The Impacts of Climate Variability on Near-Term Policy Choices and the Value of Information

Robert J. Lempert[a], Michael E. Schlesinger[b],
Steven C. Bankes[a] and Natalia G Andronova[b]

a RAND, 1700 Main St., Santa Monica, CA 90407-2138.

b Department of Atmospheric Sciences, University of Illinois at Urbana-Champaign, Urbana, IL 61801.

Submitted to Climatic Change

April 1998

ABSTRACT

Variability is one of the most salient features of the earth's climate, yet quantitative policy studies have generally ignored the impact of variability on society's best choice of climate-change policy. This omission is troubling because an adaptive emissions-reduction strategy, one that adjusts abatement rates over time based on observations of damages and abatement costs, should perform much better against extreme uncertainty than static, best-estimate policies. However, climate variability can strongly affect the success of adaptive-abatement strategies by masking adverse trends or fooling society into taking too strong an action. This study compares the performance of a wide variety of adaptive greenhouse-gas-abatement strategies against a broad range of plausible future climate-change scenarios. We find that: i) adaptive strategies remain preferable to static, best-estimate policies even with very large levels of climate variability; ii) the most robust strategies are innovation sensitive, that is, adjust future emissions reduction rates on the basis of small changes in observed abatement costs but only for large changes in observed damages; and iii) information about the size of the variability is about a third to an eighth as valuable as information determining the value of the key parameters that represent the long-term, future climate-change state-of-the-world.

1. INTRODUCTION

With the extreme uncertainty posed by the threat of climate change, policies for abatement of greenhouse-gas emissions should be adaptive. That is, decision-makers should pursue abatement policies with the expectation that they or their successors will make midcourse corrections based on observations of the relevant environmental and economic systems. A diverse body of theoretical literature supports this conclusion, including the sequential decision-making work of Manne and Richels (1992), the stochastic optimization work of Kelly and Kolstad (1996), and our own exploratory modeling studies [Lempert, Schlesinger and Bankes (1996), henceforth LSB]. Adaptive decision-making is also declared policy[1], as well as the actual practice of international and domestic policies.

Climate variability, however, poses a challenge to the design of adaptive-decision strategies. Variability is one of the most salient features of the Earth's climate. As described by other articles in this volume, global and regional temperatures, precipitation and storm patterns all exhibit natural fluctuations over many time-scales, from year-to-year changes, to variations over decades, centuries, and millennia. This variability can complicate the identification of anthropogenic climate change. Since the climate varies, it is not immediately clear whether or not any specific observed changes are due to natural causes or human intervention. Thus, variability can degrade the performance of adaptive-decision strategies by reducing their ability to learn quickly from observations. Variability can fool society into taking premature action because a variation is mistaken for a trend, or can mask an adverse trend until it is too late to respond effectively.

This study is an initial attempt to examine the impact of climate variability on the effectiveness and design of adaptive-decision strategies for greenhouse-gas abatement. Despite its importance, this question has largely been ignored in the analytic policy literature, largely because the commonly-used tools for decision-making under uncertainty are poorly equipped to address the impact of climate variability on adaptive-decision strategies. Stochastic optimization techniques provide an integrated framework that includes endogenous learning, that is, learning based on the observations an adaptive-decision strategy makes of the time-evolution of the climate and economic systems. However, stochastic optimization techniques require restrictive assumptions on the mathematics that describe the phenomena being studied. Sequential-decision methods allow a richer mathematical description of potential decision strategies and the climate and economic systems. Sequential-decision strategies, however, are designed to consider exogenous information, that is, information obtained without reference to the time-evolution of the climate system.

In this study we employ a new analytic framework for decision-making under conditions of extreme uncertainty – exploratory modeling – that is well-suited to addressing the impacts of climate variability on the choice of abatement policies for greenhouse-gas emissions. Rather than focus on a single best-estimate of the future, expressed as either a point prediction or a known probability distribution, we assume that there are a wide range of plausible futures of unknown likelihood. We simulate the performance of a number of alternative decision strategies over a wide range of these plausible futures and then search for strategies that best meet decision-makers goals (Bankes, 1993). Exploratory modeling combines some of the best features of both quantitative cost-benefit analysis with those of scenario-based planning methods (Schwartz, 1996). Exploratory modeling is a useful framework for addressing the impacts of climate variability on policy decisions because it can easily employ models of non-linear phenomena and provides a means for extracting useful policy conclusions, even when we are uncertain about important features of the problem.

In LSB we used the exploratory modeling approach[2] to compare the performance of a single, simple adaptive-decision strategy to the two 'best-estimate' policies, 'Do-a-Little' and 'Emissions-Stabilization', that are commonly advocated in today's climate change debate. As the names imply, the “Do-a-Little” policy has no near-term emissions reductions and is similar to that advocated by many opponents of the commitments negotiated at the Conference of Parties in Kyoto in December 1997. The ‘Emissions-Stabilization” policy returns and holds global emissions close to their 1990 levels through 2060 and is similar to the policies favored by many advocates of the Kyoto agreement. Using a linked system of simple climate and economic models, we compare the performance of these strategies against a wide range of plausible climate-change futures, including the possibility of large and abrupt climate change, technology breakthroughs that radically reduce abatement costs, and the possibility that climate change turns out not to be a significant problem. We found that even a very simple adaptive-decision strategy on average significantly outperforms either of the best-estimate policies unless society is highly certain about the climate-change future (on the order of 95% confident about key parameter values). This result is not surprising since the 'Do-a-Little' and 'Emissions-Stabilization' policies perform well if their underlying assumptions turn out to be valid, but can fail severely in those cases where their assumptions turn out to be wrong. In contrast, the adaptive-decision strategy can make midcourse corrections and avoid significant errors.

This previous work did not, however, consider the effects of climate variability. In this study we add a simple representation of climate variability and the damages due to this variability to the simulation models used in LSB. We also consider several thousand alternative adaptive-decision strategies, rather than the single adaptive-decision strategy of LSB, which are distinguished by differing choices for such factors as the rate of near-term emissions reductions and the sensitivity to trends in the observed damage and abatement-cost time series. We compare the performance of these alternative strategies against a large number of plausible climate-change futures to address three key policy questions. First, we want to verify that the conclusion of LSB, that adaptive-decision strategies for climate change dominate non-adaptive ones, still holds in the presence of climate variability. Second, we want to understand how expectations about the impacts of future climate variability should affect a policy-maker's choice of an adaptive-decision strategy, in particular, the near-term features of these strategies. Third, we want to estimate the value of exogenous information about the variability in comparison to information about other important climate and economic variables.

As one of the first studies in this area, this study provides only preliminary answers to these questions. This study employs a crude treatment of the factors -- such as climate variability, the damages due to variability, and adaptive-decision strategies that deal with variability -- that are most critical to addressing these questions. For instance, we model climate variability by red-like noise and the damages due to climate variability by a simple phenomenological damage function. We also consider, similarly to LSB, two-period adaptive-decision strategies that can make only a single midcourse correction, rather than consider more flexible and sophisticated adaptive strategies that can make multiple adjustments in abatement rates over the next century. As we will discuss in detail later, these simplifications may affect our judgments about the effectiveness of alternative adaptive-decision strategies.

Nonetheless, we believe we can draw a variety of interesting and important conclusions based on this preliminary work. First, we confirm that the results of LSB -- that adaptive-decision-strategies dominate prescriptive policies -- are valid even in the presence of the climate variability considered in this study. Second, we show that policy-makers can choose among a range of adaptive strategies that perform on average nearly equally well, but differ in their balance between near-term abatement rate and the rate at which the strategy tolerates false alarms due to misreading climate trends. In general, the slower the near-term reductions, the less conservative a successful strategy should be in its criteria for responding to observations. Third, we find that the most robust adaptive-decision strategies in the face of climate variability tend to be innovation-sensitive rather than damage-sensitive. That is, these strategies will change their emissions abatement rates in response to small changes in abatement costs and only very large changes in the damages. Fourth, we find that the value of exogenous information about the magnitude of the climate variability is worth about a third to an eighth as much as exogenous information that definitively determines the optimum long-term climate change policy. Finally, we believe that this initial work provides a powerful framework for further studies to determine the best climate-change emissions-abatement policies in the presence of climate variability.

2. DEFINING AN UNCERTAINTY SPACE

In order to address the impacts of climate variability on the effectiveness and design of adaptive-decision strategies under conditions of extreme uncertainty, we will compare the performance of a large number of potential strategies against a large number of plausible climate-change futures. Accordingly, we need to specify a model and a set of input parameters that define a plausible set of future states of the world.

Traditionally, one uses available data to fix best estimates and/or probability distributions for such parameters and then predicts the best policy with the resulting model. However, the extreme uncertainty surrounding key aspects of the climate-change problem, such as the variability and the damages due to climate change, can make this approach unreliable. Instead, we look for sets of input parameters that give model outputs consistent with past trends and other available information, and then use visualization and search strategies to examine a variety of adaptive-decision strategies across these sets of plausible parameters in order to: i) craft strategies that are robust against the uncertainties, and/or ii) isolate the key uncertainties on which policy-makers should focus.

Similarly to our previous work [LSB; Lempert, Schlesinger and Hammitt (1994), henceforth LSH; and Hammitt, Lempert and Schlesinger (1992), henceforth HLS], we consider a linked system of simple climate and economic models designed to compare the performance of adaptive strategies against a wide range of plausible climate scenarios. We use our energy-balance-climate/upwelling-diffusion-ocean (EBC/UDO) model (Schlesinger and Jiang, 1991) to simulate the change in global-mean surface temperature, as a function of the climate sensitivity ΔT_{2x}, due to anthropogenic emissions of greenhouse-gas (GHGs) and sulfur dioxide (SO_2). The basecase emissions of these gases are given by the IS92a scenario (Leggett et. al. 1992, henceforth IPCC92), reduced by logistic diffusion at some policy-determined rate 1/R of "fuel-switching" technologies that decrease the emissions intensity of society's energy-related capital stock. Our abatement-cost model (LSB, LSH, HLS) tracks the emitting capital that must be prematurely retired (assuming a normal lifetime of 30 years) due to emissions constraints, and determines the incremental costs of non-emitting capital by $K(t) = \kappa_o + (\kappa_1 - \kappa_o)(1-d)^{t-t_{tech}}$, where κ_1 and κ_o are the projected costs of emitting and non-emitting stock (Manne and Richels, 1991; Nordhaus, 1991), d is the rate at which innovation reduces the incremental cost of the emitting stock, and t_{tech} is the year in which technological innovation begins.

The models used in this study have two key additions to those in our previous work, a representation of climate variability and the damages due to climate variability, which we describe here. A detailed description of the complete model can be found in Lempert, Schlesinger, Bankes and Andronova (1998).

2.1. Climate Variability

In order to represent climate variability, we write the radiative forcing as

$$\Delta Q(t) = 6.3334 \ln\left(\frac{ECD(t)}{ECD(1765)}\right) + \Delta F_{SO_4}\left(\frac{E_{SO_2}(t)}{E_{SO_2}(1990)}\right) + g(t) \quad , \qquad (1)$$

where $ECD(t)$ is the effective carbon dioxide concentration that would give the same forcing as the actual concentration of carbon dioxide, methane, and other greenhouse gases; $E_{SO_2}(t)$ is the emission rate of SO_2 which is converted to sulfate (SO_4) aerosols in the atmosphere; and ΔF_{SO_4} is the sulfate-aerosol radiative forcing in 1990. Climate variability is given by the Gaussian-distributed noise $g(t) \approx n(0, \sigma_Q)$. Schlesinger and Andronova (1997) have calculated a large set of climate sensitivity/sulfate forcing pairs, using a bootstrap method with their energy-balance-climate/ upwelling diffusion-ocean model, that reproduce the instrumental temperature record from 1856 to 1995. As in LSB, we choose a range of climate sensitivities, $0.5°C \leq \Delta T_{2x} \leq 4.5°C$; the resulting sensitivity/sulfate-forcing pairs are shown by the cloud of gray dots in Figure 1. The bootstrap calculations also produce a small number of cases with climate sensitivities great than 4.5°C, which are not shown here.

There are, however, many potential sources of error that could affect the calculation of these $\left(\Delta T_{2x}, \Delta F_{SO_4}\right)$ pairs. For instance, volcanoes or solar oscillations could cause the actual sensitivity/sulfate pair to fall outside the point cloud shown in Figure 1.[3] It is thus convenient to summarize the bootstrap results with a simple analytic expression and write the sulfate forcing parameter as

$$\Delta F_{SO_4}(1990) = \begin{cases} -0.78\gamma_{SO_4}\sqrt[3]{\Delta T_{2x} - 1.8°C} & \text{for } \Delta T_{2x} \geq 1.8°C \\ 0 & \text{for } \Delta T_{2x} < 1.8°C \end{cases}, \qquad (2)$$

where γ_{SO_4} represents the uncertainty in the sulfate forcing. The value $\gamma_{SO_4} = 1$ reproduces the bootstrap results.

