Climatic Change (2008) 88:251–265
DOI 10.1007/s10584-007-9391-8

Atmospheric stabilization and the timing of carbon mitigation

Bryan K. Mignone · Robert H. Socolow · Jorge L. Sarmiento · Michael Oppenheimer

Received: 19 September 2006 / Accepted: 6 November 2007 / Published online: 13 February 2008

Abstract Stabilization of atmospheric CO_2 concentrations below a pre-industrial doubling (~550 ppm) is a commonly cited target in climate policy assessment. When the rate at which future emissions can fall is assumed to be fixed, the peak atmospheric concentration – or the stabilization "frontier" – is an increasing and convex function of the length of postponement. Here we find that a decline in emissions of 1% year^{-1} beginning today would place the frontier near 475 ppm and that when mitigation is postponed, options disappear (on average) at the rate of ~9 ppm year^{-1}, meaning that delays of more than a decade will likely preclude stabilization below a doubling. When constraints on the future decline rate of emissions are relaxed, a particular atmospheric target can be realized in many ways, with scenarios that allow longer postponement of emissions reductions requiring greater increases in the intensity of future mitigation. However, the marginal rate of substitution between future mitigation and present delay becomes prohibitively large when the balance is shifted too far toward the future, meaning that some amount of postponement cannot be fully offset by simply increasing the intensity of future mitigation. Consequently, these results suggest that a practical transition path to a given stabilization target in the most commonly cited range can allow, at most, one or two decades of delay.

B. K. Mignone (✉)
Foreign Policy Studies Program, The Brookings Institution, Washington, DC 20036, USA
e-mail: bmignone@brookings.edu

R. H. Socolow
Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

J. L. Sarmiento
Atmospheric and Oceanic Sciences Program, Princeton University, Princeton, NJ 08544, USA

M. Oppenheimer
Woodrow Wilson School of Public and International Affairs, Princeton, NJ 08544, USA

M. Oppenheimer
Department of Geosciences, Princeton University, Princeton, NJ 08544, USA

1 Introduction

Fifteen years after the ratification of the UN Framework Convention on Climate Change (UNFCCC), stabilization of atmospheric CO_2 at levels that avoid "dangerous anthropogenic interference" (DAI) with the climate system (UNFCCC 1992) remains a central but increasingly elusive goal of climate policy. While no single factor – political, economic, technological or scientific – can be held responsible for the paucity of action, the inherent ambiguity in the concept of DAI has often been implicated as an additional barrier to effective policymaking.[1] Recently, a growing body of scientific evidence has suggested that stabilization at or below a doubling of pre-industrial values (~550 ppm) will be necessary to hedge against some of the most severe climatic outcomes, such as the widespread demise of coral reefs (Hoegh-Guldberg 1999), disintegration of the West Antarctic Ice Sheet (WAIS; Oppenheimer and Alley 2004; Oppenheimer 1998) or disintegration of the Greenland Ice Sheet (GIS; Hansen 2005, 2004; Gregory et al. 2004).[2] Despite this convergence of opinion, the simple fact that any single atmospheric outcome can be realized in many ways means that, even when the ultimate objective is clear, the exact approach may not be. In this case, some guidance can be found in Article 3 of the UNFCCC itself, which adds that "policies and measures to deal with climate change should be cost-effective so as to ensure global benefits at the lowest possible cost" (UNFCCC 1992). With some notable exceptions (e.g. Yohe et al. 2004; Webster 2002; Ha-Duong et al. 1997), most economic analyses of climate change have concluded that, whenever tradeoffs can be made, some postponement is preferable to immediate action (e.g. Wigley et al. 1996).

In general, the strong economic preference toward postponement – as opposed to the political preference toward postponement – stems from three well-established facts (see, e.g., Wigley et al. 1996 for a discussion). First, the time constant for the expiration of the capital stock is known to be quite large (at least a generation), and premature capital retirement is a costly proposition.[3] Secondly, some degree of postponement would presumably allow for the development and deployment of cheaper and less carbon-intensive technology. Finally, a positive discount rate means that the net present value of future climate mitigation benefits is lower than the net present value of immediate abatement costs, even when the actual benefits and costs are comparable in magnitude. When these preferences are included in an analysis that also assumes atmospheric stabilization, postponement clearly cannot go on indefinitely (because it would violate the stabilization constraint), but it can go on long

[1] The statement by US President George W. Bush in June 2001 that "no one can say with any certainty what constitutes a dangerous level of warming, and therefore what level must be avoided" is often cited as evidence for this claim. For an alternative view – that ambiguity can, in principle, facilitate action – see Oppenheimer (2005).

[2] Scientific studies typically identify temperature (rather than concentration) thresholds for singular events such as these. For example, Hoegh-Guldberg (1999) suggests that coral bleaching could become prevalent if the global temperature were to rise by more than 1°C (relative to 1990), Oppenheimer (1998) suggests that future disintegration of WAIS could be triggered by a global temperature increase of 2°C, while Hansen (2005, 2004) suggests that disintegration of GIS might be triggered by a global increase of only 1°C, which is roughly consistent with the 3°C *local* increase that Gregory et al. (2004) require. The conversion from temperature to concentration targets can be made by noting that stabilization at 550 ppm is projected to increase the equilibrium temperature relative to 1990 by 2.0–5.2°C, assuming a plausible range of climate sensitivities (Watson et al. 2001). The policy implications of similar DAI metrics are explored in O'Neill and Oppenheimer (2002), while a discussion of alternative metrics can be found in the reviews by Oppenheimer and Petsonk (2005), Dessai (2004) and Corfee-Morlot and Hohne (2003) among others.

[3] However, this particular benefit is mitigated somewhat by the presumed construction of new capital facilities during the period of postponement that may "lock-in" suboptimal technology and render future carbon mitigation activity more costly.

enough to defer much of the mitigation burden by several decades (Wigley et al. 1996; Richels and Edmonds 1995). However, two other considerations need to be kept in mind. First, it is possible that excessively large future declines required by the stabilization constraint may be physically infeasible, for example if developments in low-carbon technology do not keep pace with the rate at which energy is demanded by the global economy. Secondly, even before the physical constraint binds, it seems likely that the high costs associated with a future rapid decline in emissions would make modest advancements in the onset of mitigation preferable, even when costs are evaluated in terms of net present value.

