

Uncertainty, imprecision, and the precautionary principle in climate change assessment

M.E. Borsuk and L. Tomassini

Department of Systems Analysis, Integrated Assessment and Modelling (SIAM), Swiss Federal Institute of Environmental Science and Technology (EAWAG), Überlandstrasse 133, 8600 Dübendorf, Switzerland (E-mail: mark.borsuk@eawag.ch, lorenzo.tomassini@eawag.ch)

Abstract Statistical decision theory can provide useful support for climate change decisions made under conditions of uncertainty. However, the probability distributions used to calculate expected costs in decision theory are themselves subject to uncertainty, disagreement, or ambiguity in their specification. This imprecision can be described using sets of probability measures, from which upper and lower bounds on expectations can be calculated. However, many representations, or classes, of probability measures are possible. We describe six of the more useful classes and demonstrate how each may be used to represent climate change uncertainties. When expected costs are specified by bounds, rather than precise values, the conventional decision criterion of minimum expected cost is insufficient to reach a unique decision. Alternative criteria are required, and the criterion of minimum upper expected cost may be desirable because it is consistent with the precautionary principle. Using simple climate and economics models as an example, we determine the carbon dioxide emissions levels that have minimum upper expected cost for each of the selected classes. There can be wide differences in these emissions levels and their associated costs, emphasizing the need for care when selecting an appropriate class.

Keywords Climate change; cost-benefit analysis; decision theory; imprecise probability; precautionary principle; upper and lower probabilities

Introduction

In the third assessment report of the Intergovernmental Panel on Climate Change (IPCC), climate scientists predicted that global average temperature would increase between 1.4 and 5.8 °C by 2100 (IPCC, 2001). No assessment was made of the relative likelihood of intermediate warming values. This is because the participating scientists held divergent views on the magnitude of warming and believed that a single probability distribution could not capture this divergence of opinion. There was also a sense that probabilities, which are normally based on repeated experiments and frequencies of measured outcomes, could not be derived for such a singular event (Pittock *et al.*, 2001).

This situation leaves policy makers who want to plan for, or mitigate, climate change in the very difficult position of having to construe probability distributions for themselves. Expression of uncertainty in the form of a probability distribution is not simply an exercise in describing how well we know a particular quantity, but rather is the first step in reaching a rational decision. Informally, probabilities allow one to hedge decisions away from potentially large losses, with a “hedge factor” that depends on the amount of uncertainty (Reckhow, 1994). More formally, the probability distribution can be used to determine “expected cost” for decision options, calculated through integration of a cost function over the probability distribution (Raiffa and Schlaifer, 1968). The option can then be selected that minimizes the expected cost.

Application of the decision theoretic procedure described above to the climate change problem has been considered by Kann and Weyant (2000) and Tol (2003). Each has concluded that the representation of uncertainty is a critical determinant of the results of analysis. Unfortunately, as the IPCC statement above exemplifies, there is not agreement

on the interpretation of probabilities in this context, nor does the conventional theory seem to capture the variety of uncertainties encountered in practice. Methods for representing uncertainty in climate change assessment therefore require further development (Pittock *et al.*, 2001).

In this paper, we propose the use of sets of probability measures (Berger, 1984) to describe uncertainty in predictions of future climate change. This concept has the potential to capture ambiguity or disagreement in probability specification and may be a more realistic portrayal of the current state of scientific knowledge than precisely specified distributions. Kriegler and Held (2003) used sets defined as bounds on cumulative distributions to quantify imprecision in probabilistic climate change forecasts. However, other representations are also possible, and we believe that greater consideration should be given to the choice of method, especially in light of the full decision context. In this paper, we describe several useful set representations and demonstrate how each might be used to portray climate change uncertainties. We also discuss the implications for applied decision theory and suggest alternative decision criteria, such as an economic version of the precautionary principle, that may help resolve problems induced by high imprecision in probability specification.