For each of a large number of $\left(\Delta T_{2x}, \gamma_{SO_4}\right)$ values, we next find the best-estimate of the associated white-noise climate forcing, σ_Q, by regressing our climate model using each pair against the instrumental temperature record. As shown in Figure 1, the estimate of the noise varies slowly for $\Delta T_{2x} \geq 1.5°C$ but increases rapidly for $\Delta T_{2x} < 1.5°C$. The noise is smallest for sensitivity/sulfate pairs near the point cloud.

This white-noise forcing model does a good job of representing interdecadal as well as interannual climate variability. We do not consider uncertainty in the noise parameter; rather as described below we consider a wide range of parameters describing the effects of the climate variability, due to σ_Q, on the damages due to climate.

We also model the emission rate of anthropogenic sulfur dioxide in Eq. (1) by

$$E_{SO_2}(t) = \begin{cases} E_{SO_2,obs}(t) & \text{for } 1861 \leq t \leq 1994 \\ \left(1 - r_{SO_2}\right)^{t-1995} \eta(t)\, F_{CO_2,\lambda}(t) & \text{for } t > 1994 \end{cases} \quad , \qquad (3)$$

where $E_{SO_2,obs}(t)$ is the anthropogenic emission rate of sulfur in the form of SO_2 given by IPCC92 and updated in IPCC 1995, $\eta(t)$ is the ratio of SO_2 emissions to carbon dioxide emissions in the IS92a estimate and r_{SO_2} is the rate at which sulfur-dioxide emissions may deviate from the IS92a estimate. We include r_{SO_2}, the deviation from the IPCC SO_2emissions trajectory, because sulfur dioxide has significant environmental impacts, such as acid rain, independent of its role in climate change, and thus may be subject to different policy pressures and technological innovation than greenhouse gases. We use two value of the sulfur-dioxide reduction rate, $r_{SO_2} = 0\%$ which reproduces the IPCC trajectories, and $r_{SO_2} = 2\%$ which gives sulfur-dioxide emissions roughly constant at 1995 levels over the next thirty years. This later case is similar to those considered by Wigley, Richels and Edmonds (1996) and others.

2.2. Damages Due to Variability

In this study we use a simple, phenomenological damage function designed to capture, in aggregate, some of the impacts of climate variability and the ability of society to adapt to changes in variability. We write the annual damage in year t as

$$D(t) = \alpha_1 \left[\frac{\Delta\overline{T}_5(t)}{3°C} \right]^{\eta_1} + \alpha_2 \left[\frac{\Delta T(t) - \Delta\overline{T}_5(t)}{0.15°C} \right]^{\eta_2} + \alpha_3 \left[\frac{\Delta T(t) - \Delta\overline{T}_{30}(t)}{0.35°C} \right]^{\eta_3} \quad , \quad (4)$$

where $\Delta T(t)$ is the annual global-mean surface temperature change, and $\Delta\overline{T}_5(t)$ and $\Delta\overline{T}_{30}(t)$ are the 5-year and 30-year running averages of $\Delta T(t)$. The second and third terms on the right-hand side of Eq. (4) represent the damages due the variability of climate. The second term represents those impacts due to changes in the variability of the climate system that society and ecosystems can adapt to on the time-scale of a year or two. The third term represents those impacts that society and ecosystems adapt to more slowly, on the order of a few decades, and thus are sensitive to both the year-to-year variability and the secularly increasing trend in temperature. The first term represents the damages due to a change in the global-mean surface temperature and is similar to the power-law functions used in most simple damage models in the literature. In this study we can view this term as representing impacts that society and ecosystems adapt to on century-long time scales.[4]

We can better understand the effects of these two variability terms by examining the stochastic properties of the global-mean temperature as generated by our energy-balance-climate/upwelling-diffusion-ocean model for the forcing of Eq. (1). Figure 2 shows the cumulative probability distributions for $\Delta T(t)$ and the three terms in Eq. (4), $\Delta\overline{T}_5(t)$, $\Delta T(t)-\Delta\overline{T}_5(t)$, and $\Delta T(t)-\Delta\overline{T}_{30}(t)$, for the year t = 1995, over an ensemble of 500 Monte Carlo runs. The curves in this figure were calculated using $\left(\Delta T_{2x}, \Delta F_{SO_4}(1990), \sigma_Q\right) = \left(2.5°C, -0.7\ W/m^2, 3.2\ W/m^2\right)$. The curves for other values of the climate parameters are similar. Note that the average global-mean temperature change $\Delta T(t)$ varies by 0.5°C between its 5% and 95% confidence levels, consistent with the up to 0.4°C annual variation in the instrumental temperature record from which our noise term σ_Q was derived. The distribution for the five-year average $\Delta\overline{T}_5(t)$ has a somewhat smaller range, depressed by about 0.1°C at the 95% confidence level, due to the effects of the secularly increasing trend in temperature. The cumulative probability curves for $\Delta T(1995)-\Delta\overline{T}_5(1995)$ and

$\Delta T(1995) - \Delta\overline{T}_{30}(1995)$ are nearly parallel, but the latter has its 50% likelihood level at 0.2°C while the former has a 50% likelihood at the smaller value of 0.05°C. Thus we see the $\Delta T(t) - \Delta\overline{T}_5(t)$ term responds mostly to the effects of the year-to-year variability, while the $\Delta T(1995) - \Delta\overline{T}_{30}(1995)$ term is sensitive to both the year-to-year variability and the secularly increasing trend in temperature.

Little systematic information is available to constrain the choice of values for the damage function parameters in Eq. (4). A large body of research on impacts due to climate change suggests that damages due to climate change would not be larger than about 1% to 2% of GWP. Other researchers argue, however, that the impacts might be significantly larger, either because of effects inadequately treated in most analyses or because future generations may place a high value on non-economic loses, such as changed ecosystems, due to climate change. At least some of the political concern about the climate-change problem is motivated by worry about the risk of such unpredicted and potentially severe damages. In this study we use a very wide range of plausible damage estimates in order to support the argument that a simple adaptive-decision strategy can be robust against both very small and very large damages.

We can place very rough constraints on the parameters in Eq. (4) by requiring the function to be consistent with past observations. We first note that year-to-year the damages associated with climate-related phenomenon such as El Niños and large-scale natural disasters are on the order of a few tenths of a percent of GWP. Thus, whatever damages due to climate change have occurred in the last few years and decades, they cannot have been more than a few tenths of a percent of GWP, otherwise we would have observed unambiguous evidence of damages to date. This information places only loose constraints on the first term in Eq. (4). With $\Delta\overline{T}_5(1995) = 0.5°C$, we can write $\alpha_1 \leq 0.1\% \cdot 6^{\eta_1}$. As in LSB, we choose a cubic term for the damages due to the change in temperature, $\eta_1 = 3$, which corresponds to a range for the damage coefficient $0\% \leq \alpha_1 \leq 20\%$ GWP.

We can also use available time series data on economic losses from large-scale natural disasters to place rough constraints on the variability coefficients and exponents, α_2, α_3, η_2 and η_3. In recent years the number of large-scale natural disasters causing widespread economic loss has been on the rise, though the year-to-year variation in damages varies considerably (Munich Re Reinsurance, 1997). Since 1960 such disasters have shown only a small trend upwards, measured as a fraction of GWP. In 1996 over 500 large-scale events (excluding earthquakes) caused $60 billion (0.2% GWP) in damages, one of the largest years on record.[5] We find that the cumulative probability distribution produced by the last two terms in Eq. (4), using the parameters $\alpha_2 + \alpha_3 = 0.2\%$ GWP and $\eta_2 = \eta_3 = 1$, provides a reasonable fit to the 35 data points of the Munich Re time series of economic losses due to large-scale natural disasters.

Figure 3 compares the Munich Re time data to the distribution of damages generated by our model over an ensemble of 500 Monte Carlo runs for the year 1995 using the climate parameters $\left(\Delta T_{2x}, \gamma_{SO_4}, \sigma_Q\right) = \left(2.5°C, 1, 3.2\ W/m^2\right)$ and three sets of damage parameters. We define "Low

Variability" as the parameters $(\alpha_2,\alpha_3,\eta_2,\eta_3) = (0.2\%, 0\%, 1, \mathrm{na})$ which, as seen in Figure 3, closely matches the Munich Re data. We define the "High Variability" and "Increasing Variability" cases to have parameters $(0.4\%, 0\%, 2, \mathrm{na})$ and $(0\%, 0.33\%, \mathrm{na}, 3)$, respectively. The resulting damage distributions are nearly identical in 1995 and have the same 50% likelihood value as the Munich Re data, though the damages for high cumulative probabilities are much larger. We argue that the 'High' and 'Increasing' damage cases are plausible fits to the observed data because: i) this extreme-event data may significantly underestimate the damages due to climate variability, or ii) chance may have spared society over the last 35 years the low-probability/high-impact damages represented by the right-hand side of the curves in Figure 3. The damages in the "High Variability" case are largely insensitive to increasing concentrations of greenhouse gases, and thus cannot be affected by any emissions abatement policy. The damages in the "Increasing Variability" case grow with greenhouse-gas concentrations and thus can be affected by policy choices. Figure 3 also shows the damages in 2020 for the "Increasing Variability" case given the IS92a emissions.

Our damage model does have important shortcomings. Among the most important for this study is that the white-noise forcing, the driver of the variability in our model, is a fit to the instrumental temperature record and does not change as we run our simulations into the future. Thus, the damage distribution due to the $\Delta T(t) - \Delta\overline{T}_5(t)$ variability term does not change over time and the damage distribution due to the $\Delta T(t) - \Delta\overline{T}_{30}(t)$ term increases only due to increases in the rate of change in the 30-year average temperature, though in fact we expect that changes in variability are more important than changes in the mean (Katz and Brown, 1992; Mendelsohn et. al., 1994). In addition, the extreme-event data used to constrain our damage-function parameters are due at least in part to trends in society's susceptibility to natural disasters rather than to any change in the size or frequency of the natural disasters themselves and, conversely, are only one component of the damages due to climate change. Nonetheless the crude phenomenological damage function in Eq. (4) provides a sufficient foundation to support our initial explorations of alternative abatement strategies and the impacts of variability on near-term policy choices.

3.4 Scenarios

So far we have defined a model and identified the space of input parameters such that model outputs are consistent with available data. Given our computational limitations,[6] we now need to choose a set of about 60 uncertainty-space points that reasonably span the space of all plausible points, consistent with our purpose of comparing the performance of potential adaptive-decision strategies. Table I shows the points we use, chosen to best address the key policy questions posed in this study. We consider 20 different combinations of climate sensitivity, sulfate forcing, innovation parameters, and damage coefficient α_1, where each point is labeled with a vector $\left(\Delta T_{2x}, \gamma_{SO_4}, \alpha_1, d\right)$. For each of these uncertainty space points we consider three damages due to variability cases, "Low", "High" and "Increasing".

We use two criteria in selecting this sample of uncertainty-space points. First, we demand that the points sample the full range of plausible values for each parameter. For instance, among the 20 points shown in Table I there is one instance of ΔT_{2x} = 0.5°C, three instances of 1.5°C, ten of 2.5°C, four of 3.5°C, and two of 4.5°C; seven instances of the innovation parameters, d = 0% and 2%, and six instances of d= 5%; and six different values of α_1: three instances each of 0%, 3.5%, 5%, four instances of 7%; five instances of 2%; and two of 10%. We also include two values of the sulfate forcing, γ_{SO_4} = 0.25 and 1.0, for each value of the climate sensitivity greater than 1.5°C. (For ΔT_{2x} = 0.5°C and 1.5°C, Eq. (2) gives the same sulfate forcing $\Delta F_{SO_4}(1990)$ for any value of γ_{SO_4}.)

Second, we demand that our sample of uncertainty-space points be spread relatively uniformly across a large range of different future states of the world. It is convenient to group such futures according to the optimum climate-change policy society would follow if it had perfect information. In LSB we considered the performance of three static (non-adaptive) policies: 'Do-a-Little', similar to the policy favored by many climate-change skeptics; 'Emissions-Stabilization Policy', which returns and holds global emissions to their 1990 levels through 2060; and a Drastic Reductions Policy of immediate and draconian emissions reductions that would be appropriate for some combination of very large damages due to climate change and very rapid innovation-induced reductions in abatement costs. In this study we compare the performance of these three static policies for each of the 60 points and label each point with the lowest cost of the three policies. These groupings are shown in Table I. The 60 uncertainty-space points give 15, 20, and 25 instances of DAL, ES and DR futures, respectively. We believe that the results presented in the next section are relatively insensitive to the precise choice of uncertainty-space points considered here.

3. CHOOSING A STRATEGY

We can now address the problem of a decision-maker trying to craft a policy response to the threat of climate change in the face of climate variability and extreme uncertainty about this variability and the costs and benefits of various climate-change strategies. In the language of Section 2, the decision-maker is uncertain about which of the uncertainty-space points is the most accurate representation of the future or what probability distribution to lay across these points. In reality the "decision-maker" is also a collection of many individuals who hold a wide variety of very different expectations about the relative likelihood of these uncertainty-space points. In the presence of this uncertainty and of climate variability, should the decision-maker use an adaptive strategy? If so, what should be the near-term features of such a strategy? What information would be most valuable in improving the performance of an adaptive-decision strategy?