Our primary goal in this paper is to examine the question of timing using only a simple model of the global carbon cycle (c.f. Socolow and Lam 2007). The main discussion (in Section 3 below) consists of two parts. In the first, we consider a world in which the future rate of emissions decline is fixed so that the minimum stabilization target – or the stabilization "frontier" – is uniquely determined (for a fixed near-term "business as usual" emissions trajectory) by the number of years beyond today that mitigation is postponed. Assuming a decline rate of 1% year^{-1} and a neutral terrestrial biosphere (in which net terrestrial emissions are assumed to be zero), we find that immediate mitigation would place the frontier near 475 ppm and that each additional year of delay increases it by about 9 ppm on average, meaning that stabilization below a pre-industrial doubling will require the onset of dedicated mitigation within about a decade. In the second part of our analysis, we introduce the emissions decline rate as an additional parameter in order to show how near-term postponement and future increases in the intensity of mitigation can be traded off, that is, how they can be varied in tandem to achieve the same final atmospheric goal. We find that the marginal rate of substitution between future and present mitigation (i.e. the increase in future decline rate needed to offset a given amount of delay) becomes quite large when the decline rate increases beyond 1 or 2% per year, meaning that small increases in delay necessitate very large increases in the intensity of future mitigation. We claim that the strong convexity of the stabilization frontier makes it unlikely that additional delays will yield positive economic benefits in a world that is committed to stabilization below a doubling.

2 Methods

In this study, we use a simple, but well-tested box-diffusion model of the world ocean (HILDA) (Shafer and Sarmiento 1995; Siegenthaler and Joos 1992) to study the carbon uptake response to a diverse set of prescribed CO_2 emissions scenarios. The model includes a parameterized treatment of ocean circulation (with explicit contributions from advection and diffusion), surface carbonate chemistry comparable to that used in OCMIP-2 (Orr et al. 2000) and gas exchange at the surface parameterized according to the Wanninkhof (1992) formulation.

After spinning up the model for 10,000 years, we force the atmosphere with the historical CO_2 record between 1750 and 2004 (so that atmospheric pCO_2 in 2004 approaches 378 ppm) and tune the upper ocean diffusivity profile so that predicted ocean uptake in 1995 falls near the central estimate (2.3 Pg C year^{-1}) of several recent observational and modeling studies (see, e.g., Mignone et al. 2006 and references therein). This process yields a future uptake response to prescribed forcing in our model that is consistent with the uptake response in more sophisticated ocean general circulation models.

In propagating the model forward in time, net fossil fuel emissions from CO_2 are prescribed, beginning at approximately 6.8 Pg C year^{-1} in 2004 and rising at the rate of 0.2 Pg C year^{-1} (i.e. increasing by 1 Pg C year^{-1} every 5 years) under "business as usual"

(BaU) to approximately 16 Pg C year^{-1} in 2050, thus roughly tracking the IPCC A1B scenario (Houghton et al. 2001). The future atmospheric accumulation is calculated as the difference between net fossil emissions and the model-predicted ocean uptake. In all cases, we assume a "neutral" terrestrial biosphere in which the total source from deforestation exactly balances the sink from afforestation and fertilization. In other words, we assume that net terrestrial emissions are zero over the course of the simulation.

3 Results

3.1 Atmospheric option value of carbon mitigation

We begin our analysis by examining the sensitivity of the atmospheric CO_2 response to the timing of mitigation. If near-term fossil fuel emissions climb linearly at a rate of ~1 Pg C year^{-1} every 5 years until future mitigation forces them to fall at a constant rate of 1% year^{-1} (beginning at some point during the next half century), then atmospheric CO_2 concentrations will peak within approximately the next two centuries, when the emissions rate first falls below the ocean uptake rate (see Fig. 1). This peak effectively defines a frontier below which stabilization – when taken to imply a hard ceiling constraint on concentration – will be impossible.[4] Our goal in this section is to show how the stabilization frontier depends on a single macroscopic policy decision variable, namely the number of years beyond today that mitigation is postponed.

A representative subset of the scenarios considered in this study is shown in Fig. 1. Panel a shows six emissions trajectories for which the end members are immediate mitigation (dark blue curve) and 50-year postponement of mitigation (black curve). Intermediate trajectories representing postponement of 10, 20, 30 and 40 years are also shown, while intermediate trajectories representing postponement of 5, 15, 25, 35 and 45 years are not shown, although these trajectories are implicitly included in the rest of our analysis. In all of these simulations, mitigation occurs in two phases. In the first phase, which lasts a decade, emissions remain constant at the value obtained at the onset of mitigation; in the second phase, emissions decline from this value at a constant rate of 1% per year, so that all trajectories asymptotically approach zero in the distant future (the convergence is apparent by 2300). The initial "platform" duration of 10 years is arbitrary, as is the functional form of the emissions release itself, but the two-phase trajectory avoids having to assume an immediate reversal of emissions at the onset of mitigation. In effect, it replaces a single large discontinuity in the emissions growth rate (that occurs when the growth rate flips from positive to negative) with two smaller discontinuities (a switch from positive to zero growth followed by a switch from zero to negative growth) spread over the period of a decade.[5]

Panel b of Fig. 1 shows the ocean uptake as a function of time for each of the scenarios discussed above, while panel c shows the resulting atmospheric concentration trajectories.

[4] In a world in which concentrations are allowed to overshoot their final values, the concept of a frontier, as it is defined here, is less meaningful. However, the trouble with defining acceptable overshoots may limit the applicability of these types of scenarios.

[5] Qualitatively similar emissions trajectories were explored in Pacala and Socolow (2004). However, their 500 ppm stabilization scenario allows emissions to platform for 50 years (beginning today) before declining. See also Footnote 15, which discusses the difficulty one faces in choosing an appropriate emissions decline rate.

