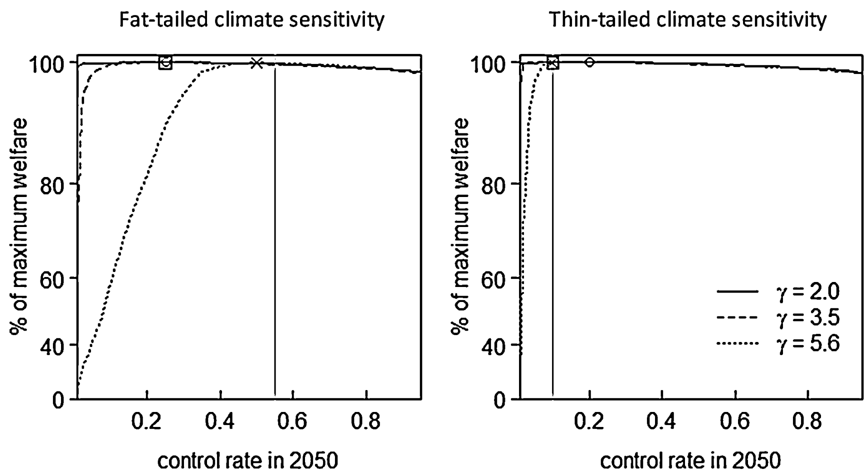

Fig. 4. Relative social welfare as a function of the emissions control rate in 2050 given fat-tailed (left) and thin-tailed (right) uncertainty concerning climate sensitivity. Relative welfare is expressed as the percentage of the highest observed stationary consumption level that arises for each level of risk aversion. The control rate with the highest net benefit is marked by a circle for $\gamma = 2.0$, a square for $\gamma = 3.5$, and an $\times$ for $\gamma = 5.6$. Note that a non-linear vertical scale is used to better show the differences among parameter settings. A solid vertical line demarcates the boundary for the occurrence of catastrophe. Scenarios to the left experienced at least one catastrophe in the 100 million model runs considered in the analysis, while those to the right did not. Thin-tailed climate sensitivity fat-tailed climate sensitivity

emissions reductions, an outcome that serves to reduce the stationary consumption level by up to 1.4%. Intuitively, this is akin to giving up 1.4% of consumption in the short to medium term to obtain benefits that are too far off in the future to strongly influence decisions taken today.

In contrast, calibrating preferences to reflect decision-makers' high observed risk aversion ($\gamma = 5.6$) favors relatively high control rates and suggests that even overly aggressive reductions are strongly preferable to little or no control. In this case, under-abatement can impose welfare costs equivalent to a reduction of up to 79% in the stationary consumption level. The highest level of welfare arises when the emissions control rate increases to 50% in the year 2050. Welfare, however, is nearly invariant across the set of scenarios that achieve rapid decarbonization. Switching to policies that achieve a 90% abatement level in 2050, for example, reduces the stationary consumption level by only 1.1%. This is true despite the relatively conservative approach taken in this analysis, which – contrary to the work of Gerst et al. (2010) and Kopp and Mignone (2012) – assumes a narrow range of uncertainty concerning the climate change damage function. Relaxing this assumption would clearly increase perceived climate risks in a manner that would support earlier and more stringent abatement.

These qualitative results holds true regardless of whether uncertainty about climate sensitivity is represented by a fat- or thin-tailed distribution. Although the fat-tailed distribution implies a need for higher emissions control rates to maximize net benefits, failing to cut emissions adequately still leads to very large welfare losses under thin-tailed uncertainty. Clearly, these results are driven by risk aversion. With the level of risk aversion implied by the observed difference in returns between risky and safe assets, decision makers attach an especially high weight to the value of precautionary actions that reduce the threat of catastrophic outcomes.

It is worth noting, however, that although our core results are driven by high risk aversion, our model nonetheless yields finite levels of social welfare in all policy scenarios. Moreover, climate catastrophes still occur with low probability in the emissions control scenarios that yield the highest levels of welfare. Thus, our setup avoids the paradox associated with Weitzman's (2009) "Dismal Theorem," in which exceedingly small risks completely dominate policy decisions by driving social welfare to minus infinity (Nordhaus, 2012). In the present analysis, the social welfare function is bounded from below at zero as described in Section 2 above. It is therefore bad – but not infinitely bad – to consider policy outcomes in which consumption is driven down to subsistence.