Representation of climate change uncertainties

Much effort has been made recently to quantify uncertainties in climate change prediction (see reviews by Kann and Weyant, 2000; Katz, 2002). A typical approach is to represent uncertainty in future emissions of greenhouse gases (GHGs) by scenarios representing various assumptions about key socio-economic drivers such as population, economic growth, and energy technology. Emissions scenarios are then converted into atmospheric concentrations of GHGs and resultant temperature changes using one or more general circulation models (GCMs). Uncertainty in the predictions of these models is generally estimated by developing probability distributions for key parameters, which are then propagated through the models using a Monte Carlo method. Model structural uncertainty is usually assessed by generating and comparing results from multiple model formulations. For computational reasons, simplified models that emulate the behavior of more complex GCMs are often used. The result of such an analysis is usually a probability distribution for global or regional temperature increase corresponding to each emissions scenario. A representative example of this approach is the analysis of New and Hulme (2000).

Of course, the results of such uncertainty analyses are sensitive to the choice of probability distributions for the model parameters. Often, subjective judgment is the main determinant of this choice, and the expert elicitations of Morgan and Keith (1995) are a common source of such judgment. More recently, expert judgment has been combined with data-based estimates using Bayesian updating (Webster and Sokolov, 2000; Forest *et al.*, 2002; Webster *et al.*, 2003). However, due to the limited ability of historical data to resolve the value of key parameters, the distributions that result are still highly sensitive to the expert “priors.” This is disconcerting, especially considering the disagreement often found among experts.

Keith (1996) argues that, rather than combine the judgments of multiple experts, as has been done for most uncertainty analyses, the divergence of opinion should be maintained by propagating the distributions of the individual experts separately. While conceptually this may be an ideal solution, practically the presentation to policy makers of dozens of probability distributions for each emissions scenario is likely to overwhelm their analytical capabilities. It is not even clear how such sets of results can be used to reach a decision consistent with conventional decision criteria.

Bounding analysis has been proposed as a solution to this problem of disagreement and ambiguity (Keith, 1996). One form of bounding as it relates to probability distributions involves the concept of imprecise probabilities (Reichert, 1997). Rather than choosing a single, precisely defined distribution to describe an uncertain quantity, a set of distributions is employed. For example, Kriegler and Held (2003) chose to summarize the divergent results of the Bayesian analysis of Forest *et al.* (2002) using upper and lower bounds on the cumulative distribution function (CDF) of two key climate model parameters. These were then used to define random sets which were projected onto the range of global mean temperature increase using a simplified climate model. The resulting estimates of warming were very imprecise.

One reason for high imprecision in the results of Kriegler and Held (2003) may be their choice of set representation. Sets defined by upper and lower bounds on the CDF allow probability density functions that have sharp peaks at specific values. These densities may be implausible to climate experts and can be avoided with other representations. For example, sets might be defined by upper and lower bounds on densities, rather than cumulative distributions. The choice of an appropriate representation should depend on more detailed statements by the climate experts that provide the prior distributions, or on the specific assumptions employed in the Bayesian inference procedure. The intended use of the results, such as in a decision theoretic analysis, should also be considered in deciding what assumptions about sets are appropriate.

Because the assumptions behind the various set representations, or classes, are not always immediately clear, in the remainder of this section we outline some of the more useful classes. In the next section, we then apply these classes to the problem of climate change assessment in an effort to explore the implications of imprecision and class selection for decision making.

Parametric families

The parameters (e.g. mean, standard deviation, scale) of any distributional family can be specified using intervals, rather than precise values, thus defining a set of distributions in the family. Computations with parametric families are relatively straightforward and the results can be conveniently communicated. The main disadvantage of parametric classes is that they may fail to capture a wide range of realistically possible distributions. The assumptions that underlie the use of a specific parametric family are rather strong.

Probability box

Let L and U be nondecreasing functions from the real line to $[0,1]$ with $L(\theta) \leq U(\theta)$. Then a set of probability distributions called a probability box (or distribution band) can be defined by:

$$\Gamma_{L,U}^{PB} = \{ \pi : L(\theta) \leq F(\theta) \leq U(\theta) \}$$

where F denotes the cumulative distribution function (CDF) of the probability distribution π . In words, a probability box is a set of probability distributions, the CDFs of which are contained by given upper and lower bounds. Probability boxes will include distributions that have point masses, that is, probability distributions with discontinuous CDFs. Such distributions may be considered unreasonable in some cases. Methods for calculating bounds on expectations resulting from distribution bands are given by Basu and DasGupta (1995).