3.1. When to Choose an Adaptive Strategy

In LSB we compared the performance of a single adaptive-decision strategy to that of two static policies as a function of a decision-maker's expectations about the climate-change future. We found that in the absence of climate variability even a very simple adaptive strategy dominated the static policies unless one was virtually certain about the future state of the world. In this work we compare the performance of a large number of adaptive-decision strategies to each other and to the performance of two static policies. We find that in the presence of climate variability, adaptive-decision strategies generally dominate static policies, but that in some cases the variability so degrades the performance of the adaptive strategies that a decision-maker should choose a static policy.

We consider simple, two-period, threshold-based adaptive-decision strategies similar to those used in LSB that can respond to a policy-maker's estimate of any trend in damages based on annual observations of the noisy time series D(t). We calculate this estimate, $D_{est}(t)$, using a linear, discrete-time Kalman filter (Lewis, 1986) -- a Bayesian estimator that rapidly detects any statistically significant trend in the damage time series. As shown in Figure 4, our adaptive strategies begin with a pre-determined abatement rate $1/R_1$ and can switch to a second-period abatement rate $1/R_2$ in the year, t_{trig}, when either the damages exceed, or the abatement costs drop below, some specified target values. The logistic half-life R represents the years needed to reduce emissions to one-half the basecase. The damage target (in % GWP) is given by $D_{est}(t) > D_{trend}$ and the abatement-cost target (in $/ton carbon abated) is given by $K(t) < K_{thres}$. The second-period rate depends on the year t_{trig}. If $t_{trig} < T_1$, then the second-period abatement is given by $R_2 = R_{2/1}$. If $T_1 \le t_{trig} < T_2$, the second-period abatement is $R_2 = R_{2/2}$. If neither condition is met by the year T_2, the strategy switches to a second-period abatement $R_2 = R_{2/3}$. In practice such strategies would certainly be more complex, but the simple version used here is sufficient for an initial treatment of the effects of climate variability on the design and effectiveness of adaptive strategies.

We express the decision facing policy-makers as a choice among the eight parameters defining the adaptive-decision strategies used in this study. We chose from among a discrete set of values for each parameter that allows us to explore a large range of potentially interesting adaptive strategies while remaining within our computational limitations. In order to focus on the near-term features of the adaptive strategies, as shown in Table II we examine 4 possible first-period rates, 8 damage triggers, 5 innovation triggers, 2 alternative combinations of second-period rates, 4 possible earliest trigger years, and 4 possible latest trigger years, for a total of 5,120 different adaptive-decision strategies. Our results in this study appear insensitive to additional parameter combinations. For instance, we examined many cases in which we included additional parameters for the second-period rate and for the earliest and latest trigger years. In no case was there any significant change in the results reported here.

Figure 5 compares the regret for a number of adaptive-decision strategies to the regret for the static DAL and ES policies as a function of p_{ES}, the likelihood ascribed to the ES future, in the case with low SO_2 emissions (r_{SO_2} = 2%) and where the decision-maker ascribes no likelihood to a DR future (p_{DR} = 0%). We measure the performance of an adaptive-decision strategy using its 'regret', defined as the difference between the expected cost of the strategy for a given expectation about the future and the expected cost of choosing the best strategy if the future state of the world suddenly became known. (See the appendix for the mathematical description of the regret). The upper, middle and lower panels show the cases with "Low Variability", "High Variability" and "Increasing Variability", respectively. Each adaptive-decision strategy is labeled with its $(R_1, D_{thres}, K_{thres})$ values. We show only those strategies that have the smallest regret for some value of p_{ES}.[7] The best choice for each value of p_{ES} is shown by the thick line.

This figure shows that the best choice of climate-change policy depends strongly on the damages due to climate variability. For 'Low' damages due to variability, at least one of the adaptive-decision strategies is dominant over the DAL and ES static policies over a wide range of expectations about the future, $2\% \leq p_{ES} \leq 81\%$. However for 'Increasing' damages due to variability, the static ES policy dominates any adaptive-decision strategy over a wide range of expectations, $p_{ES} > 20\%$.

Climate variability degrades the performance of the adaptive strategies relative to the static policies because it masks any trend in the observed damage time series. For instance, Figure 6 compares the estimated damage $D_{est}(t)$ and the actual damage D(t) for two cases, the first with 'Low' damages due to variability and significant trend (α_1 = 3.5%) and the second with 'High' damages due to variability and no trend (α_1 = 0%). The Bayesian estimator does a reasonably good job of tracking both damage time series, but because of the high variability the estimates for the trend and no-trend cases do not diverge until about 2020. Thus, an adaptive-decision strategy attempting to distinguish between these cases based on observations of the damage time series would have to wait at least two decades before being able to act. In contrast, the static DAL and ES policies are insensitive to the variability. Since they make no observations, these policies cannot be affected by false signals, and thus the variability only affects the variance of their regret (not shown here), not the mean.

Perhaps not surprisingly, as the uncertainties proliferate, the adaptive-decision strategies become the dominant choice independent of the variability. Figure 7 compares the regret for a number of adaptive-decision strategies to the regret for the static DAL and ES policies as a function of p_{ES} in the case where uncertainties about the sulfate forcing and DR future now play a role. This figure assumes high SO_2 emissions ($r_{SO_2} = 0\%$), so that different values of $\Delta F_{SO_4}(1990)$ will affect the results, and a small chance ($p_{DR} = 2\%$) of a DR future. The regrets for the adaptive strategies are similar to those in Figure 5, but the regrets of the static policies are significantly larger. For instance, the upward shift of the $p_{ES} = 0\%$ and 100% intercepts for the DAL and ES policy regrets is due to the fact that the static policies are more sensitive to a small probability of a DR future than the adaptive-decision strategies, because the latter are able to respond to the onset of a very large damage trend. We thus find that the results of LSB hold, even when climate variability is large, if there is significant uncertainty about sulfate forcing and/or even a small chance of a drastic future.

3.2. Which Adaptive Strategy to Choose

Figures 5 and 7 also show that the particular adaptive-decision strategy a decision-maker should choose depends strongly on the climate variability and his or her expectations about the relative likelihoods of the DAL and ES futures. Not surprisingly, the first-period rate of the best adaptive-decision strategy depends most strongly on the probability ascribed to the ES future. For instance, for 'Low' damages due to variability in Figure 5, $R_1 = \infty$ for $2\% \le p_{ES} \le 10\%$, $R_1 = 100$ years for $10\% \le p_{ES} \le 57\%$, and $R_1 = 60$ years for $57\% \le p_{ES} \le 81\%$. We see similar patterns in all the panels of Figures 5 and 7 (except for the bottom panel of Figure 5).

The damages due to climate variability also affect the first-period rate of the best strategy, decreasing R_1 from 100 years to 60 years as we go from 'Low' to 'High' variability in the region near p_{ES} = 50% in Figure 5, and - from ∞ to 60 years as we go from 'High' to 'Increasing' variability in the region near p_{ES} = 10%. However variability has its largest impact on the damage threshold, D_{thres}, of the best adaptive-decision strategies. For 'Low' variability case in Figures 5 and 7, all the lowest regret strategies have D_{thres} = 0.4% GWP, which nearly doubles (Figure 5) or triples (Figure 7) for 'High' damages due to variability. These higher thresholds are necessary to reduce the risk that the strategy will mistake large random fluctuations as the onset of an adverse trend in the damages. However, this higher threshold delays the detection of real damage trends, and thus degrades the performance of the strategies. For instance, $(60, 0.75\%, \$65)$, the best strategy for 'High' variability and $24\% \le p_{ES} \le 56\%$ in Figure 5, has a p_{ES}=56% regret of \$22 billion/year, over a third larger than that of $(60, 0.4\%, \$40)$, the most cost-effective strategy for 'Low' variability and $57\% \le p_{ES} \le 81\%$, with a p_{ES}=57% regret of \$16 billion/year.[8]

The most significant difference between the strategies in Figures 5 and 7 is the height of the damage thresholds, D_{thres}. This difference is due to the high SO_2 emissions in the latter figure, rather than to a plausible drastic future. Higher SO_2 emissions reduce the rate of temperature increase for a given climate sensitivity, and thus reduce the damages due to climate change. All other things being equal, a decision-maker uncertain about the size of the sulfate forcing (ascribes relatively equal probabilities to $\gamma_{SO_4} = 0.25$ and 1.0) will expect a lower probability of very high damages in the case of high SO_2 emissions than in the case of low SO_2 emissions. Thus, in the former case it is better on average to raise the damage threshold, sacrificing some ability to rapidly detect large trends, in order to reduce the rate of false alarms.

3.3. Robust Adaptive Strategies

Until now we have considered the situation where the variability is known and where the decision-maker has a definite expectation about the relative likelihood of the DAL and ES futures. In practice, decision-makers do not know the damages due to climate variability, nor is there widespread agreement about the relative likelihood of different climate-change futures. Much of the political debate about climate-change policy is dominated by stakeholders who hold very different expectations about these issues.

In such a situation it may be useful to employ a climate-change strategy that is insensitive to the probabilities a decision-maker ascribes to the size of the damages due to variability or to the various climate-change futures. Such robust adaptive-decision strategies should perform relatively well in any possible climate-change future.

In the case of high SO_2 emissions, the most robust strategy is $(60, 1.2\%, \$65)$.[9] Figure 7 shows the regret for this strategy (heavy dashed lines) as a function of p_{ES}. This strategy is the best choice for a decision-maker who ascribed equal likelihoods to the DAL and ES futures in the 'High' and 'Increasing' variability cases and close to the best in the 'Low' variability case. With its high damage threshold, D_{thres} = 1.2%, the strategy responds primarily to trends in abatement costs, switching rapidly to a second-period abatement rate of R_2 = 40 years when innovation reduces the costs of abatement (innovation parameter d = 2% and 5%). If abatement costs do not fall, the strategy can increase its abatement rate if damages rise rapidly to the high damage threshold. Otherwise, the strategy reduces its abatement rate to R_2 = 150 years. The choice of this strategy is relatively insensitive to the probability ascribed to the DR future, which can range from $0\% < p_{DR} \approx 15\%$ before a decision-maker should choose a different strategy, $(40, 0.75\%, \$50)$. We also find a similar robust strategy, $(60, 0.75\%, \$65)$, in the case of low SO_2 emissions. As is generally the case for low SO_2 emissions, this strategy adopts a lower damage threshold, D_{thres} = 0.75%, compared to the case with high SO_2 emissions.

Given the intense political focus on near-term emissions reductions, it is important to note that a decision-maker can also craft a robust strategy with a slower near-term abatement rate. As shown in Figure 7, the strategy $(100, 1.2\%, \$65)$ performs nearly as well as the faster-starting $(60, 1.2\%, \$65)$ in the case of high SO_2 emissions. Similarly for the case of low SO_2 emissions, the strategy $(100, 0.6\%, \$65)$ performs nearly as well as the faster-starting $(60, 0.75\%, \$65)$. The former strategy compensates for slower near-term abatement by slightly increasing its sensitivity to large, adverse trends at the cost of slightly increasing its susceptibility to false alarms. Thus, policy-makers can choose among strategies that perform similarly, on average, by trading off among near-term emissions abatement rate and the risk of erroneously responding to a false trend, as determined by the sensitivity of the damage threshold.

It is also interesting to note that in the face of potentially large damages due to climate variability, the most robust adaptive-decision strategies are "innovation-sensitive" rather than "damage-sensitive." That is, these robust strategies will change to their second-period abatement rate based on small observed changes in innovation-reduced abatement costs, but only for large observed changes in climate damages.[10]

3.4. Value of Information

Adaptive-decision strategies use endogenous-information-based observations to make their decisions. However, other sources of information are also available to aid decision-makers. In particular, several studies have examined the value of exogenous information in the context of choosing fixed or sequential strategies for abating climate change (Peck and Teisberg, 1993; Nordhaus, 1994). In contrast to the endogenous observations shown in Figure 4, exogenous information generally represents new scientific information, gained independently of any induced changes in the systems being studied, that improves our estimates of key model parameters.

Here we examine the value exogenous information has for a decision-maker choosing an adaptive-decision strategy in the face of uncertain climate variability. In contrast to previous studies, we can estimate the value of exogenous information as a function of climate variability, which affects the rate of endogenous learning, as well as estimate the value of information about the variability itself. The value of information about the variability is an important question, because at present most climate-change impacts research is focused on predicting the long-term damages from climate change rather than on characterizing the damages due to variability.