4. The social cost of carbon

The final step in our analysis is to calculate the social cost of carbon as a function of the stringency of emissions control policies and decision-makers' risk aversion. As noted in the introduction, it is well-known that, all else equal, the use of low discount rates leads to a relatively high value for the social cost of carbon, which reflects the marginal social cost imposed by short-run greenhouse gas emissions (see Johnson and Hope, 2012). In addition, authors such as Hope (2006) have argued that the social cost of carbon does not depend strongly on the timing and stringency of greenhouse gas emissions reductions. As we shall see, this hypothesis does not hold in the model under consideration; in particular, accounting for uncertainty and risk aversion fundamentally changes the results.

To see this, we calculate the social cost of carbon numerically using the formula:

$$\mathrm{SCC} = -\frac{\partial W/\partial E_0}{\mathrm{MUC}_0}. \tag{7}$$

In this expression, $\partial W/\partial E_0$ is the change is social welfare that arises when the level of greenhouse gas emissions at date $t = 0$ is exogenously increased by one incremental unit. This measure – which is negative because current emissions reduce future consumption and therefore welfare – reflects the impacts of current emissions on future climate, the impacts of climate change on the economy, and decision-makers' attitudes toward time and risk. Because $\partial W/\partial E_0$ is measured in utility units, it is necessary to convert it to monetary units by dividing through by the marginal utility of consumption at date $t = 0$ (MUC_0). Because we increase emissions exogenously without adjusting the emissions control rate, this calculation captures the marginal present-value costs of climate impacts without accounting for the marginal benefits of current emissions.

As shown in Fig. 5, assuming the low value of risk aversion that Nordhaus (2008) employs in DICE ($\gamma = 2.0$) generates a social cost of carbon in the initial period of the analysis (2010) that is close to \$10 per metric ton of carbon dioxide (measured in 2005 dollars).

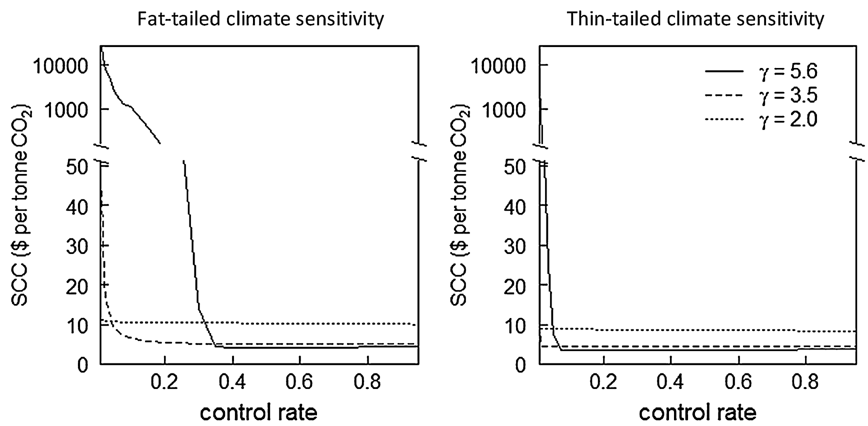

Fig. 5. The social cost of carbon ($ per metric ton of carbon dioxide) in 2010 as a function of the emissions control rate in 2050 with uncertainty about climate sensitivity represented by a fat-tailed distribution (left) and thin-tailed distribution (right). Note that a broken vertical axis is used with a logarithmic scale above the break.

This value is nearly independent of the emissions control path and does not depend on whether uncertainty about climate sensitivity is fat- or thin-tailed. This is consistent with the welfare evaluation shown in Fig. 4, where the various policy scenarios yield relatively similar levels of social welfare.