Quantile classes

A continuous parameter space Θ is partitioned into m disjoint segments, such that $\Theta = I_1 \cup \dots \cup I_m$. For $i \in \{1, \dots, m\}$, let l_i and u_i satisfy $l_i \leq u_i$ and $\sum l_i \leq 1 \leq \sum u_i$. Then a quantile class of distributions (Lavine, 1991) is defined as:

$$\Gamma_{l,u}^Q = \{ \pi : l_i \leq \pi(I_i) \leq u_i \}$$

In words, quantile classes are defined by placing upper and lower bounds on the probability that a parameter value lies within each of a finite number of intervals. Quantile classes are relatively easy to interpret, elicit from experts, and use to calculate bounds on expectations (Lavine, 1991). However, they too admit distributions which contain point masses at specific parameter values. They therefore tend to be “too broad,” especially in higher dimensions.

Density ratio classes

A density ratio class (DeRobertis and Hartigan, 1981) is defined as:

$$\Gamma_{l,u}^{DR} = \left\{ \pi : \frac{f(\theta)}{f(\theta')} \leq \frac{u(\theta)}{l(\theta')} \right\}$$

where f denotes the probability density of π (assuming it exists), and l and u are two bounded nonnegative functions, such that $l(\theta) \leq u(\theta)$. In words, density ratio classes are specified by placing bounds on the ratio of probabilities. They can also be interpreted as sets of measures with (unnormalized) densities between given upper and lower bounds. A set of probability measures is then obtained by normalizing these measures. In few dimensions, the calculation of bounds on expectations for density ratio classes is relatively straightforward (DeRobertis and Hartigan, 1981). They enjoy convenient mathematical properties such as marginalization invariance, which makes it possible to reduce high dimensional problems to the one dimensional case and to propagate classes through functions while maintaining class structure (Wasserman and Kadane, 1992). The major disadvantage of this class is that it may be difficult to elicit.

Density bounded classes

A density bounded class (Lavine, 1991) is defined as:

$$\Gamma_{l,u}^{DB} = \{ \pi : l(\theta) \leq f(\theta) \leq u(\theta) \}$$

where again f denotes the density of the probability distribution π , and l and u are two bounded nonnegative functions, such that $l(\theta) \leq u(\theta)$. In words, density bounded classes are specified by placing upper and lower bounds on densities. The resulting class is the set of all normalized densities between these bounds, thus avoiding densities with extreme peaks or other possibly unreasonable features. Bounds on expectations can be calculated according to Lavine (1991). Elicitation and interpretation is simpler than for the density ratio class.

ϵ -contaminated classes

For a fixed distribution π_o and $\epsilon \in [0, 1]$, an ϵ -contaminated class is defined as:

$$\Gamma_{l,u}^{EC} = \{ \pi = (1 - \epsilon) \cdot \pi_o + \epsilon \cdot q : q \in Q \}$$

where Q is an alternate set of distributions. The ϵ -contaminated classes are easy to work with and have a rather logical interpretation. However, if Q is a large class of distributions (for example the set of all probability distributions), the ϵ -contaminated class is

also large, especially in high dimensions and with large values of ε . Elicitation can be straightforward (depending on Q) because only the reference distribution π_0 and the value of ε have to be determined, but this makes it not very flexible.

Climate change example

To simplify comparison with other published studies, we follow the typical procedure for uncertainty analysis exemplified by [New and Hulme \(2000\)](#), [Webster et al. \(2003\)](#), and [Kriegler and Held \(2003\)](#). Explicit emissions scenarios and key model parameters are used as inputs to a simplified deterministic climate model which calculates atmospheric concentrations and global mean temperature increase. Emissions scenarios correspond to a “business as usual” scenario ([Table 1](#)) and various reduction scenarios, assumed to be a constant percentage of the business as usual scenario in each of the next 100 years. Only emissions of carbon dioxide are considered, this being the main GHG and a possible proxy for all other GHGs.