We calculate the value of exogenous information for an adaptive-decision strategy by: i) searching for the best strategy for each of a set of particular climate-change futures, and ii) searching for the best adaptive-decision strategy assuming a specific probability distribution over the set of futures. The expected value of information is then given by the expected value (assuming the probability distribution) of the difference in regrets for each future between the strategies found in i) and ii). For instance, imagine a decision-maker sure that SO_2 emissions will be low, damages due to variability will be low, and a drastic future is implausible, but who ascribes equal likelihood to the DAL and ES futures. From the upper panel of Figure 5 we see that information distinguishing the DAL and ES futures would be worth $15 billion/year, since without it, the decision-maker's best choice is the strategy $(100, 0.4\%, \$65)$, with a p_{ES}=50% regret of $15 billion/year.

In general the value of information distinguishing the DAL and ES future states of the world is about $22.5 billion/year for a decision-maker who ascribes equal likelihoods to 'Low', 'High', and 'Increasing' variability (labeled 'Unknown Variability' in Figure 8), in the case of low SO_2

emissions, and is $16 billion/year in the case of high SO_2 emissions. We focus in this figure on the case of a decision-maker who ascribes equal likelihoods to the DAL and ES futures, who knows whether SO_2emissions are high or low, and who ascribes a 0% probability to a drastic climate-change future. We also focus in this study on exogenous information that becomes available immediately, as opposed to that which becomes available at some later time in the future. The value of information is higher in the case of low SO_2 emissions because, as discussed above, the adaptive-decision strategies work relatively better in the case of high SO_2 emissions. As shown in Figure 8, when the decision-maker knows the damages due to climate variability, we see that the value of information distinguishing the DAL and ES futures is about 30% higher for 'High' variability than for 'Low' variability, because the adaptive strategies perform better in the latter case.

Figure 9 shows the expected value of information (EVOI) about the climate variability to a decision-maker who ascribes equal likelihood to both the DAL and ES futures, as well as equal likelihood to the Low, High and Increasing Variability cases. Comparing the EVOI of the Low and High SO_2 Emissions/No Drastic cases in Figure 9 with the EVOI of the Low and High SO_2/Unknown Variability cases in Figure 8, we see that information distinguishing the variability cases is about one-third as valuable as information distinguishing between the DAL and ES futures in the case of low SO_2 emissions, and one-eight as valuable in the case of high SO_2 emissions The possibility that the variability may be 'Low' makes a significant contribution to the expected value of information, independent of the decision-maker's expectations about the SO_2 emissions or the possibility of a DR future, because proof of 'Low' variability would allow the decision-maker to choose a strategy with a low damage threshold. For instance, in the low SO_2 emissions case, a decision-maker who learns the variability is 'Low' can choose the strategy $(100, 0.4\%, \$65)$ (Fig. 5), which performs \$9 billion/year better than the high-damage-threshold strategy, $(60, 0.75\%, \$65)$, that hedges against all three variability cases. Similarly, information which proves the variability is 'Increasing' is also valuable in the case of low SO_2 emissions.

In contrast, the possibility that the variability is 'High', or that the variability may be 'Increasing' in high SO_2-emissions cases, does not make a significant contribution to the expected value of information, since proof of 'High' of 'Increasing' variability does not change the decision-maker's best strategy. Information that guarantees a drastic climate-change future is implausible ($p_{DR} = 0\%$ rather than 2%) is also worth little except in the case of low SO_2 emission and 'Increasing' damages due to variability, because only in this case does this information make a significant difference to a decision-maker's choices. However, information that guarantees that the probability of a drastic future is less than approximately 15% is valuable because this information will make a difference in the choice of an adaptive-decision strategy.

4. CONCLUSIONS

This study is one of the first that examines the choice among alternative, adaptive climate-change strategies in the face of climate variability. We find that adaptive-decision strategies remain preferable to static, best-estimate policies even with very large levels of climate variability; that the most robust strategies are innovation-sensitive, that is, adjust future emissions-reduction rates on the basis of small changes in observed abatement costs but only for large changes in observed damages; and that information about the size of the variability is about a third to an eighth as valuable as information determining the value of the climate sensitivity, sulfate forcing, damage, and technology innovation parameters that represent the long-term, future climate-change state-of-the-world.

We have made a number of important simplifications in this study that place caveats on these conclusions and influence any suggestions we can provide to policy-makers about the design of effective and robust adaptive-decision strategies. Perhaps most importantly, our representation of climate variability is invariant over time and may under-represent longer period fluctuations, while our representation of damages due to variability clearly ignores the effects of variations in climate factors other than temperature (such as precipitation). These omissions will likely understate the effects of variability on adaptive strategies. In contrast, we have considered adaptive-decision strategies that make only a single midcourse correction. This assumption will likely overstate the impact of variability on adaptive-decision strategies because in reality, a strategy can make at least several important midcourse corrections over the next century. On balance, it is unclear whether more sophisticated adaptive strategies will perform better or worse for a more complete representation of variability than do the simple adaptive strategies against the simple variability considered here. This question, among others, is an important one for further research.

Despite these shortcomings, we believe this study provides two important inputs for climate-change policy-makers. First, the study suggests that the climate-change-impacts research community should place relatively more effort on characterizing the variability due to climate change and its effects. There is a large, worldwide research effort devoted to predicting the long-term damages due to climate change. However, relatively little effort seems to be devoted to collecting and assessing information on variability. We can only offer anecdotal evidence on this score. But while developing the crude representation of the damages due to climate variability used in this study, we were unable to find the answer to this straight-forward and extremely policy-relevant question:

> "Assume decision makers came to an agreed-upon definition of damages due to climate change and the scientific

> community reported annual estimates of these damages based on the data collected by satellites, terrestrial monitoring stations, and all other sources. Assume further there were no anthropogenic climate change. What would be the magnitude and stochastic properties of the damage estimates reported each year?"

If we had the answer to this question, it would be relatively easy to conduct an analysis such as the one in this study and find the adaptive-decision strategies (or the prescriptive policies) that perform best against this level of variability. Much of the scientific information needed to address this question may already exist, but no one in the integrated-assessment community has yet compiled it into a useful form. This study suggests that this would be a very useful thing for policy-makers to ask the integrated-assessment community to do.

Second, this study provides some tentative suggestions about the design of the particular adaptive-decision strategies policy-makers should employ. The current Framework Convention on Climate Change requires that the developed nations reduce their greenhouse-gas-emissions to 5% below their 1990 levels by 2010. It is difficult to relate this target directly to the results of our model because we consider global aggregate emissions while Kyoto distinguishes between the emissions of developed and developing nations, and because our focus on long-lived capital stock limits the range of emissions reductions we can consider in 2010. Nonetheless, the Kyoto targets roughly correspond to a near-term abatement rate in our model of $R_1 = 100$ years, which we find to be a reasonable choice for decision-makers who believe there is any significant probability that the best climate-change strategy will ultimately require emission stabilization.

In addition, this study suggests that the rate of near-term reductions, while important, is not the sole, dominant factor determining the success of an adaptive-decision strategy. In particular, for any given set of expectations about the future, there is a trade-off between the near-term emissions-abatement rate and the sensitivity with which the strategy responses to potential trends in damages due to climate change. Thus, faced with stakeholders reluctant to commit to aggressive near-term reductions, policy-makers may have the flexibility to craft a robust adaptive-decision strategy by encouraging those stakeholders to agree to a program of environmental and technology monitoring, and to commit beforehand to a particular aggressive response if in the future certain trends were detected by such monitoring.

Our study also suggests that variability might mask any damage trends for at least several decades. Thus, independent of the rate of near-term emissions abatement, the most robust adaptive-decision strategies in the presence of climate variability will be innovation sensitive, that is, will change future abatement rates based primarily on observations of small changes in abatement costs, and be sensitive only to relatively large trends in the damages due to climate change.

[1] The United Nations Framework Convention on Climate Change recognizes that steps to address climate change will be most effective if continually re-evaluated in the light of new information.

[2] In addition to LSB, exploratory modeling has been widely used in a number of policy studies, including studies of science and technology investment portfolios (Lempert and Bonomo, 1998), the future of California higher education (Park and Lempert, 1998), weapons procurement decisions for the Air Force (Brooks, Bankes and Bennett, 1997), and military strategy (Davis, 1997a, 1997b).

[3] It is also plausible that sensitivity and sulfate forcing are not constant, but depend on temperature.

[4] We have chosen the normalization factors in the denominators in Eq. (4) so that the coefficients have straightforward interpretations. The coefficient α_1 represents the damages due to a 3°C increase in the global-mean temperature, and the coefficients α_2 and α_3 represent the maximum damages due to climate variability in 1995 at the 90% confidence level, as seen from Figure 2. All three coefficients are measured in units of percent of GWP.

[5] In 1996, large-scale events caused $180 billion (0.6% GWP) in damages, but much of this was due to the Kobe earthquake.

[6] The results in this paper required about three weeks of CPU time, or six million calls to our linked system of models. We considered 20 Monte Carlo runs for each strategy-uncertainty space combination and (as described in Section 3) 5,120 strategies. This allows us about 60 scenarios. We have verified that our results are insensitive to this relatively small number of Monte Carlo runs.

[7] We find these strategies with a two-step screening process. First, we screen for strategies that are relatively robust by searching for the strategy with the lowest ES regret, $R\left(Z_j|100\%,0\%\right)$, among those strategies with a DAL regret, $R\left(Z_j|0\%,0\%\right)$, less than some given value. (These symbols are defined in the Appendix.) We use a series of DAL regrets in $5 billion/year increments starting with $2.5 billion/year. We use this first screen because we assume that a decision-maker, choosing among alternative strategies that perform nearly equally well in the future they expect to be most likely, will prefer that strategy that will also perform well in the future they consider least likely. Second, we plot the regrets of the strategies remaining after the first screen and discard those that do not give the lowest regret for some values of p_{ES}. In the case of low SO_2 emissions, the results for γ_{SO_4} = 0.25 and 1.0 are similar, so we only use the 36 uncertainty space points with γ_{SO_4} = 1.0.

[8] Using the notation in the appendix, the former regret is $R_{High}\left(Z_j|56\%,0\%\right)$ = $22 billion/year and the later is $R_{Low}\left(Z_j|57\%,0\%\right)$ = $16 billion/year.

[9] Our robustness criteria is defined in the Appendix.

[10] A strategy with relatively large numbers for its damage and abatement cost thresholds is sensitive to small changes in abatement costs and only large changes in damages because we model damages as currently small and potentially increasing and abatement costs as currently high and potentially decreasing.

APPENDIX

In Section 3 we consider the regret costs of adaptive-decision strategies and static policies. This appendix describes how we calculate these regrets.

As shown in Table I, we begin by assigning each uncertainty-space point $X_i = \left(\Delta T_{2x}, \gamma_{SO_4}, \alpha_1, d\right)$ to one of the exclusive sets of 'Do-a-Little' (DAL), 'Emissions-Stabilization' (ES) or 'Drastic Reductions' (DR) future states of the world, . We write the regret for strategy Z_j as

$$
\begin{aligned}
R_V\left(Z_j \middle| p_{ES}, p_{DR}\right) = \\
&\frac{(1-p_{ES})(1-p_{DR})}{\vartheta_{DAL,V}} \sum_{X_i \in DAL \cup V} \left[\overline{C}\left(Z_j \middle| X_i\right) - \overline{C}\left(Z_{DAL} \middle| X_i\right)\right] \\
&+ \frac{p_{ES}(1-p_{DR})}{\vartheta_{ES,V}} \sum_{X_i \in ES \cup V} \left[\overline{C}\left(Z_j \middle| X_i\right) - \overline{C}\left(Z_{ES} \middle| X_i\right)\right] \\
&+ \frac{p_{DR}}{\vartheta_{DR,V}} \sum_{X_i \in DR \cup V} \left[\overline{C}\left(Z_j \middle| X_i\right) - \overline{C}\left(Z_{DR} \middle| X_i\right)\right]
\end{aligned}
\qquad \text{(A1)}
$$

where $\overline{C}\left(Z_j \middle| X_i\right)$ is the cost of the strategy Z_j (given by one of the sets of parameters in Table II) at the uncertainty-space point X_i; p_{ES} is the relative likelihood ascribed to an ES future as opposed to a DAL future, and p_{DR} as the probability ascribed to a DR future; Z_{DAL}, Z_{ES}, and Z_{DR} are the DAL, ES, and DR static policies defined in LSB; V = Low, High, or Inc labels the 'Low Variability', 'High Variability', and 'Increasing Variability' cases; where the sums are over the set of uncertainty-space points X_i

representing the appropriate variability (Low, High, or Inc) and climate-change future (DAL, ES, or DR); and where $\vartheta_{DAL,V}$, $\vartheta_{ES,V}$, and $\vartheta_{DR,V}$ are the total number of points in each of these sets. The $\overline{C}\left(Z_j|X_i\right)$ are annualized global sums of the abatement and damage cost time-series from the present to 2140, calculated using a 5% discount rate, averaged over a set of Monte Carlo draws for the stochastic radiative forcing time-series in Eq. (1). We use the same set of time-series for all the strategies Z_j we compare at each uncertainty-space point X_i.