In contrast, the social cost of carbon assumes the markedly high value of $25,700 per metric ton of carbon dioxide in the initial period of the model when: (a) greenhouse gas emissions remain uncontrolled; and (b) the coefficient of relative risk aversion is set equal to its central value ($\gamma = 5.6$) as determined by Ding et al.'s (2012) analysis of data from financial markets given sequentially rational investment behavior. This result fits well with the welfare analysis presented above. Under these circumstances, incremental cuts in greenhouse gas emissions lead to especially large risk reductions and welfare gains. In addition, the results show the social cost of carbon then falls sharply with the level of emissions control, converging to the neighborhood of $4–5 per ton when the emissions control rate in year 2050 exceeds 40% – the point at which most of the risk reduction value has been accomplished.

With thin-tailed uncertainty concerning climate sensitivity, the risk of climate catastrophes is reduced, and the social cost of carbon in the "no control" scenario assumes a value of $1690 per metric ton of carbon dioxide when the coefficient of risk aversion is high ($\gamma = 5.6$). This figure – although lower than the estimate we obtain given fat-tailed uncertainty – is in itself is considerably higher than typical values for the social cost of carbon reported elsewhere in the literature. This suggests that our paper's key qualitative conclusion – that high risk-aversion leads to an especially high value of the social cost of carbon that is sensitive to the emissions control rate – is not driven by our assumptions concerning climate sensitivity.

Substantially lower shadow prices arise when the level of risk aversion is lowered to a value of $\gamma = 3.5$. Under fat-tailed uncertainty concerning climate sensitivity, this reduces the maximum value of the social cost of carbon to $44 per metric ton of carbon dioxide. Under thin-tailed climate sensitivity, it yields a social cost of carbon that is actually below the value generated by the standard parameterization of DICE. Although perhaps surprising at first blush, this point may be explained by the fact that in our model the coefficient γ measures both the level of risk aversion and decision-maker's elasticity of intertemporal substitution. When $\gamma = 3.5$, the level of risk aversion is too low to attach importance to reductions in the likelihood of low-probability risks that occur in the relatively distant future. At the same time, this parameter value implies that the marginal utility of consumption falls rapidly in "good" states of nature characterized by high consumption growth. Hence assuming that $\gamma = 3.5$ is analogous to using a relatively high effective monetary discount rate in the model under consideration, except for policy scenarios involving very low levels of emissions reduction (see Broome, 2008).

This analysis may shed light on Nordhaus' (2008) finding that accounting for uncertainty can in some cases lead to reductions in the social cost of carbon (see Newbold and Daigneault, 2009). As Howarth (2003) explains, uncertainties in the baseline rate of economic growth can induce a positive correlation between climate change damages and future consumption levels, generating a negative risk premium in which certainty-equivalent damages are lower than expected damages. In the same breath, uncertainties about the damage function can reverse this correlation, especially if catastrophic damages would lead to particularly low consumption levels in some states of nature. As the results depicted in Fig. 5 suggest, the balance between these effects is sensitive to both the degree of risk and to decision-makers' risk aversion. A high risk premium arises in our model when climate sensitivity is fat-tailed, risk aversion is high, and/or the rate of emissions control is low, leading to a higher risk of catastrophic climate impacts.

5. Conclusions

This paper has assessed the social cost of carbon in a stochastic growth model adapted to account for the costs of greenhouse gas emissions reductions, the relationship between emissions and future mean global temperature, and the economic impacts of climate change. The model is novel because it provides an internally consistent approach to analyzing the net benefits of climate stabilization under strong uncertainty. It allows for fat-tailed uncertainty concerning climate sensitivity, and its assumptions concerning time and risk preferences are grounded in a formal representation of intertemporal investment behavior under uncertainty that is applied to empirical observations of the rate of consumption growth and the market rates of return on safe and risky assets. Importantly, we assign a higher value to the coefficient of relative risk aversion than many other studies in the integrated assessment literature. Our calibration, however, is supported by both the principle of revealed preference and by experimental studies concerning people's attitudes toward climate risks.