Climate response model

Because of our interest in incorporating cost functions into our analysis, the climate model we use was derived from the cost-benefit analysis of [Maddison \(1995\)](#). We programmed the dynamic non-linear model of Maddison and summarized the results as simple functions through which all the probability classes described in the previous section could be easily propagated. In this way, atmospheric carbon dioxide concentration in the year 2100, $[CO_2]_{2100}$, is estimated as

$$[CO_2]_{2100} = 650 - 3.126 \cdot P \tag{1}$$

where $[CO_2]_{2100}$ is in ppv, and P is the annual percent emissions reduction. Equation (1) implies that under the “business as usual” case, the carbon dioxide concentration in 2100 will be 650 ppmv. The corresponding increase in global mean surface air temperature, ΔT_{2100} , is then estimated as

$$\Delta T_{2100} = \frac{\Delta T_{2x}}{\ln 2} \ln \left(\frac{[CO_2]_{2100}}{[CO_2]_{1990}} \right) + 0.255 \tag{2}$$

where $[CO_2]_{1990}$ is the carbon dioxide concentration in 1990 (approximately 350 ppmv), and ΔT_{2x} is the climate sensitivity, defined as the increase in temperature resulting from a doubling in atmospheric CO_2 concentration relative to 1990 levels. Climate sensitivity

Table 1 Gross world national product (trillion \$US/yr) and global carbon dioxide emissions (GtC/yr) assumed for the business as usual scenario (from [Maddison, 1995](#))

Year	GNWP (trillion 1990\$US/yr)	Emissions (GtC/yr)
1990	22.92	7.5
2000	30.97	8.5
2010	39.77	10
2020	51.06	11.4
2030	66.55	12.6
2040	81.94	13.6
2050	100.89	14.5
2060	122.73	15.6
2070	149.31	16.8
2080	180.64	18
2090	217.35	19.2
2100	261.52	20.3

has been identified as the most important uncertain model parameter and will be the focus of our uncertainty analysis.

Economic model

The costs of climate change consist of two principal components: emissions abatement costs and warming-induced damage costs. Using Maddison's (1995) survey of various abatement cost modeling studies and his projected global economic growth (Table 1), we derived the following summary relationship, expressed as a function of annual percent emissions reduction, P ,

$$ABCOST_{2100} = 0.001 \cdot P^3 \quad (3)$$

where $ABCOST_{2100}$ is the cumulative abatement cost by the year 2100 in trillions of constant 1990 US dollars, assuming no future discounting. Total damage costs were also derived from a survey of Maddison (1995) and can be approximated as

$$DAMAGE_{2100} = 14.729(\Delta T_{2100})^2 + 24.636(\Delta T_{2100}) + 23.924 \quad (4)$$

where $DAMAGE_{2100}$ is the cumulative damage cost by the year 2100 in trillions of constant 1990 US dollars, again assuming no discounting.

Without doubt, the climate and cost models represented by equations (1) – (4) are exceedingly simple. We do not intend the present analysis to be realistic. Rather, our goal is to demonstrate the development, propagation, and use of various imprecise probability representations within the general context of climate change. For these purposes, the simple models given above suffice. We leave it to experienced climate scientists and economists to substitute their own more complex, and possibly more realistic, models into this framework.

Class specification

As mentioned above, the climate sensitivity, ΔT_{2x} , is a critical model parameter because it is uncertain and model results are very sensitive to it. Therefore, as in most other uncertainty analyses of climate change (New and Hulme, 2000; Andronova and Schlesinger, 2001; Krieger and Held, 2003; Webster *et al.*, 2003), uncertainty about this parameter will be the focus of our study. All other assumptions, including future economic growth, carbon dioxide emissions, and other model parameter values will be assumed to be known with certainty. Also, as with other studies, we use the elicitations of Morgan and Keith (1995) as the basis for constructing sets of distributions on ΔT_{2x} . Morgan and Keith (1995) interviewed 16 climate experts, assessing points on the cumulative distribution function of ΔT_{2x} . However, expert #5 gave responses that differed fundamentally from the others, both in magnitude and degree of uncertainty. Therefore, to reduce complications in this simple example, we exclude the assessment of this expert from our analysis. Some of the experts estimated points corresponding to every 5th percentile, while others only estimated every 10th percentile. For consistency, we only used the 10th percentile estimates for all experts (Figure 1).