In Section 3 we also find the most robust adaptive-decision strategies. In order to find such strategies we define a measure of robustness given by

$$\begin{aligned}
&\text{Robustness}\left(Z_j|p_{DR}\right) = \\
&\quad 0.5\left(1-p_{DR}\right)\left\{\frac{1}{\vartheta_{DAL}}\sum_{X_i\in DAL}\left[\overline{C}\left(Z_j|X_i\right)-\overline{C}\left(Z_{DAL}|X_i\right)\right]\right\}^2 \\
&\quad +0.5\left(1-p_{DR}\right)\frac{1}{\vartheta_{ES}}\left\{\sum_{X_i\in ES}\left[\overline{C}\left(Z_j|X_i\right)-\overline{C}\left(Z_{ES}|X_i\right)\right]\right\}^2 \\
&\quad +p_{DR}\frac{1}{\vartheta_{DR}}\left\{\sum_{X_i\in DR}\left[\overline{C}\left(Z_j|X_i\right)-\overline{C}\left(Z_{DR}|X_i\right)\right]\right\}^2 .
\end{aligned} \tag{A2}$$

and search for the strategy Z_j which gives the smallest value for $\text{Robustness}\left(Z_j|p_{DR}\right)$ as a function of p_{DR}. This expression is similar to Eq. (A1) with p_{ES} = 50%, except we take the square of the regret terms and sum over all three variability cases so that our robustness measure is analogous to a least-squares fit across the uncertainty space.

REFERENCES

Bankes, S. C.: 1993, 'Exploratory Modeling for Policy Analysis', Operations Research, vol. 41, no. 3, pp. 435-449, (also published as RAND RP-211).

Brooks, A., Bankes, S., Bennett, B.: 1997, "Weapon Mix and Exploratory Analysis," RAND, DB-216/2-AF.

Davis, P. K. and J. Carrillo, M. J.: 1997a, "Exploratory analysis of the halt problem: a briefing on methods and initial insights," RAND, DB-232.

Davis, P. K., R Kugler, R., and J. Hillestad, R. J.: 1997b, "Strategic Issues and Options for the Quadrennial Defense Review", RAND, DB-201-OSD.

Hammitt, J. K., Lempert, R. J., and Schlesinger, M. E.: 1992, "A sequential-decision strategy for abating climate change," Nature, 357, 315–318.

Houghton, J. T., Jenkins, G. J. and Ephraums, J. J. (eds.): 1990, Climate Change The IPCC Scientific Assessment. Cambridge University Press, Cambridge, 364 pp.

Katz, R., and Brown, B. G.: 1992, "Extreme Events in a Changing Climate: Variability is More Important than Averages," Climatic Change 21: 289-302.

Kelly, D. L. and Kolstad, C. D.: 1996, "Tracking the Climate Change Footprint: Stochastic Learning About Climate Change," University of California Santa Barbara Economics Working Paper 3-96R.

Leggett, J.A., Pepper, W.J., and Swart, R.J. in Climate Change, 1992: The Supplementary Report to the IPCC Scientific Assessment, eds. Houghton, J.T., Callander, B.A. and Varmey, S.K.) 69-95, Cambridge University Press.

Lempert, R. J. and Bonomo, J.: 1998, "New Methods for Robust Science and Technology Planning," RAND, DB-238-DARPA.

Lempert, R. J., Schlesinger, Bankes, S. C., and Andronova, N.G., 1998: (RAND working paper, in preparation).

Lempert, R. J., Schlesinger, Bankes, S. C.: 1996, "When We Don't Know the Costs or the Benefits: Adaptive Strategies for Abating Climate Change," Climatic Change, 33: 235-274.

Lempert, R. J., Schlesinger, M. E., Hammitt, J. K.: 1994, "The Impact of Potential Abrupt Climate Changes on Near-Term Policy Choices," Climatic Change, 26:351-376.

Lewis, F. L.: 1986, Optimal Estimation, John Wiley & Sons.

Manne, A. S. and Richels, R. G.: 1992, Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emissions Limits, MIT Press, Cambridge, 182 pp.

Manne, A. S. and Richels, R. G.: 1991, "Global CO_2 emission reductions: The impact of rising energy costs," The Energy Journal, 12, 88-107.

Mendelsohn, R., W. Nordhaus, and D. Shaw: 1994, "The Impact of Global Warming on Agriculture: A Ricardian Analysis" American Economic Review 84: 753-771.

Munich Re (Muncherner Ruckversicherungs-Gesellschaft, D-80791 Munchen, Germany (http://www.munichre.com), 1997).

Nordhaus, W. D.: 1991, "The Cost of Slowing Climate Change: A Survey," The Energy Journal, 12, 37-65.

Nordhaus, W. D.: 1994, Managing the Global Commons: The Economics of Global Change, MIT Press, 213 pp.

Park, G, and Lempert, R. J., 1998: "The Class of 2014: Preserving Access to California Higher Education," RAND, MR-971.

Peck, S. C. and Teisberg, T. J., 1993: "Global warming uncertainties and the value of information: An analysis using CETA", Resource and Energy Economics, 15, 71-97.

Schlesinger, M. E., and N. G. Andronova, 1998: (in preparation).

Schlesinger, M. E., and Andronova, N. G.: 1996, The standard error of observationally based estimates of climate sensitivity and sulfate forcing (in preparation).

Schlesinger, M. E., and Jiang, X.: 1991, "Revised projection of future greenhouse warming," Nature, 350, 219-221.

Schwartz, P.: 1996, The Art of the Long View, Currency Doubleday, NY, 272 pp.

Wigley, T. M. L., Richels, R. & Edmonds, J. A.: 1996, Nature 379, 240-243.

TABLES

Climate Sensitivity ΔT_{2x}	Sulfate forcing γ_{SO_4}	Damage coefficient α_1	Innovation Rate d	Climate Change Future
0.5°C	1.0	3.5%	5%	DAL
1.5°C	1.0	0%	2%	DAL
1.5°C	1.0	2%	0%	DAL
2.5°C	1.0	0%	0%	DAL
2.5°C	0.25	0%	0%	DAL
1.5°C	1.0	5%	5%	ES
2.5°C	1.0	3.5%	2%	ES
2.5°C	1.0	2%	5%	ES/ES/DR
2.5°C	0.25	3.5%	2%	ES
2.5°C	0.25	2%	5%	ES
3.5°C	1.0	2%	0%	ES
3.5°C	0.25	2%	0%	ES
2.5°C	1.0	7%	2%	DR
2.5°C	1.0	10%	5%	DR
2.5°C	0.25	7%	2%	DR
2.5°C	0.25	10%	5%	DR
3.5°C	1.0	5%	2%	DR
3.5°C	0.25	5%	2%	DR
4.5°C	1.0	7%	0%	DR
4.5°C	0.25	7%	0%	DR

For each $(\Delta T_{2x}, \gamma_{SO_4}, \alpha_1, d)$ point we also consider three cases for the damages due to climate variability: $(\alpha_2, \alpha_3, \eta_2, \eta_3)$ = (0.2%,0%,1,na), (0.4%,0%,2,na), and (0%,0.33%,na,3), which we label 'Low', 'High', and 'Increasing' damages due to climate variability.

Table I: Uncertainty Space Points.

Parameter	Description	Values
R_1	First period rate	∞, 100, 60, 40 years
D_{thres}	Damage trigger	0.1%, 0.25%, 0.4%, 0.6%, 0.75%, 1.2%,1.5%, and 2% GWP
K_{thres}	Innovation trigger	4, 25, 40, 50, and 65 $/ton-carbon
$R_{2/1}$/ $R_{2/2}$/ $R_{3/2}$	Alt. Second Period Rates	20/40/150 years, 20/70/150 years
T_1	Last year for early trigger	2005, 2010, 2020, 2030
T_2	Last year for late trigger	2017, 2025, 2030, 2040

Table II: Strategy parameters.

FIGURE CAPTIONS

Figure 1: Climate sensitivity and sulfate forcing pairs $\left(\Delta T_{2X}, \Delta F_{SO_4}\right)$ (point cloud) and the magnitude of white-noise forcing, σ_Q, as a function of $\left(\Delta T_{2X}, \Delta F_{SO_4}\right)$ (contour lines) estimated from the 1856 to 1995 instrumental temperature record.

Figure 2: Variability in 1995 of the global-mean temperature as measured by the cumulative probability distributions of the annual temperature $\Delta T(1995)$, the five-year running average $\Delta T_5(1995)$, and the differences between the temperature and the five- and thirty-year running averages, $\Delta T(1995) - \Delta T_5(1995)$ and $\Delta T(1995) - \Delta T_{30}(1995)$. All distributions calculated using the climate parameters $\left(\Delta T_{2X}, \Delta F_{SO_4}, \sigma_Q\right) = \left(2.5°C, -0.7\,W/m^2, 3.2\,W/m^2\right)$.

Figure 3: Damages due to climate variability as measured by the cumulative probability distribution (solid lines) for 1995 with two sets of assumptions about the damage function parameters, $(\alpha_2, \alpha_3, \eta_2, \eta_3)$ = $(0.2\%, 0\%, 1, na)$ and $(0.4\%, 0\%, 2, na)$, and calculated for 1995 and 2020 with the damage function parameters $(0\%, 0.33\%, na, 3)$. All four distributions calculated with the climate parameters $\left(\Delta T_{2X}, \Delta F_{SO_4}, \sigma_Q\right) = \left(2.5°C, -0.7\,W/m^2, 3.2\,W/m^2\right)$. Solid circles show observed damages due to extreme events from 1960 to 1995.

Figure 4: Flow chart describing the adaptive-decision strategies.

Figure 5: Regret for adaptive-decision strategies and static policies as a function of the likelihood ascribed to the "Emissions Stabilization" future (p_{ES}), assuming low SO_2 emissions (r_{SO_2} = 2%) and no likelihood of a "Drastic Reductions" future (p_{DR} = 0%). Robust adaptive strategies shown with dashed lines. All strategies labeled with their $\left(R_1, D_{thres}, K_{thres}\right)$ values. Thick lines show strategies with lowest regret.

Figure 6: Annual estimated, $D_{est}(t)$, and actual, $D(t)$, damages for two cases, the first with 'Low' damages due to variability and significant trend (α_1 = 3.5%) and the second with 'High' damages due to variability and no trend (α_1 = 0%), both calculated using the climate parameters $\left(\Delta T_{2x}, \Delta F_{SO_4}, \sigma_Q\right) = \left(2.5°C, -0.7W/m^2, 3.2W/m^2\right)$ and the IS92a estimate for SO_2 emissions ($r_{SO_2} = 0$).

Figure 7: Regret for adaptive-decision strategies and static policies as a function of the likelihood ascribed to the "Emissions-Stabilization" future

(p_{ES}), assuming high SO_2 emissions ($r_{SO_2} = 0\%$) and a 2% likelihood of a "Drastic-Reductions" future ($p_{DR} = 0\%$).

Figure 8: Expected value of exogenous information distinguishing 'Do-a-Little' (DAL) and "Emissions-Stabilization" (ES) future states of the world assuming equal likelihoods for the DAL and ES futures and no likelihood of a "Drastic Reductions" future ($p_{DR} = 0\%$), as a function of the SO_2 emissions and the variability. Unknown variability assumes an equal likelihood for the 'Low', 'High' and 'Increasing' variability states of the world.

Figure 9: Expected value of exogenous information distinguishing the 'Low', 'High' and 'Increasing' variability states of the world as a function of future SO_2 emissions ($r_{SO_2} = 0\%$ or 2%) and the likelihood ascribed to a "Drastic-Reductions" future ($p_{DR} = 0\%$ or 2%), assuming equal likelihood for the three variability cases. The three segments in each bar show the contribution from the expected value of each variability case.