The analysis finds that the social cost of carbon attains the particularly high numerical value of $25,700 per metric ton of carbon dioxide given a baseline policy scenario in which greenhouse gas emissions remain unregulated. This shadow price falls rapidly with the stringency of emissions abatement policies. For policy scenarios that lead to rapid emissions reductions aimed at stabilizing the global climate, the shadow price of carbon – i.e.

the present-value marginal benefit of emissions abatement – is just \$4–5 per metric ton. These results are in sharp contrast with the findings of previous studies (Hope, 2006; Interagency Working Group, 2010), which concluded that the social cost of carbon is largely independent of the time path for emissions. This suggests that accounting for the threat of catastrophic climate impacts in a model assuming a realistically high degree of risk aversion – can strongly influence the results of integrated assessment models.

In an important sense, we believe that our results show that the logic and intuition of the familiar "diamond-water" paradox (see Farber et al., 2002) are applicable to the climate problem under consideration. Our model, for example, suggests that moving from the no-abatement baseline to aggressive emissions abatement can produce net social benefits equivalent to an almost four-fold increase in per capita consumption at all points in time and in each uncertain state of nature. This is true because deep emissions cuts are sufficient to reduce the relatively low probability of catastrophic climate impacts to essentially zero. In our model, risk-averse decision-makers attach especially high value to averting the risk of climate catastrophes. Since these risks affect aggregate economic activity, they cannot be resolved through risk pooling or risk spreading (Arrow and Lind, 1970). Thus social evaluations should reflect the same preferences held by private economic actors. It is perhaps fair to say, then, that the key problem in climate change policy is to manage risk by stabilizing a biophysical system characterized by fundamental uncertainties (in this case concerning climate sensitivity).

Our results also show, however, that the *marginal* benefit of emissions reductions is quite low in scenarios that successfully abate risk. Indeed, the level of social welfare in our model is virtually the same for all policy scenarios that achieve emissions control rates of at least 40% by the year 2050. This means that under-abatement can generate large, avoidable net social losses, while over-abatement would have only minor effects on the net benefits of climate change policies. Like water, climate stabilization may be crucial to sustained human flourishing in ways that generate a very large total economic value. Yet at the margin, the benefits of incremental water supplies or greenhouse gas emissions reduction are quite low once the core goal of adequate provisioning is attained. This implies that economists' normal emphasis on equating marginal costs and benefits may be a problematic framing in the context of climate change policy, especially given the very large uncertainties that exist in characterizing and valuing climate change impacts in a way that properly values risk and precaution.

We close, then, by observing that the social cost of carbon is contingent on a thick set of assumptions regarding empirical analysis, modeling methods, ethics, and the interdependence between short-run and long-run policy decisions in determining the time path of greenhouse gas emissions and, therefore, the risks and damages that today's emissions will impose on future generations. There is no fact-of-the-matter concerning the social cost of carbon that can provide an objective and value-free guide for policy evaluation. Instead, we agree with Dietz (2012) that it may be desirable to shift the debate away from its current emphasis on calculating present-value marginal benefits and toward a focus on the carbon prices needed to achieve stipulated emissions reductions at the lowest possible social cost. This is especially salient since climate governance involves core moral considerations that do not reduce easily to the language and metrics of cost-benefit analysis (Howarth, 2011).

Acknowledgements

We thank the referees and the editors for their helpful and engaging comments on the initial version of this paper. We also thank the Yale School of Forestry and Environmental Studies for hosting a seminar on this work.