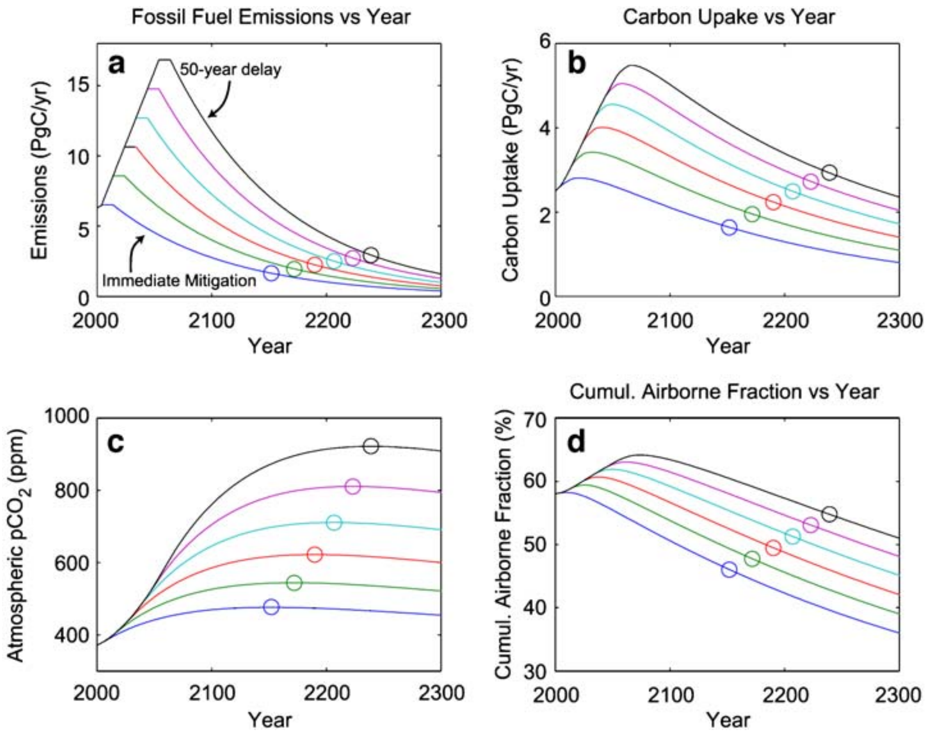

Fig. 1 Several illustrative scenarios of mitigation postponement. **a** Shows emissions trajectories for a subset of scenarios considered in this study. In all scenarios, emissions increase at the rate of 1 Pg C year^{-1} every 5 years until mitigation forces them to fall at some point in the future. The onset of mitigation varies between 0 and 50 years, but all scenarios assume that the reductions take the form of a 10-year emissions platform followed by a steady 1% year^{-1} decline. **b** shows the evolution of the ocean carbon uptake response, and **c** shows the resulting atmospheric trajectory. **d** gives the cumulative airborne fraction (defined as the ratio of the cumulative atmospheric accumulation to the cumulative emissions release) for the same scenarios. In each of the four panels, the open circles indicate the point at which atmospheric pCO_2 peaks in the relevant simulation. A comparison between **a**, **b** and **c** shows that the atmospheric concentration reaches its peak exactly when global emissions equal global uptake

A quick inspection of the latter makes it clear that postponing mitigation increases both the peak atmospheric CO_2 concentration and the amount of time it takes to reach this peak (the open circles in these figures indicate when the peak is reached). For example, the immediate mitigation scenario (blue curve) achieves a maximum atmospheric CO_2 concentration of 475 ppm in roughly 2150, while the 50-year postponement scenario (black curve) achieves a maximum concentration of 922 ppm in approximately 2240. Clearly, even if the total airborne fraction did not change with the addition of CO_2 (panel d shows that it does), one would still expect the atmospheric inventory to increase with the cumulative release. The non-linearity of this increase (i.e. the fact that the difference in accumulation between two scenarios increases with the cumulative release) results from the change in relative atmosphere–ocean partitioning that will be discussed in greater detail below. Qualitatively, the shift in the peak year can also be explained rather easily. If the peak is achieved when emissions first fall below the ocean uptake and if differences in ocean uptake are relatively small, then the time to peak is essentially the time it takes the emissions curve to reach the relatively low value of 2–4 Pg C year^{-1} that the ocean uptake approaches in later years. A closer look at panel a suggests that the amount of time it takes the emissions curve to fall below these values increases substantially with postponement, a result of the long exponential tail of the assumed emissions trajectories.[6]

The same results are presented more compactly in Fig. 2. Panel a of this figure shows the stabilization frontier as an explicit function of delay, while panel b shows the *marginal* atmospheric increase (MAI) in peak concentration as a function of delay. Panel b is effectively the time derivative of panel a; at any given time, it shows the *change* in peak concentration that would result from a 1 year shift in the onset of mitigation. Thus, an MAI of 6.6 ppm $year^{-1}$ at year 5 means that the peak atmospheric concentration would increase by 6.6 ppm if mitigation were delayed 1 year (i.e. if mitigation were begun in year 5 instead of year 4), while an MAI of 11.4 ppm $year^{-1}$ at year 50 means that the peak concentration would increase 11.4 ppm if the onset of mitigation were delayed from 49 years beyond today to 50 years beyond today.

The model results shown in panel a of Fig. 2 can be used to construct, by interpolation, a single continuous function that relates the stabilization frontier to the number of years of mitigation postponement (the solid curve in panel a shows a quadratic fit), which in turn makes it possible to associate with any particular stabilization target a hard upper bound on the amount of delay that can be tolerated. From this relationship, it is immediately apparent that stabilization in the most commonly cited range of 450–550 ppm will require the reversal of emissions growth within approximately the next decade. Stabilization at the higher end of this range (the canonical pre-industrial doubling of about 550 ppm) might allow postponement of up to 15 years; stabilization at the lower end of this range (450 ppm) appears to be virtually impossible even if aggressive mitigation were to begin today.[7]

Several other important conclusions can be drawn from these results. First, panel b of Fig. 2 suggests that relatively small increases in the amount of delay yield relatively large increases in the stabilization frontier. Because each increment of delay decreases the number of options available to society, it seems natural to equate delay with lost "option value." On average, a 1-year delay increases the stabilization frontier by about 9 ppm, leading us to associate a negative atmospheric option value to delay of approximately 9 ppm/year, or conversely, to associate a *positive* atmospheric option value of the same amount to the deliberate advancement of mitigation.[8] It is perhaps worth contrasting the average option value of ~9 ppm $year^{-1}$ with the recent annual increase in atmospheric concentrations of ~2 ppm $year^{-1}$ (Keeling and Whorf 2005). The former is much larger because it is essentially an indicator of additional future commitment, assuming a reasonable amount of inertia in the energy system. Stated differently, option value is related to the integral of emissions, while the observed annual increase is related more closely to instantaneous emissions.

Secondly, panel b of Fig. 2 shows that MAI – or equivalently, option value – increases over time, meaning that as mitigation is postponed, the rate at which societal options

[6] Scenarios in which future emissions decline linearly, rather than exponentially with time are considered briefly in the concluding section.

[7] Stabilization at 450 ppm is still possible if (1) the rate of future emissions decline can exceed 1% $year^{-1}$ or (2) the global carbon sink is larger than we have anticipated. In Section 3.2, we show that a decline of ~1.5% $year^{-1}$ beginning today would be sufficient to achieve stabilization at 450 ppm. We briefly discuss issues related to possibility (2) in the conclusion.

[8] Note that the currency of option value in this paper is ppm not dollars. Also, an "average" option value is not terribly well-defined in this context. The implicit assumption here is that delays between 0 and 50 years represent the entire universe of possibilities. If we included scenarios that allowed for even greater delay, the "average" option value would be larger than 9 ppm $year^{-1}$.