Figure 1

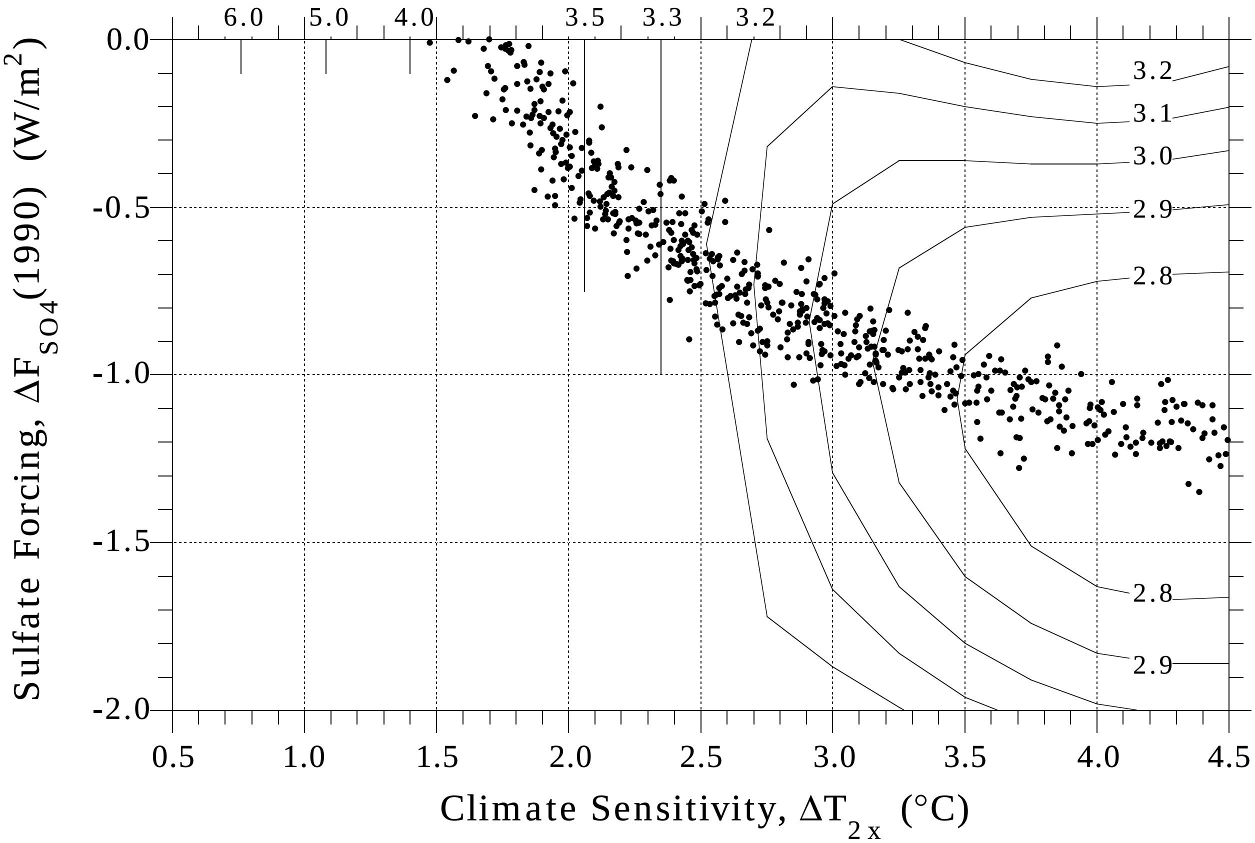

Figure 2

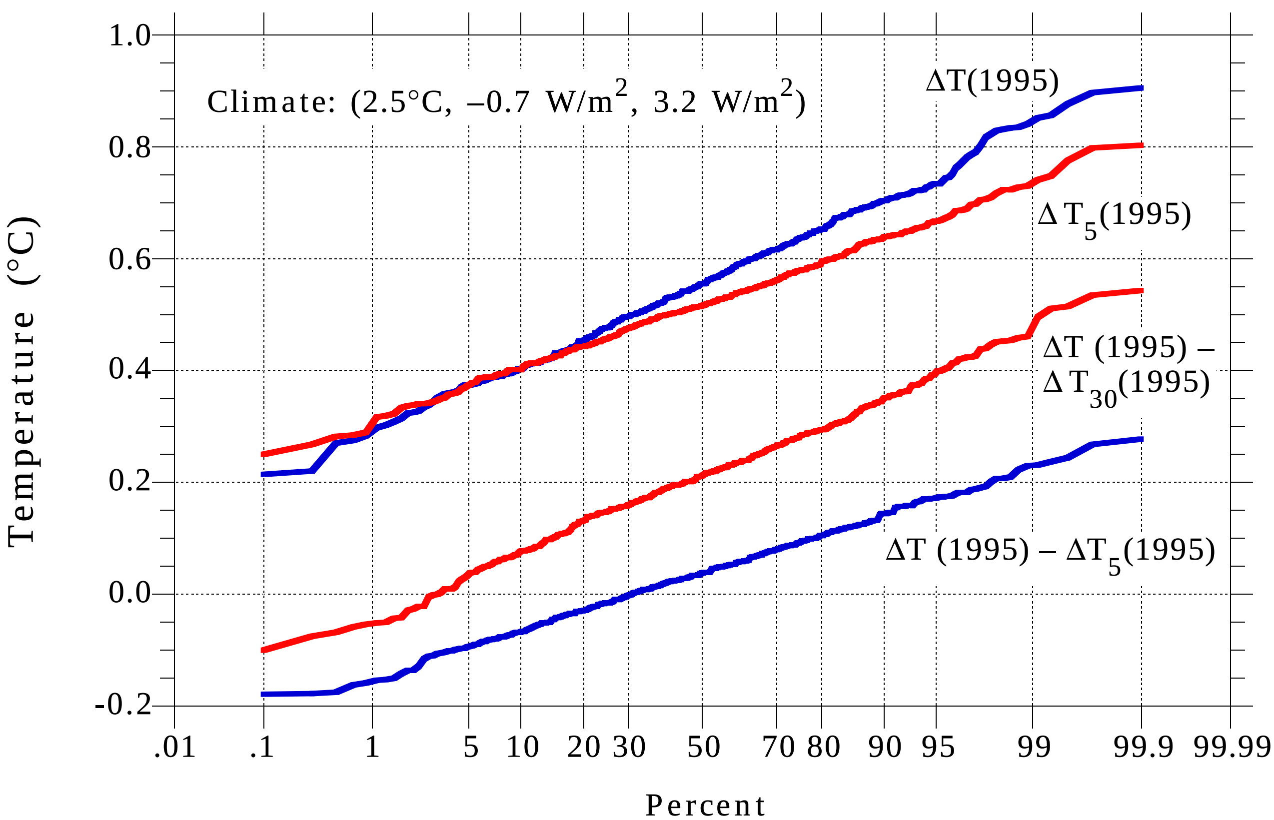

Figure 3

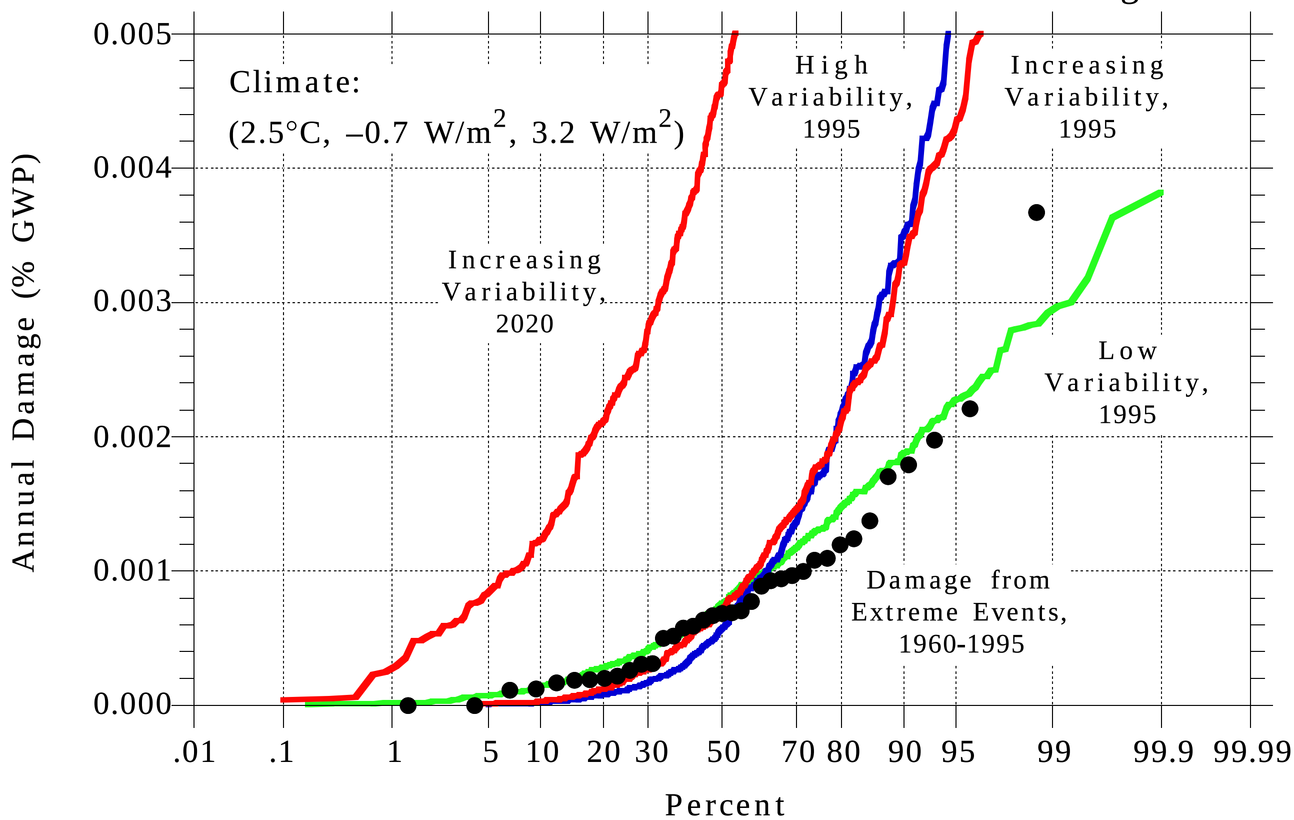

Figure 6

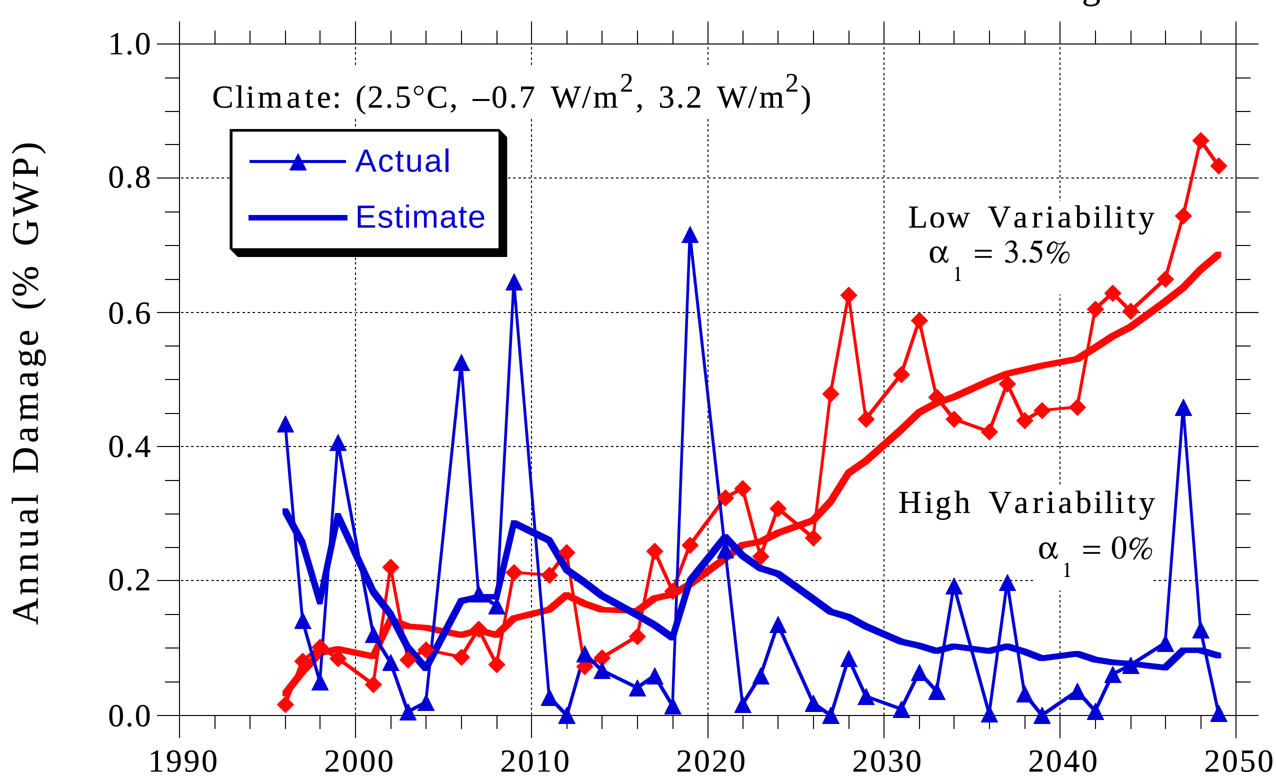

Figure 4

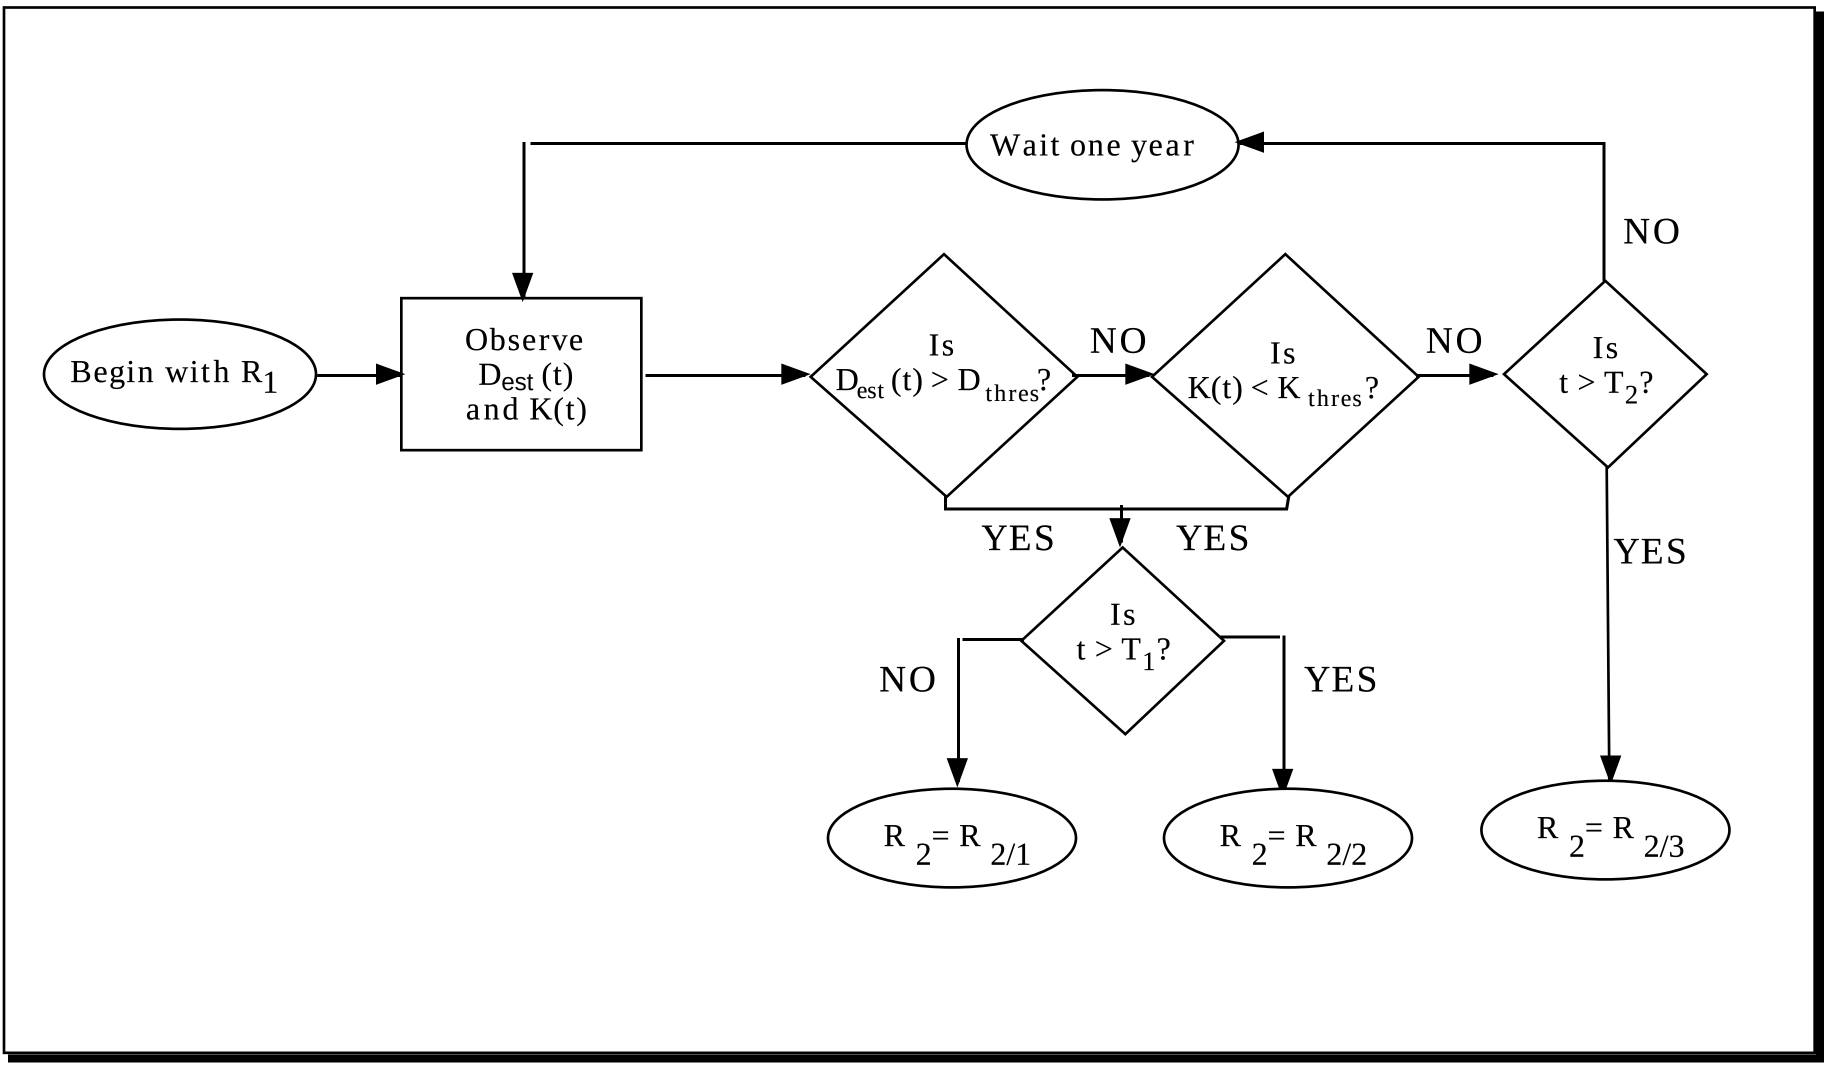

Sulfate Trigger

Damage and Abatement Cost Trigger

Figure 5

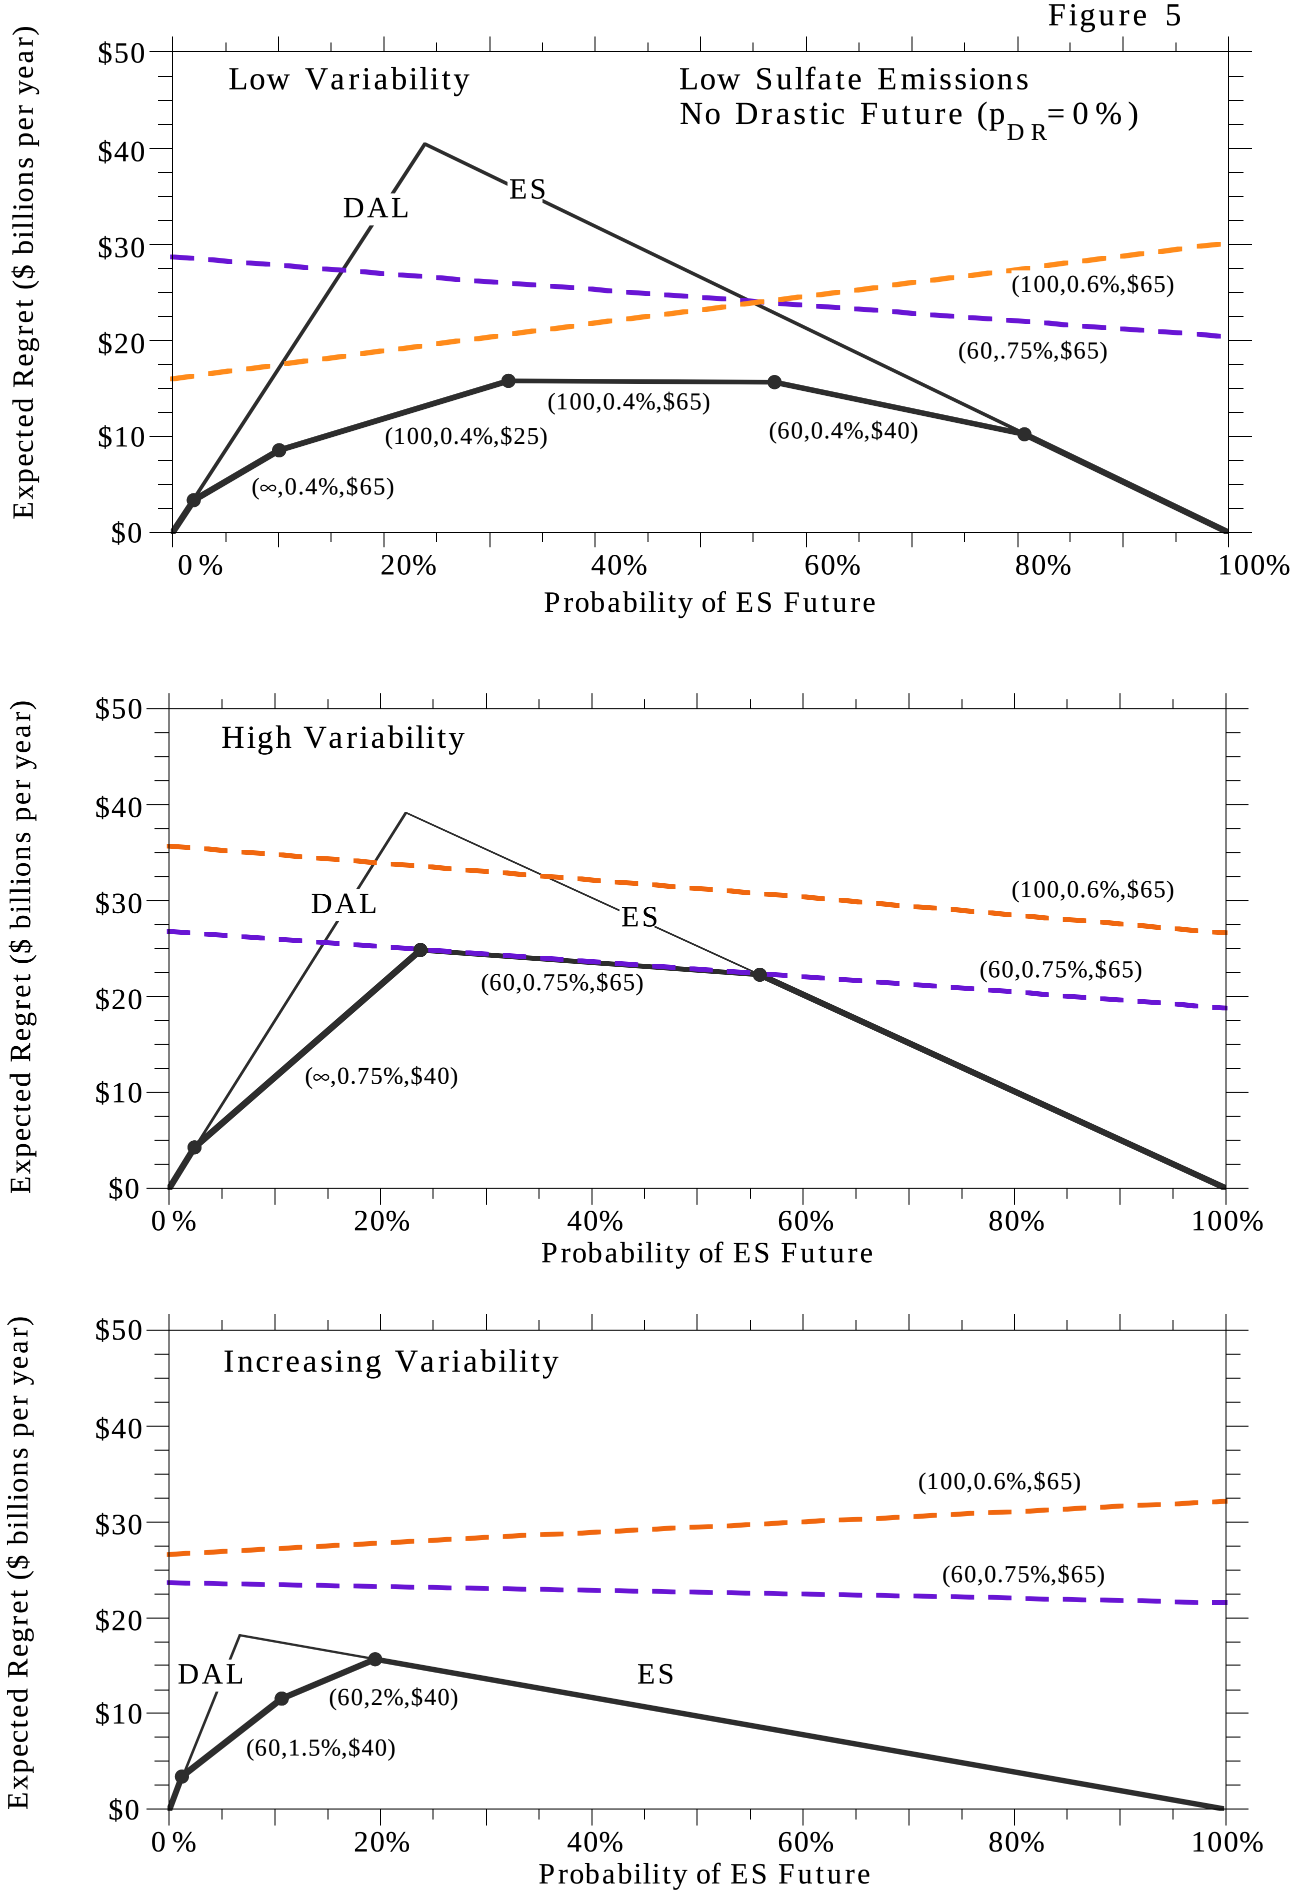

Low Variability
Low Sulfate Emissions
No Drastic Future (p_{DR} = 0 %)
DAL
ES
(100,0.6%,$65)
(60,.75%,$65)
(100,0.4%,$65)
(100,0.4%,$25)
(60,0.4%,$40)
(∞,0.4%,$65)
Expected Regret ($ billions per year)
$50
$40
$30
$20
$10
$0
0 %
20%
40%
60%
80%
100%
Probability of ES Future
High Variability
DAL
ES
(100,0.6%,$65)
(60,0.75%,$65)
(60,0.75%,$65)
(∞,0.75%,$40)
Expected Regret ($ billions per year)
Probability of ES Future
Increasing Variability
(100,0.6%,$65)
(60,0.75%,$65)
DAL
ES
(60,2%,$40)
(60,1.5%,$40)
Expected Regret ($ billions per year)
Probability of ES Future