Application of the various classes to summarizing the responses of the experts should incorporate information from the experts themselves about the appropriateness of the defining assumptions. However, in our case, we only had access to the information contained in Morgan and Keith (1995). Therefore, each of the classes should be viewed simply as attempts to use the salient features of the experts' distributions to construct a set which contains all distributions that have those features. The rationale could be that the interviewed experts are representative of the set of all experts who might provide

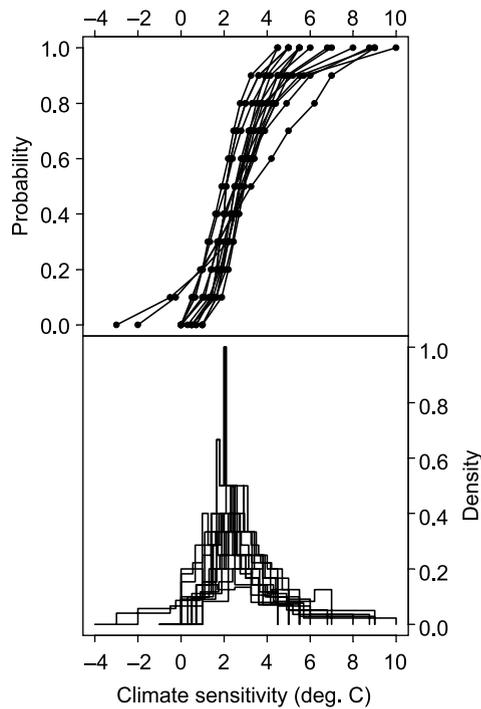


Figure 1 Elicited probability distributions of climate sensitivity shown as cumulative distribution functions (top) and probability densities (bottom). The solid points represent assessed values corresponding to every 10th percentile. (Data from Morgan and Keith, 1995)

assessments and that we would like to extend our results to represent that broader set. Alternatively, the various classes can be seen as attempts to “robustify” the assessments that are available. In either case, each of the different classes focuses on different salient features of the assessments to define the set, as follows.

- **Parametric family** – It was assumed that all distributions in the set are Gaussian, and means and standard deviations were fitted to the elicited percentiles by minimizing the Kolmogorov-Smirnov measure for each expert. The maximum and minimum values of each parameter across the experts ($\mu = [1.86, 3.48]$, $\sigma = [0.95, 1.97]$) were then used as bounds to specify a set of Gaussian distributions, assuming independence between the parameters.
- **Probability boxes** – It was assumed that the minimum and maximum values of ΔT_{2x} across the experts for each assessed percentile provide bounds on a set of cumulative distributions, which could be fully defined by linear interpolation between the assessed points.
- **Quantile classes** – Quantile classes assume that the experts’ assessments provide bounds on the probability that the value of climate sensitivity is within certain defined intervals. In this case, the ΔT_{2x} range was divided into one degree intervals and the maximum and minimum CDF difference across the expert responses was calculated for each interval. These values were then taken to represent the upper and lower bounds on the probability for each interval.
- **Density ratio classes** – The density ratio class was specified by linearly interpolating between the points on the CDF assessed by each expert and differentiating to yield corresponding densities. Maximum and minimum densities across the experts were then calculated at each point and interpreted as upper and lower bounds on a set of unnormalized densities.

- **Density bounded classes** – The bounds used to specify the density ratio class were also used to define the density bounded class. However, for this class only normalized densities between these bounds were considered.
- **ϵ -contaminated classes** – To specify an ϵ -contaminated class, the average density at each value of ΔT_{2x} was calculated and used as the reference distribution π_0 . The set of all distributions between the minimum and maximum values of ΔT_{2x} deemed possible by the experts, $[-3, 10]$, was then used as the alternate set, Q . The contamination factor, ϵ , was set at an arbitrary level of 0.25.

Computational methods

A number of properties of our climate model simplify the necessary computations. The monotonic relation between climate sensitivity and global warming (Eq. 2) means that class types are preserved upon propagation through the climate model. That is, if uncertainty in climate sensitivity is described by a particular class, then the resulting uncertainty in global temperature can be described by the same class. This means that the methods described by Lavine (1991) can be used to calculate the upper and lower bounds on expected cost over the sets of probability measures derived for global temperature. Additionally, in our model, only the damage cost is subject to uncertainty in temperature, as the abatement cost is fixed for a given level of emissions reductions. Because the damage cost function is monotonic, identification of the “critical” distributions within each class that determine the upper and lower bounds on expectations is relatively straightforward (Lavine, 1991). While the model assumptions leading to these simplifications may not be entirely realistic, they improve the didactic value of this example.