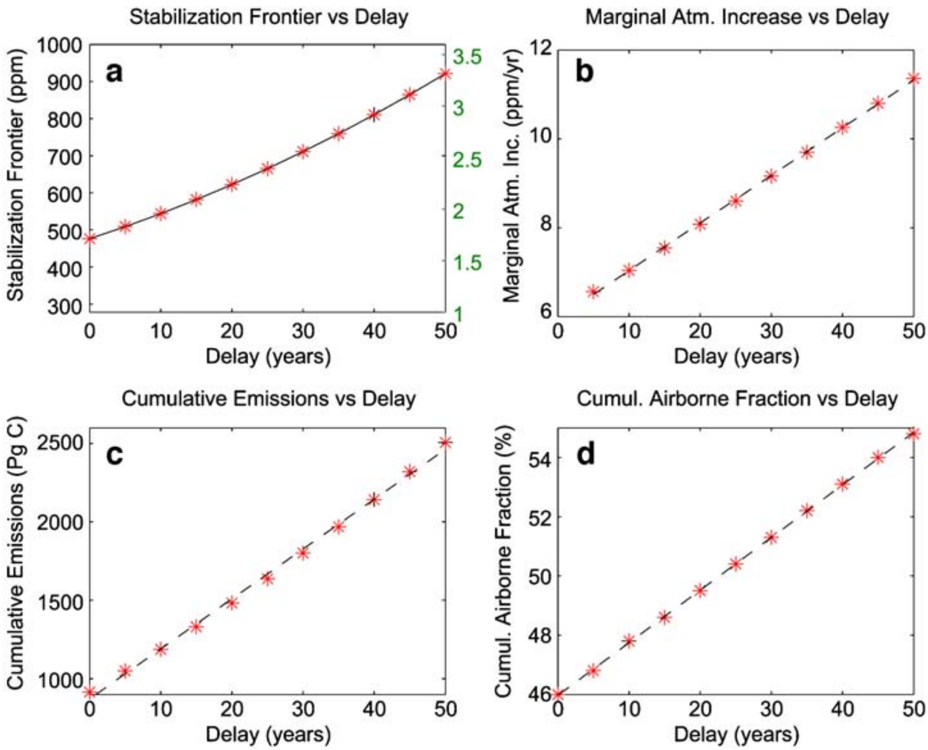

Fig. 2 Sensitivity of the atmospheric response to the length of mitigation postponement when emissions are assumed to decline at a fixed rate of 1% year^{-1}, as in Fig. 1. **a** shows the peak atmospheric concentration achieved – the stabilization "frontier" – as a function of delay time. The scale on the right-hand side of this panel indicates the factor by which atmospheric pCO_2 exceeds its pre-industrial value. The black curve represents a quadratic fit to the model output (*red points*). **b** shows the *marginal* increase in the stabilization frontier (in ppm per year) as a function of delay time, calculated by differencing consecutive points in **a**. The *dashed line* is a linear fit to the model output. **c** and **d** show the cumulative emissions release in the peak year (includes emissions released during the historical period) and the cumulative airborne fraction in the peak year, respectively, as a function of delay time, with the *dashed lines* providing linear fits to each

disappear actually increases. To see this, consider the atmospheric burden between two scenarios in which mitigation begins in the near term (say, delays of 0 and 5 years) and between two scenarios in which mitigation is postponed by several decades (e.g. delays of 40 and 45 years). The former scenarios reach peak concentrations of 475 and 509 ppm, yielding a difference of 34 ppm, while the latter reach peak concentrations of 811 and 865 ppm, yielding a difference of 54 ppm. We believe that these results add an additional degree of complexity to analyses that assign implicit or explicit benefits to delay. Normally, advocates of a "wait-and-see" approach appeal to the informational benefits of delay and to the avoidance of premature retirement of the energy capital stock. On the other hand, the cost of delay is clearly associated with the disappearance of relatively benign climate options that derive from low atmospheric CO_2 concentrations. Our results quantitatively link postponement to the foreclosure of options and suggest that the rate at which options disappear increases over time.

The linear increase in option value (panel b of Fig. 2), or equivalently the non-linear (ostensibly quadratic) dependence of the stabilization frontier on delay (panel a) reflects not only the specific modeling assumptions of this study (for example, the functional form of the emissions trajectory) but also the response of the ocean carbon system to increasing CO_2 concentrations.[9] A few additional insights can be gained by examining panels c and d of Fig. 2, which show, respectively, the *cumulative* emissions in the peak year and the *cumulative* airborne fraction in the same year as a function of delay (cumulative airborne fraction is defined as the ratio of the cumulative atmospheric accumulation to the cumulative emissions release). Since the increase in the atmospheric CO_2 concentration at any given time can be expressed as the product of the cumulative emissions release and the cumulative airborne fraction and since the stabilization frontier is defined as the concentration at the peak, it follows that the stabilization frontier should be linearly related to the product of the airborne fraction and the cumulative release in the peak year. Since both of these increase nearly linearly with delay, the stabilization frontier must increase approximately quadratically with delay, as seen in panel a of Fig. 2.[10]

To make this equivalence more concrete, consider the case of immediate mitigation. Panel c of Fig. 2 shows that the cumulative release by 2150 (the peak year) under this scenario is ~900 Pg C. This number includes contributions from both the historical release and the future release found by integrating the emissions curve from the present until 2150. Panel d shows that the cumulative airborne fraction in 2150 is approximately 46%, or that a little over 400 Pg C of the 900 Pg C release remains in the atmosphere at peak. Since each additional Pg C increases the atmospheric concentration by approximately 0.5 ppm, the additional atmospheric burden at peak must be approximately 200 ppm, which means that the frontier for immediate mitigation must be ~478 ppm (assuming pre-industrial concentrations of ~278 ppm). As noted above, this is indeed what we find in panel a of Fig. 2.

3.2 The tradeoff between timing and intensity of carbon mitigation

In the analysis above, we examined the sensitivity of the atmospheric response to the timing of mitigation, assuming that the rate of emissions decline was permanently fixed at 1% $year^{-1}$. While the qualitative result (i.e. the linearity of the option value curve when plotted as a function of delay) is robust to the choice of decline rate, the actual position of the frontier depends quite strongly on this parameter. In what follows, we extend the previous analysis to scenarios in which the rate of emissions decline varies between 0.5 and 3% $year^{-1}$, so that both the decline rate and delay time (but not the plateau period of 10 years) are assumed to be controllable by policy.[11] We find that when the assumed rate of decrease is small (~0.5% $year^{-1}$), relatively

[9] The convexity of the stabilization frontier in panel a of Fig. 2 also depends on the metric one uses to assess the environmental outcome. For example, if the outcome is evaluated in terms of radiative forcing (W/m^2) rather than in terms of CO_2 concentrations, then the frontier increases nearly linearly with the amount of delay (not shown), because forcing is a *concave* function of the CO_2 concentration. However, the key point – that the MAI is much greater than the instantaneous increase – still holds, because the large values for MAI result directly from the assumed inertia in the energy system.