References

Ackerman, F., DeCanio, S.J., Howarth, R.B., Sheeran, K., 2009. Limitations of integrated assessment models of climate change. Climatic Change 95, 297–315.
Ackerman, F., Stanton, E.A., 2012. Climate risks and carbon prices: revising the social cost of carbon. Economics E-Journal 6, 2012–2110.
Ahmed, A., Vargas-Hill, R., Smith, L., Wiesmann, D., Frankenberger, T., 2007. The World's Most Deprived: Characteristics and Causes of Extreme Poverty and Hunger. International Food Policy Research Institute, Washington, D.C..
Anthoff, D., Tol, R.S.J., 2010. The Climate Framework for Uncertainty, Negotiation and Distribution (FUND), Technical Description, Version 3.5. Economic and Social Research Institute, Dublin.
Anthoff, D., Tol, R.S.J., Yohe, G., 2009. Risk aversion, time preference, and the social cost of carbon. Environmental Research Letters 4, 1–7.
Arrow, K.J., Lind, R.C., 1970. Uncertainty and the evaluation of public investment decisions. American Economic Review 60, 364–378.
Atkinson, G., Dietz, S., Helgeson, J., Hepburn, C., Sælen, H., 2009. Siblings, not triplets: social preferences for risk, inequality and time in discounting climate change. Economics E-Journal 3, 2009–2026.
Baker, M.B., Roe, G.H., 2009. The shape of things to come: Why is climate change so predictable? Journal of Climate 22, 4574–4589.
Barro, R.J., 2006. Rare disasters and asset markets in the twentieth century. Quarterly Journal of Economics 121, 823–866.
Barsky, R.B., Juster, F.T., Kimball, M.S., Shapiro, M.D., 1997. Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement study. Quarterly Journal of Economics 112, 537–579.
Bellman, R.E., 1957. Dynamic Programming. Princeton University Press, Princeton, NJ.
Broome, J., 2008. The ethics of climate change. Scientific American 298, 97–102.
Cline, W.R., 1992. The Economics of Global Warming. Institute for International Economics, Washington, D.C..
Clinton, W.J., 1993. Executive Order 12866 of September 30, 1993 – Regulatory Planning and Review. The White House, Washington, D.C..
Department for Environment Food and Rural Affairs, 2002. Valuing the Social Cost of Carbon Emissions: DEFRA Guidance. DEFRA, London.
Dietz, S., 2012. The treatment of risk and uncertainty in the US social cost of carbon for regulatory impact analysis. Economics E-Journal 6, 2012–2018.
Ding, P., Gerst, M.D., Bernstein, A., Howarth, R.B., Borsuk, M.E., 2012. Rare disasters and risk attitudes: international differences and implications for integrated assessment modeling. Risk Analysis 32, 1846–1855.
Farber, S.C., Costanza, R., Wilson, M.A., 2002. Economic and ecological concepts for valuing ecosystem services. Ecological Economics 41, 375–392.
Gerst, M.D., Howarth, R.B., Borsuk, M.E., 2010. Accounting for the risk of extreme outcomes in an integrated assessment of climate change. Energy Policy 38, 4540–4548.
Gerst, M.D., Howarth, R.B., Borsuk, M.E., 2013. The interplay between risk attitudes and low probability, high cost outcomes in climate policy analysis. Environmental Modelling and Software 41, 176–184.
Hahn, R.W., Sunstein, C.R., 2002. A new executive order for improving federal regulation? Deeper and wider cost-benefit analysis. University of Chicago Institute for Law and Economics: Olin Working Paper No. 150.
Hall, R.E., 1988. Intertemporal substitution in consumption. Journal of Political Economy 96, 339–357.
Hansen, J.E., Sato, M., Kharecha, P., Beerling, D., Masson-Delmotte, V., Pagani, M., Raymo, M., Royer, D.L., Zachos, J.C., 2008. Target atmospheric CO_2: where should humanity aim? Open Atmospheric Science Journal 2, 217–231.
Hope, C., 2006. The social cost of carbon: what does it actually depend on? Climate Policy 6, 565–572.
Hope, C., 2008. Optimal carbon emissions and the social cost of carbon under uncertainty. Integrated Assessment Journal 8, 107–122.
Howarth, R.B., 2003. Discounting and uncertainty in climate change policy analysis. Land Economics 79, 369–381.
Howarth, R.B., 2011. Intergenerational justice. In: Dryzek, J.S., Norgaard, R.B., Schlosberg, D. (Eds.), The Oxford Handbook on Climate Change and Society. Oxford University Press, Oxford.
Interagency Working Group on Social Cost of Carbon, 2010. Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866. U.S. Environmental Protection Agency, Washington, D.C..
Iverson, T., Perrings, C., 2011. Precaution and proportionality in the management of global environmental change. Global Environmental Change 22, 161–177.
Johnson, L.T., Hope, C., 2012. The social cost of carbon in U.S. regulatory impact analyses: an introduction and critique. Journal of Environmental Studies and Sciences 2, 205–221.
Kopp, R.E., Golub, A., Keohane, N.O., Onda, C., 2012. The influence of the specification of climate change damages on the social cost of carbon. Economics E-Journal 6, 2012–2013.
Kopp, R.E., Mignone, B.K., 2012. The U.S. government's social cost of carbon estimates after their first two years: pathways for improvement. Economics E-Journal 6, 2012–2015.
Kousky, C., Kopp, R.E., Cooke, R., 2011. Risk premia and the social cost of carbon: a review. Economics E-Journal 5, 2011–2021.