Results

If the relationship between carbon dioxide emissions and global warming were precisely determined, and economic cost were the only consideration, then the optimal emissions reduction (24%) would be at the minimum of the total cost function (Figure 2, vertical dotted line). When uncertainty in climate sensitivity is considered and represented by a conventional probability distribution on global temperature

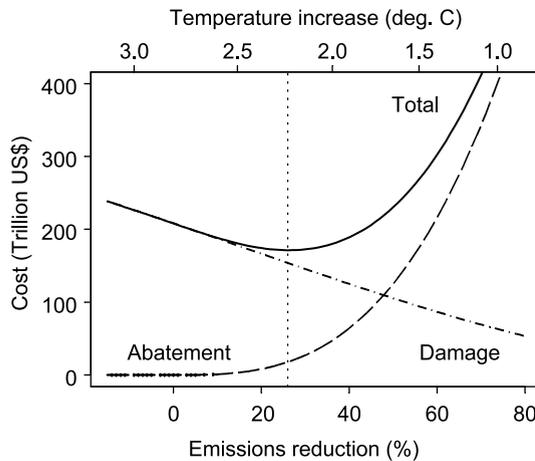


Figure 2 Abatement cost as a function of emissions reduction (dashed line, bottom axis), damage cost as a function of global temperature increase (dot-dashed line, top axis), and total cost as a function of both (solid line) assuming no uncertainty between emissions and temperature. The vertical dotted line represents the emissions reduction and corresponding temperature increase with minimum total cost, assuming no uncertainty

rise can be estimated and used to calculate expected damage cost for each emissions level. This expected damage cost can then be added to the deterministic abatement cost to yield the expected total cost. The optimal emissions reduction level would then be the one that minimizes this expected total cost. In this case, the result of representing ΔT_{2x} by a precise density function (the density, π_0 , used in the ε -contaminated class) is that a slightly greater emissions reduction (26%) would be chosen to offset the damage costs associated with the possibility of large temperature increases. The expected total cost for this case would be \$197.2 trillion.

When sets of probability measures are used to describe imprecision in the distribution of ΔT_{2x} , the expected total cost for each level of abatement is no longer unique. Rather, upper and lower bounds on expectations are the result, and these bounds may be highly sensitive to the choice of class representation (Figure 3). This is because each class is more or less liberal in the probability distributions it allows. For example, the quantile class, which leads to the widest interval between upper and lower expected costs, allows relatively high probability to be assigned to very low or high values of ΔT_2 , in part because of the imprecision inherent in discretization of the elicitation process (Figure 4). Similarly, the density ratio class can assign greater relative probability to the extreme values, because it is not restricted to containing normalized densities. In this case, the critical density function is even bimodal. The critical distributions of the ε -contaminated class assign all of the “contamination probability” to the allowable extreme points, also leading to very wide expectation bounds. The Gaussian family, however, is relatively narrow tailed, and the critical upper density – the one with the greatest mean and greatest variance – does not assign much likelihood to high values of ΔT_2 , thus limiting the upper expected cost. In this case, the probability box also leads to tight bounds relative to some of the other classes. However, this may not be true in general when more complex cost functions are used that are sensitive to the possibility of point masses.

For the purposes of decision-making, when a given class leads to intervals on expectations that overlap for two or more emissions reduction levels, then conventional decision theory does not provide a basis for choosing between them. Alternate decision criteria are required to arrive at a unique solution. Cheve and Congar (2002) recommend using a criterion of maximum lower expected utility, or, as it applies in this case,

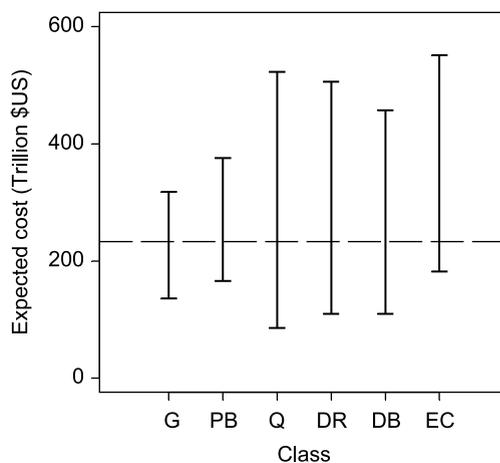


Figure 3 Upper and lower expected total cost for a 0% emissions reduction for each of the six classes (G = Gaussian, PB = probability box, Q = quantile, DR = density ratio, DB = density bounded, EC = ε -contaminated). The horizontal dashed line represents the expected cost resulting from a precise distribution on climate sensitivity, and a 0% emissions reduction