[10] The linearity of the cumulative release can be extracted from the geometry of the emissions curves in Fig. 1a. The linearity of the airborne fraction results from the unique response of the surface ocean carbon chemistry system to increased CO_2 loading (see any textbook treatment of ocean carbon chemistry such as the one found in Zeebe and Wolf-Gladrow 2003).

[11] Several other implicit assumptions are worth making explicit. First, our analysis does not consider regional variations in the decline rate, and in this sense, the values we explore should be considered global averages. In reality, a given global decline rate might be achieved by combining more strenuous commitments from developed countries with less strenuous commitments from developing countries, as the FCCC envisioned. Secondly, the reported decline rates are enacted after emissions have already stabilized. Since the act of stabilizing emissions is itself a difficult feat, one should not underestimate the work required to achieve the one or two percent annual reductions discussed here, that, by themselves, might not sound ambitious.

modest changes in either of the policy levers (delay period or future decline rate) drive large changes in the stabilization frontier. On the other hand, when the assumed rate of decrease is greater than about 1.5% year^{-1}, the stabilization frontier is relatively insensitive to changes in these quantities. As a result, we conclude that future intensification of mitigation can offset current postponement, but only up to a point; delays of more than two or three decades permanently remove the option to stabilize atmospheric concentrations below a pre-industrial doubling.

Panel a of Fig. 3 shows the stabilization frontier as a function of delay under several different assumptions about the emissions decline rate (ranging from 0.5 to 2% year^{-1}), and panel b shows the marginal atmospheric increase (option value) as a function of delay for the same scenarios. Those that assume that mitigation takes the form of a 1% year^{-1} decline in emissions (open green circles) are thus identical to the scenarios shown in panels a and b of Fig. 2. While the qualitative response is similar in all cases, the frontier is consistently higher in the scenarios that assume lower rates of future emissions decline, a simple consequence of the fact that the cumulative release is much higher in these simulations. It is worth noting, however, that once the decline rate reaches about 1.5% year^{-1}, the marginal gains from further mitigation intensification are rather small, and the stabilization frontier converges toward a fixed lower bound.[12]

To provide a better sense of the potential role for the future intensification of mitigation, we show the same results as an explicit function of the decline rate in panels c and d of Fig. 3 for several different assumptions about delay. The marginal increases are reported as the increase in concentration (negative values imply decrease) associated with an increase in the decline rate of 0.1% year^{-1}. As above, we find that most of the differences occur when the decline rate is relatively small (<~1.5% year^{-1}). In a 30-year delay scenario (open black diamonds), increasing the decline rate from 0.5 to 1% year^{-1} decreases the frontier by more than 200 ppm (from 933 to 711 ppm), while increasing it from 1 to 1.5% year^{-1} decreases it by less than 100 ppm (from 711 to 640 ppm). The strong nonlinear response in these curves reflects the asymptotic approach to the frontier defined by instantaneous mitigation.

The interaction between the timing and intensity of mitigation is explored more fully in Fig. 4. Panel a explicitly shows the stabilization frontier as a function of the two policy levers, the delay time and future emissions decline rate. It is convenient to view the contour lines in this figure as “indifference curves” that show how current postponement and future intensification of mitigation can be traded off in such a way that the final outcome is preserved.[13] An inspection of this figure reveals two interesting facts. First, it is clear that near-term postponement can in fact be offset by increases in the intensity of mitigation later on (represented by movement along a contour toward the upper right in panel a).[14] However, it is also clear that, at relatively high rates of emissions decline, the marginal reduction in peak concentration associated with further increases in the decline rate becomes negligible, cautioning us against viewing future technology as a panacea for the additional accumulation that results from ongoing near-term postponement.

[12] As the decline rate gets very large, emissions effectively turn off after the plateau period and the frontier asymptotically approaches the concentration achieved at the end of the plateau period.

[13] An important subtlety is ignored here. Previous work has shown that future impacts may be sensitive to both the concentration endpoint *and* the rate at which this endpoint is approached (O’Neill and Oppenheimer 2004; Stocker and Schmittner 1997), implying that one would not be truly indifferent between two scenarios on the contour in Fig. 4a. However, we will ignore the implications of such differences here.

[14] The fact that many different emissions trajectories can lead to the same concentration endpoint has been well understood for at least a decade (Wigley et al. 1996) and probably for longer. More recently, O’Neill and Oppenheimer (2002) have explicitly examined the tradeoff between delays in mitigation and future rates of emissions decline, but they considered only a handful of possible scenarios.