Lenton, T.M., Held, H., Kriegler, E., Hall, J.W., Lucht, W., Rahmstorf, S., Schellnhuber, H.J., 2008. Tipping elements in the earth's climate system. Proceedings of the National Academy of Sciences 105, 1786–1793.

Lucas, R., 1978. Asset prices in an exchange economy. Econometrica 46, 1429–1445.

Lutz, W., Sanderson, W., Scherbov, S., 2008. The coming acceleration of global population. Nature 451, 716–719.

Masur, J.S., Posner, E., 2011. Climate regulation and the limits of cost-benefit analysis. California Law Review 99, 1557–1599.

McInerney, D., Lempert, R., Keller, K., 2012. What are robust strategies in the face of uncertain climate threshold responses? Climatic Change 112, 547–568.

Mehra, R., Prescott, E., 1985. The equity premium: a puzzle. Journal of Monetary Economics 15, 145–161.

Neumayer, E., 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, 297–301.

Newbold, S.C., Daigneault, A., 2009. Climate response uncertainty and the benefits of greenhouse gas emissions reductions. Environmental and Resource Economics 44, 351–377.

Nordhaus, W.D., 2008. A Question of Balance: Weighing the Options on Global Warming Policies. Yale University Press, New Haven, CT.

Nordhaus, W.D., 2012. Economic policy in the face of severe tail events. Journal of Public Economic Theory 14, 197–219.

Pearce, D.W., 2003. The social cost of carbon and its policy implications. Oxford Review of Economic Policy 19, 362–384.

Pindyck, R.S., 2013. Climate change policy: what do the models tell us? Journal of Economic Literature 51, 860–872.

Roe, G.H., Baker, M., 2007. Why is climate sensitivity so unpredictable? Science 318, 629–632.

Romer, D., 2006. Advanced Macroeconomics, third ed. McGraw-Hill, New York.

Rose, S.K., 2010. The role of the social cost of carbon in policy. Wiley Interdisciplinary Reviews: Climate Change 3, 195–212.

Sinnott-Armstrong, W., 2012. In: Zalta, E.N. (Ed.), Consequentialism. The Stanford Encyclopedia of Philosophy, http://plato.stanford.edu/archives/win2012/entries/consequentialism.

Stern, N., 2007. The Economics of Climate Change. Cambridge University Press, Cambridge.

Stern, N., 2013. The structure of economic modeling of the potential impacts of climate change: grafting gross underestimation of risk onto already narrow science models. Journal of Economic Literature 51, 838–859.

Tol, R.S.J., 1997. On the optimal control of carbon dioxide emissions – an application of FUND. Environmental Modelling and Assessment 2, 151–163.

Tol, R.S.J., 2011. The social cost of carbon. Annual Review of Resource Economics 3, 419–443.

Tomassini, L., Reichert, P., Knutti, R., Stocker, T.F., Borsuk, M.E., 2007. Robust bayesian uncertainty analysis of climate system properties using markov chain Monte Carlo methods. Journal of Climate 20, 1239–1254.

Weitzman, M.L., 2009. On modeling and interpreting the economics of catastrophic climate change. Review of Economics and Statistics 91, 1–19.

Zickfeld, K., Morgan, M.G., Frame, D.J., Keith, D.W., 2010. Expert judgments about transient climate response to alternative future trajectories of radiative forcing. Proceedings of the National Academy of Sciences 107, 12451–12456.