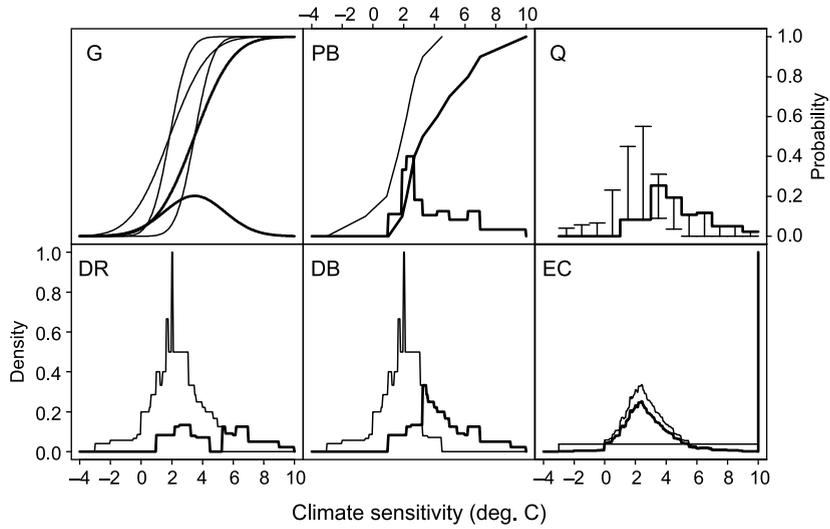


Figure 4 Representations of each of the selected classes (G = Gaussian, PB = probability box, Q = quantile, DR = density ratio, DB = density bounded, EC = ϵ -contaminated). Thin lines represent bounds defining the sets, and thick lines represent the critical distributions determining the upper expected damage cost for a 0% emissions reduction. Both cumulative probability and probability density are shown for the critical distributions for G and PB. For EC, the critical distribution assigns point mass of 0.25 to a value of 10 (shown as a vertical line)

minimum upper expected cost. They show that adoption of this criterion leads to a decision that satisfies an economic interpretation of the precautionary principle. That is, if the chosen decision is optimal with respect to this criterion, then whatever distribution may eventually realize, the decision maker cannot be reproached *a posteriori* for lack of precaution. Cheve and Congar (2002) also show that no other decision exists that satisfies this principle. That is, if the decision chosen is not optimal with respect to this criterion, then there exists at least one probability distribution in the set for which the decision-maker can be shown to exhibit a lack of precaution if it realizes. As the precautionary principle is often cited as being appropriate for situations of environmental risk (CEU, 2000), we adopt the criteria of minimum upper expected cost as the basis for choosing the appropriate emissions reduction level for each class representation (Figure 5).

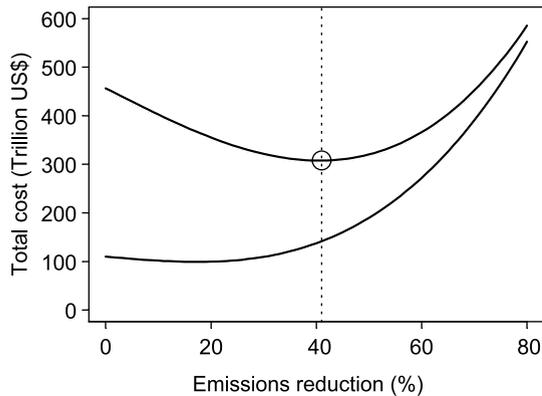


Figure 5 Upper and lower bounds on expected total cost for the density bounded class, as a function of emissions reduction. The vertical dotted line represents the emissions reduction with minimum upper expected cost (shown by an open point)

More liberal classes generally lead to higher upper expected damage costs for a given emissions level. These damage costs can be at least partially reduced through further emissions reduction. This means that according to the precautionary principle, the chosen reduction will be somewhat greater for these liberal classes, and the minimum upper expected total cost somewhat higher (Table 2). There is as much as a 10.6% emissions and \$106.8 trillion difference between the chosen reduction levels under the various classes, and an 18.5% emissions and \$146 trillion difference between the ε -contaminated class and a precise distribution. This last difference can be considered the price of precaution.