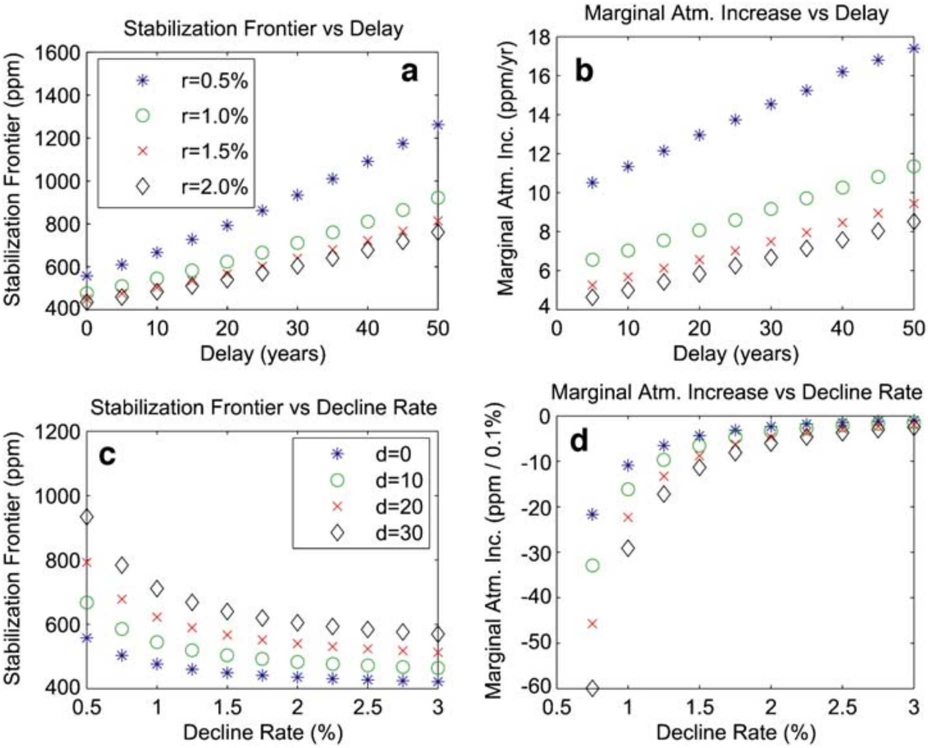

Fig. 3 Sensitivity of the atmospheric response when both the delay time and future decline rate are parameterized. **a** shows the stabilization frontier as a function of delay time under four different assumptions about the future decline rate (ranging from 0.5 to 2.0% year^{-1}), while **b** shows the marginal increase in the frontier for the same scenarios (the open green circles are identical to the results given in **a** and **b** of Fig. 2). **c** shows the stabilization frontier as a function of decline rate under four assumptions about delay (ranging from 0 to 30 years), while **d** gives the marginal increase (in ppm per 0.1% increase in decline rate) for the same scenarios. *Negative values* in **d** indicate that the peak atmospheric concentration decreases as the decline rate increases

As an example, consider atmospheric stabilization at 550 ppm. This atmospheric outcome could be achieved with fairly modest emissions reductions (declines of ~0.5% year^{-1}) if such reductions were to begin today. If the onset of mitigation were postponed by 15 years, however, then future declines would need to proceed at the considerably faster pace of 1.2–1.5% year^{-1} in order to achieve the same outcome, and if mitigation were delayed by 30 years, then stabilization at 550 ppm would effectively become impossible even if emissions could be made to fall at rates exceeding 3% year^{-1}, which seems challenging at best.[15] The emissions implications for other stabilization targets are shown in Table 1.

Panel b of Fig. 4 shows the marginal atmospheric increase from delay – the option value – as a function of both delay and emissions decline rate (effectively the partial derivative of

[15] The literature is relatively silent about what rates of decline are feasible, because such values cannot be accurately derived from theory or from historical analysis. Nevertheless, the range we consider is comparable to assumptions made in other studies. For example, Alcamo and Kreileman (1996) suggest that 2% year^{-1} is a "reasonable upper limit" for the future decline rate.

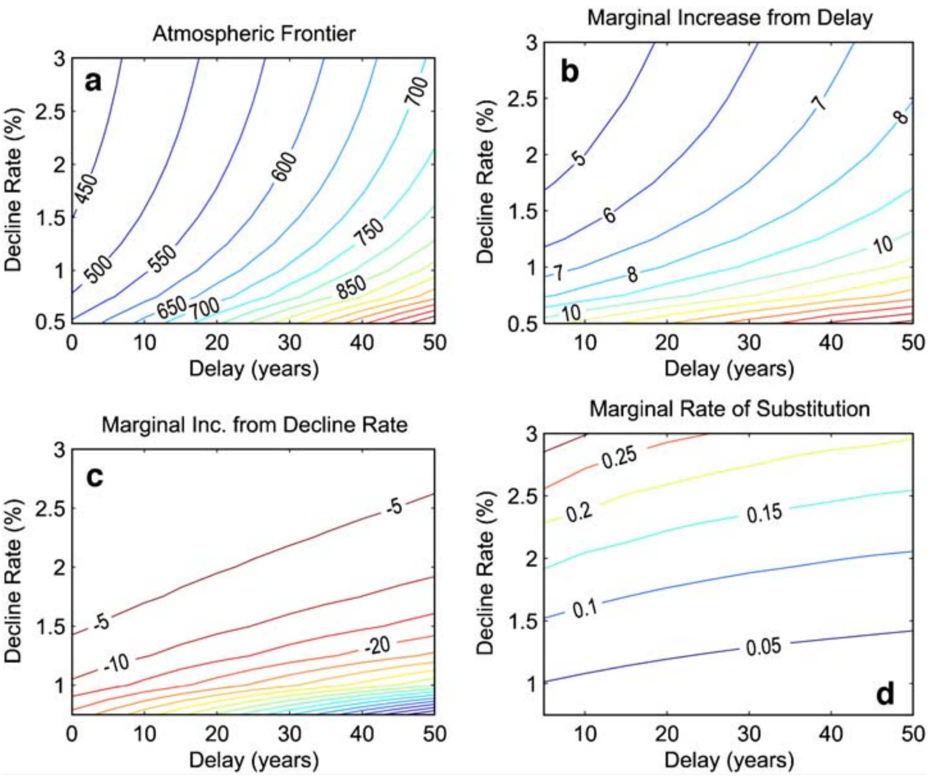

Fig. 4 Sensitivity of the stabilization frontier to delay time and emissions decline rate. **a** shows the atmospheric frontier (in ppm) as a function of both parameters, **b** shows the marginal increase from delay (in ppm year^{-1}) and **c** shows the marginal increase from decline rate changes (in ppm per 0.1% increase in decline rate). **d** gives the marginal rate of substitution between future mitigation and present delay (the slope of the contours in **a**). A value of 0.1 means that a 0.1% increase in decline rate would be needed to offset a 1-year delay in mitigation

panel a with respect to the delay variable). We find that the option value is largest when the decline rate is smallest, meaning that if future declines are constrained to be small, then small delays in mitigation will drive large increases in the stabilization frontier. For example, when emissions decline at the rate of 0.5% year^{-1}, the average option value is about 14 ppm (per year of delay), more than 50% greater than the option value when the decline proceeds at 1% year^{-1} (9 ppm per year of delay) and more than twice as great as the option value when the decline proceeds at 3% year^{-1} (6 ppm per year of delay).

Panel c shows the marginal atmospheric increase from increases in the future decline rate (the partial derivative of panel a with respect to the decline rate variable). Negative values indicate that future peak concentrations decrease when the decline rate increases, as one would expect. These changes are even more sensitive to assumptions about the decline rate. The average marginal decrease in going from 0.5 to 0.75% year^{-1} is 54 ppm (per 0.1% year^{-1} increase), in going from 0.75 to 1% year^{-1} it drops by more than 50% to 26 ppm and further increases in the growth rate (>2% year^{-1}) generally lead to changes under 10 ppm.