FIGURE 7

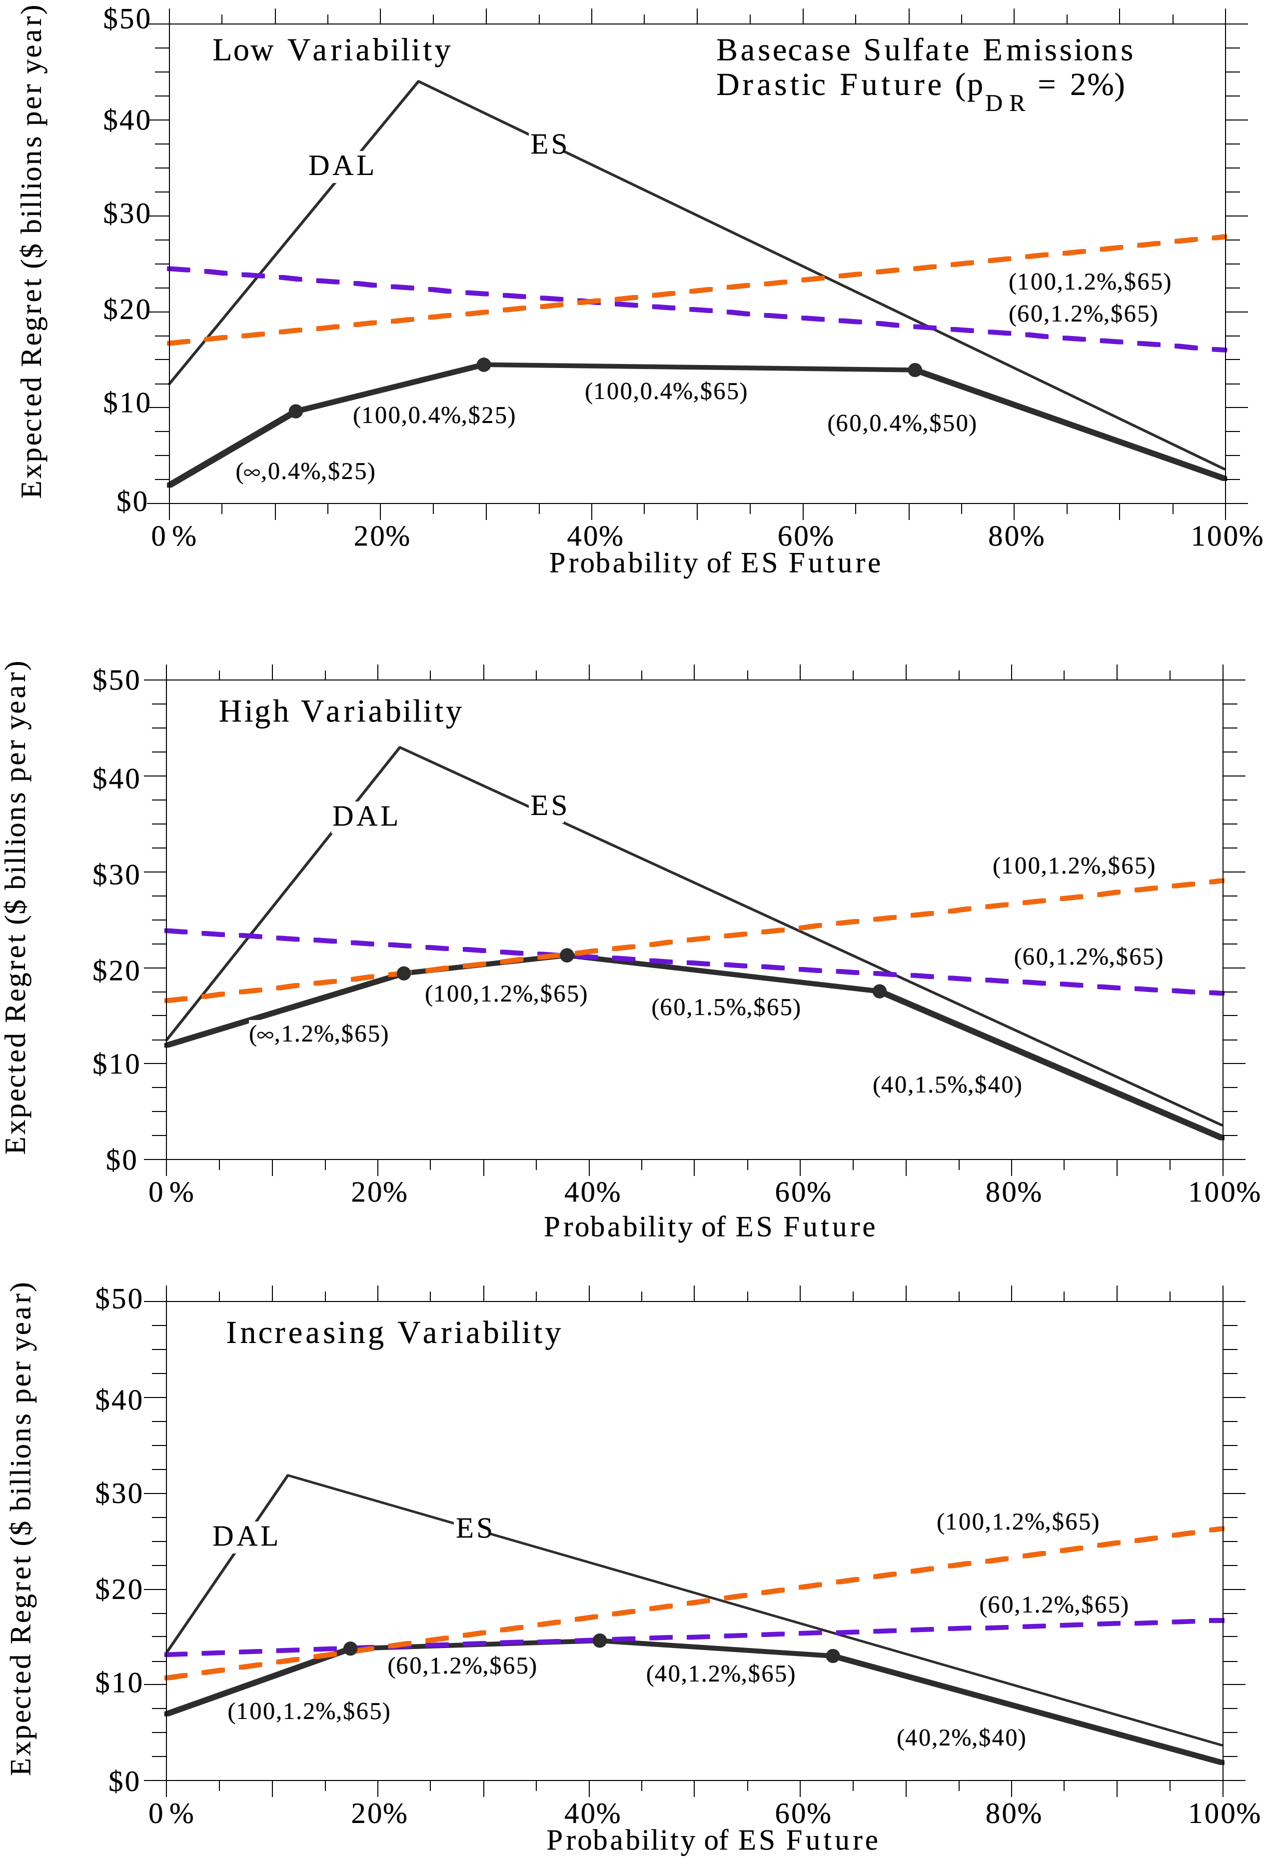

Low Variability
Basecase Sulfate Emissions
Drastic Future (p_{DR} = 2%)
DAL
ES
(100,1.2%,$65)
(60,1.2%,$65)
(100,0.4%,$65)
(100,0.4%,$25)
(60,0.4%,$50)
(∞,0.4%,$25)
Expected Regret ($ billions per year)
$50
$40
$30
$20
$10
$0
0 %
20%
40%
60%
80%
100%
Probability of ES Future
High Variability
DAL
ES
(100,1.2%,$65)
(60,1.2%,$65)
(100,1.2%,$65)
(60,1.5%,$65)
(∞,1.2%,$65)
(40,1.5%,$40)
Increasing Variability
DAL
ES
(100,1.2%,$65)
(60,1.2%,$65)
(60,1.2%,$65)
(40,1.2%,$65)
(100,1.2%,$65)
(40,2%,$40)

Figure 8

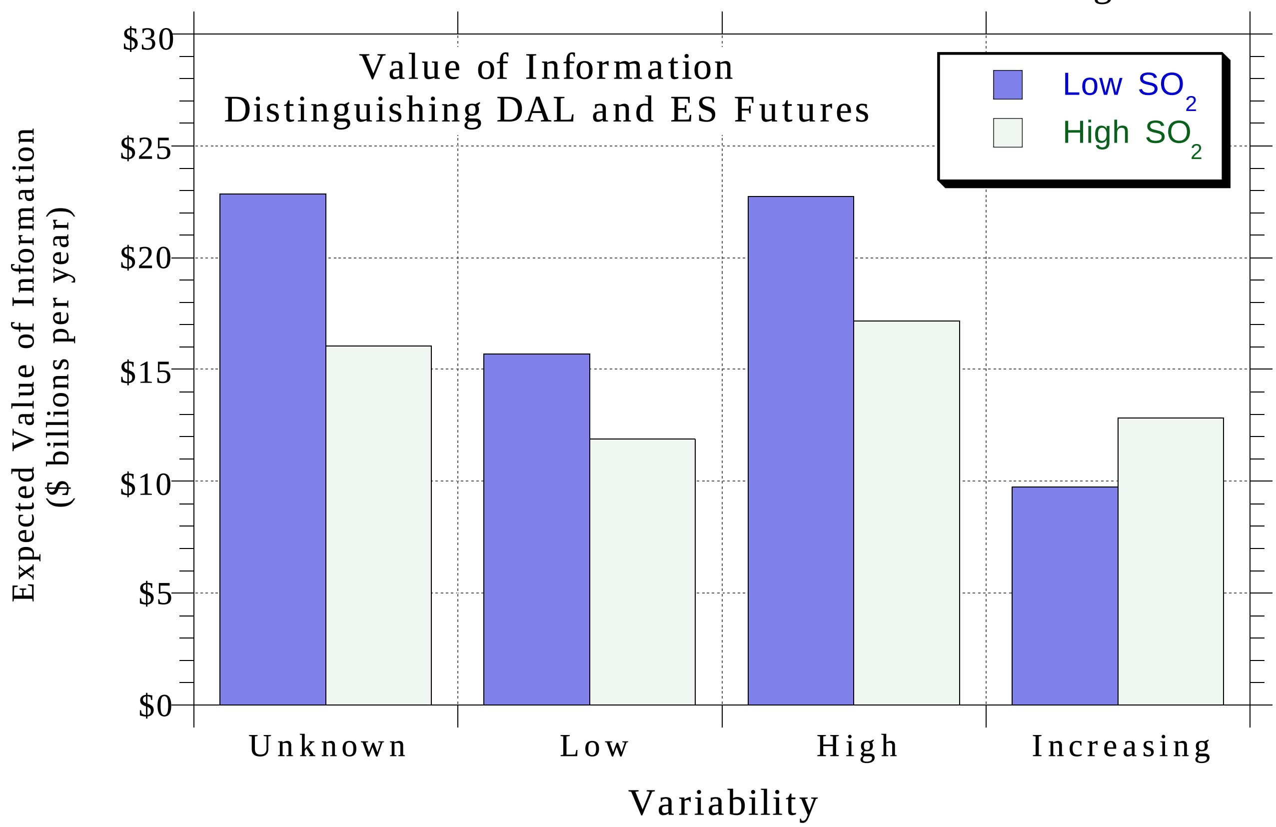

Figure 9