Panel d of Fig. 4 attempts to synthesize these results by providing a measure of how strongly changes in delay need to be compensated by changes in the decline rate to

Table 1 Emissions implications of various amounts of mitigation delay when concentration endpoints are imposed at 450, 550 or 650 ppm

Conc. target (ppm)	Equil T. (°C)	Delay (years)	Dec. rate (L) (% year^{-1})	Dec. rate (A) (% year^{-1})
450	1.5–3.9	0	1.75–2.0	1.25–1.5
		10	>3.0	>3.0
		20	>3.0	>3.0
		30	>3.0	>3.0
550	2.0–5.2	0	0.5–0.75	0.5–0.75
		10	1.0–1.25	0.75–1.0
		20	2.0–2.25	1.75–2.0
		30	>3.0	>3.0
650	2.5–6.2	0	<0.5	<0.5
		10	0.5–0.75	0.5–0.75
		20	1.0–1.25	0.75–1.0
		30	1.5–1.75	1.25–1.5

The equilibrium temperature response is taken from Watson et al. (2001). The decline rate indicates how rapidly future emissions must decline in order to meet the stabilization constraint when the carbon sink in our model is tuned to match the average (A) 1990s uptake from a set of recent studies (described in Mignone et al. 2006) or when the carbon sink is tuned to match the 1990s uptake from the lower limit (L) of these studies.

maintain a particular stabilization option. If the frontier (F) is a function of delay (t) and decline rate (r), then along a contour (F^*) of constant peak CO_2, we may write

$$F = F(t, r) = F^* \tag{1}$$

Differentiating with respect to t gives

$$\frac{\mathrm{d}F}{\mathrm{d}t} = \frac{\partial F}{\partial t} + \frac{\partial F}{\partial r} \times \frac{\mathrm{d}r}{\mathrm{d}t} = 0 \tag{2}$$

And upon dividing through by dF/dr and re-arranging, we find

$$\frac{\mathrm{d}r}{\mathrm{d}t} = \frac{\partial F / \partial t}{-\partial F / \partial r} \tag{3}$$

This gives an expression for the contour slope (dr/dt) of the curves in panel a as a function of the marginal increase from delay ($\partial F / \partial t$ as in panel b) and the marginal increase form decline rate ($\partial F / \partial r$ as in panel c). Since the latter is everywhere negative (increases in decline rate decrease the frontier), the expression for contour slope is everywhere positive, consistent with what we find in panel a. In what follows, we shall refer to dr/dt alternatively as the marginal rate of substitution (MRS) between future mitigation and near-term postponement, since it effectively measures the increase in decline rate that must accompany an increase in delay. The MRS is shown explicitly as a function of decline rate and delay in panel d. A value of 0.1 implies that an increase in delay of 1 year can be offset by an increase in the decline rate of 0.1% year^{-1}. The fact that the contours are nearly horizontal implies that the MRS is primarily a function of decline rate, a fact that is also apparent from direct inspection of the slopes of the contours in panel a. This result is consistent with the observation that the cumulative release is linearly related to the delay period but more strongly sensitive to the decline rate parameter.

4 Conclusions

In this paper, we have examined the implications of delays in mitigation using a simple, well-tested model of the ocean carbon cycle. We find that when future emissions are constrained to decline at the rate of 1% year^{-1}, the peak atmospheric concentration – the so-called stabilization frontier – increases, on average, at the rate of ~9 ppm year^{-1}. Since immediate mitigation would place the frontier near 475 ppm, our results imply that stabilization below a pre-industrial doubling would require a determined mitigation effort to begin within approximately the next decade. Secondly, we find that when both the delay time and the future decline rate are assumed to be flexible levers of policy, modest near-term postponement can be offset by future increases in the intensity of mitigation, but more significant postponement cannot be similarly offset by future increases, because the marginal rate of substitution between present and future mitigation becomes prohibitively large.

Given space constraints, we do not address in detail the sensitivity of our results to assumptions about the strength of the future carbon sink, to assumptions about (BaU) emissions growth in the near term, or to assumptions about the shape of possible mitigation trajectories in the future. However, a preliminary analysis (described below) suggests that changes in any of these assumptions do not significantly alter our fundamental qualitative conclusions, even though they do alter some of the specific quantitative results.

For example, we have run simulations in which the strength of the global carbon sink is diminished in such a way that ocean uptake in 1995 is about 0.4 Pg C lower than uptake in the same year in the baseline version (1.9 rather than 2.3 Pg C year^{-1}) and is roughly the lower bound of the published estimates discussed in Mignone et al. (2006). This exercise can be viewed as a generic attempt to evaluate the atmospheric impact of a weaker global carbon sink, but it should not be be viewed as an attempt to evaluate the response to a particular climate-carbon cycle feedback or as an attempt to diagnose the implications of uncertainty in the carbon budget more generally.

We find that the largest differences between our original and modified simulations occur when the delay is large or when the decline rate is small. For example, when a 10-year delay and 0.5% year^{-1} decline rate are assumed, the peak atmospheric concentration in the low-uptake version is 33 ppm higher than the peak concentration in the baseline version (700 vs 667 ppm). At 50 years, the difference between versions grows to 67 ppm (1,329 vs. 1,262 ppm). By contrast, when the assumed rate of decline is 2% year^{-1}, the difference between versions is 15 ppm for a delay of 10 years and 36 ppm for a delay of 50 years (see also Table 1). Given that many other plausible assumptions could be made about the response of the land and ocean sinks to future carbon loading, these results should be viewed as a first attempt to provide a mapping between changes in the carbon sink and changes in policy-relevant metrics, not as an attempt to bracket the full range of possible outcomes.

The sensitivity of our results to assumptions about near-term emissions growth can also be understood by examining simple perturbations to our existing scenarios. For example, consider the difference between an immediate mitigation scenario and a 10-year delay scenario. When emissions increase at the rate of 1 Pg C year^{-1} every 5 years, the emissions curve of the latter sits about 2 Pg C year^{-1} above the emissions curve of the former (Fig. 1, panel a). Since the peak occurs near 2150, the additional release is approximately 300 Pg C. With an airborne fraction of ~0.5, the additional burden is about 150 Pg C or about 70 ppm, and thus the marginal increase is about 7 ppm year^{-1}, as shown in Fig. 2, panel b. If the emissions growth rate were doubled, the emissions curve of the second scenario would sit roughly 4 Pg C year^{-1} above the emissions curve of the first, leading to twice the burden when differences in time to peak are neglected. This example demonstrates the strong sensitivity of our quantitative results to assumptions about near-term emissions growth. As mentioned previously, our emissions growth rate was chosen, in part, to approximate one of the central projections of the IPCC, the so-called A1B emissions scenario.

Finally, the sensitivity to assumptions about the shape of future mitigation trajectories can be understood by considering one plausible alternative. Suppose that once mitigation begins, future emissions decline at a constant absolute rate (0.2 Pg C year^{-1}) rather than at a constant relative rate (1% per year). We find that when the long exponential tail of emissions vanishes (in the former case), concentrations peak sooner and at a lower value, since less carbon is ultimately emitted to the atmosphere. For example, the maximum concentration in the 10-year delay scenario is 462 ppm when emissions decline linearly at 0.2 Pg C year^{-1} but 544 ppm when emissions decline exponentially at 1% per year. However, the marginal atmospheric increase grows with delay at nearly the same rate in each case (5–10 ppm per year when emissions decline at a constant absolute rate, as opposed to 6–11 ppm per year when emissions decline at a constant relative rate). It is worth noting that the scenarios discussed in Section 3.2 provide a better quantitative sense of how future mitigation intensification alters the position of the stabilization frontier and the relationship between MAI and delay.

These results highlight the need for continued research in several areas. First, greater mechanistic knowledge of the carbon cycle as well as greater certainty about land-use decisions and emissions growth from fossil fuel burning would help to constrain the exact position of the stabilization frontier. Secondly, our crude assumptions about the rate at which future emissions may decline suggests that additional knowledge about the costs and capabilities of technologies currently under consideration coupled with knowledge about how quickly such technological developments might be diffused in space and in time would allow us to more accurately include feasibility constraints on the rate of emissions decline and make it possible to determine the precise point at which various atmospheric options drop off the table.

Finally, our assumption that atmospheric stabilization is in some sense “optimal” essentially exogenizes the normative aspect of the problem, allowing us to derive “indifference curves” for various atmospheric outcomes in terms of near-term postponement and future increases in mitigation intensity (i.e. panel a of Fig. 4). We do not endorse particular targets or approaches in this paper, believing that the choice of concentration endpoint is a normative question that must be resolved by discussion among affected stakeholders (Dessai 2004) and that the choice of approach (i.e. the particular tradeoff between present and future mitigation that will lead to a given endpoint) is ultimately an economic question. However, if the cost-minimizing solution occurs when the demands on present and future mitigation are reasonably balanced (i.e. where the slope of the indifference curves in Fig. 4a is not too great or too small), then our results suggest that only modest delay should be tolerated when the world is committed to stabilization below a doubling.

Acknowledgements BKM gratefully acknowledges financial support from the MacArthur Foundation and the Princeton Environmental Institute and thanks B. O’Neill and two anonymous reviewers for helpful comments on the manuscript. In addition, all of the authors acknowledge financial support from BP and the Ford Motor Company under the Princeton Carbon Mitigation Initiative. RHS and BKM also thank H. Lam for several useful discussions.

References

Alcamo J, Kreileman E (1996) Emission scenarios and global climate protection. Glob Environ Change 6:305–334

Corfee-Morlot J, Hohne N (2003) Climate change: long-term targets and short-term commitments. Glob Environ Change 13:277–293

Dessai S (2004) Defining and experiencing dangerous climate change. Clim Change 64:11–25
Gregory JM, Huybrechts P, Raper SCB (2004) Threatened loss of the Greenland ice-sheet. Nature 428:616
Ha-Duong M, Grubb MJ, Hourcade J-C (1997) Influence of socioeconomic inertia and uncertainty on optimal CO_2-emission abatement. Nature 390:270–273
Hansen J (2004) Defusing the global warming time bomb. Sci Am 290:68–77
Hansen J (2005) A slippery slope: how much global warming constitutes "dangerous anthropogenic interference"? Clim Change 68:269–279
Hoegh-Guldberg O (1999) Climate change, coral bleaching and the future of the world's coral reefs. Mar Freshw Res 50:839–866
Houghton JT et al (ed) (2001) Climate change 2001: the scientific basis. Cambridge University Press, Cambridge, UK
Keeling CD, Whorf TP (2005) Atmospheric CO_2 records from sites in the SIO air sampling network. In: Trends: a compendium of data on global change, carbon dioxide information Analysis Center. Oak Ridge National Laboratory, US Department of Energy, Oak Ridge, TN. Digital data available at http://cdiac.ornl.gov/trends/co2/sio-mlo.htm
Mignone BK, Gnanadesikan A, Sarmiento JL, Slater RD (2006) Central role of Southern Hemisphere winds and eddies in modulating the oceanic uptake of anthropogenic carbon. Geophys Res Lett 33:L01604
O'Neill BC, Oppenheimer M (2002) Dangerous climate impacts and the Kyoto Protocol. Science 296:1971–1972
O'Neill BC, Oppenheimer M (2004) Climate change impacts are sensitive to the concentration stabilization path. Proc Nat Acad Sci USA 101:16411–16416
Oppenheimer M (1998) Global warming and the stability of the West Antarctic ice sheet. Nature 393:325–332
Oppenheimer M (2005) Defining dangerous anthropogenic interference: the role of science, the limits of science. Risk Anal 25:1399–1407
Oppenheimer M, Alley RB (2004) The West Antarctic ice sheet and long-term climate policy. Clim Change 64:1–10
Oppenheimer M, Petsonk A (2005) Article 2 of the UNFCCC: historical origins, recent interpretations. Clim Change 73:195–226
Orr JC et al (2000) Abiotic how-to document. Text available at http://www.ipsl.jussieu.fr/OCMIP
Pacala S, Socolow R (2004) Stabilization wedges: solving the climate problem for the next 50 years with current technologies. Science 305:968–972
Richels R, Edmonds J (1995) The economics of stabilizing atmospheric CO_2 concentrations. Energy Policy 23:373–378
Shafer G, Sarmiento JL (1995) Biological cycling in the global ocean 1. A new analytical model with continuous vertical resolution and high-latitude dynamics. J Geophys Res 100:2659–2672
Siegenthaler U, Joos F (1992) Use of a simple model for studying ocean tracer distributions and the global carbon cycle. Tellus 44B:186–207
Socolow RH, Lam SH (2007) Good enough tools for global warming policy making. Philos Trans R Soc Lond A 365:897–934
Stocker TF, Schmittner A (1997) Influence of CO2 emission rates on the stability of the thermohaline circulation. Nature 388:862–865
United Nations Framework Convention on Climate Change (UNFCCC) (1992) Text available at http://www.unfccc.int
Wanninkhof R (1992) Relationship between wind speed and gas exchange over the ocean. J Geophys Res 97:7373–7382
Watson RT et al (2001) Climate change 2001: the synthesis report. Cambridge University Press, Cambridge, UK
Webster M (2002) The curious role of learning in climate policy: should we wait for more data? Energy J 23:97–119
Wigley TML, Richels R, Edmonds JA (1996) Economic and environmental choices in the stabilization of atmospheric CO_2 concentrations. Nature 379:240–243
Yohe G, Andronova S, Schlesinger M (2004) To hedge or not against an uncertain climate future? Science 306:416–417
Zeebe RE, Wolf-Gladrow D (2003) CO_2 in seawater: equilibirum, kinetics, isotopes. Elsevier, Amsterdam