Discussion

Despite being based upon the same set of expert elicitations, the different distribution classes used in this analysis lead to very different upper and lower bounds on the expected total cost of climate change. However, each class also incorporates additional assumptions which cannot be supported or refuted using the available information. It is precisely these assumptions that are responsible for the different expectation bounds. It does not appear that, for the climate change model, the probability box representation chosen by Krieglger and Held (2003) leads to greater imprecision than some of the other possible classes. However, the choice of class should not be arbitrary, but should be justified by acknowledgment of the underlying assumptions. For example, are we willing to allow sharp peaks or point masses? Is there reason to constrain ourselves to distributions of a particular parametric family? Should we assume that the density function must be unimodal? Can we state assumptions about the ratio of possible densities? Is probability assessment facilitated by discretization of the uncertain quantity? The answers to these questions will point to particular suitable classes.

Computational complexity may be another important factor to consider. In the example presented here, the climate and cost models were relatively simple. The assumption of a monotonic dependence of global temperature on climate sensitivity preserves class type upon propagation of uncertainty through the climate model. This in turn allowed for the application of established techniques to determine the expectation bounds for the different classes. Some of the classes, such as the density ratio class, enjoy this property of marginalization invariance in general, even for non-monotonic transfer functions. Others, such as the Gaussian family, do not. The monotonicity of the damage cost function also greatly simplified the calculations of expectation bounds. However, in one dimension all the classes can be handled numerically for more general cost functions as well. Nevertheless, calculations with probability boxes in particular can become quite involved. In higher dimensions, all calculations quickly become computationally expensive or even impracticable. An assumption of independence between parameters can ease the computational burden in such situations.

Table 2 Minimum upper expected total cost and corresponding emissions reduction for each considered class

Class	Minimum upper expected total cost (Trillion \$US)	Emissions reduction (%)
Gaussian family	236.3	33.7
Distribution bands	265.5	36.7
Quantile	337.4	43.5
Density ratio	329.4	42.6
Density bounds	307.7	40.7
ε Contam.	343.2	44.5
Precise distribution	197.2	26.0

We adopted the criterion of minimum upper expected cost to resolve the decision dilemma introduced when imprecise probabilities preclude a unique choice based on minimum expected cost alone. This criterion is not a part of conventional decision theory but is consistent with an economic statement of the precautionary principle. Other criteria would also be possible, such as minimum lower expected cost or minimum upper regret (Cheve and Congar, 2002). However, these do not seem appropriate for the climate change problem. Another option would be to consider the distribution of the predicted difference between two options, rather than a distribution for each option separately. For example, the distribution of the predicted difference in temperature between a given emissions reduction scenario and the business as usual scenario could be estimated. If the two have some of the same sources of uncertainty, then the distribution of the difference may be significantly more precise than the distribution of either alone (Reichert and Borsuk, 2005). This may improve a decision maker's confidence that one option is better than the other, even if the actual outcome is not well-known.

In our analysis, the sets of distributions on the climate parameter ΔT_2 were based upon expert elicitation, without any further comparison against data. However, these sets could also serve as classes of priors in a process of Bayesian inference (Berger, 1984). That is, a comparison against actual data or model results could be used to modify prior beliefs, as was done for precise prior distributions by Forest *et al.* (2002) and Webster *et al.* (2003). In this way, representation of climate sensitivity and predictions of global warming can keep up with advances in scientific knowledge.

Conclusions

Uncertainty is an important consideration in model-based climate change decisions. However, there is often disagreement or ambiguity regarding the appropriate probability distributions to use in describing this uncertainty. The use of sets of probability measures provides a potential solution but requires careful attention to practical and conceptual considerations to choose an appropriate class. We have shown that different classes can lead to very different decisions when minimum upper expected cost is employed as the decision criterion. Clearly, more research is required into the appropriate class and set specifications for key model parameters. A more realistic climate model should also be adopted before accurate results can be obtained.

References

- Andronova, N. and Schlesinger, M. (2001). Objective estimation of the probability density function for climate sensitivity. *Journal of Geophysical Research*, **106**(D19), 22605–22611.
- Basu, S. and DasGupta, A. (1995). Robust Bayesian analysis with distribution bands. *Statistics & Decisions*, **13**(4), 333–349.
- Berger, J. (1984). The robust Bayesian viewpoint (with discussion). In *Robustness of Bayesian Analysis*, Kadane, J. (ed.), North-Holland, Amsterdam, pp. 63–124.
- CEU (2000). *Communication from the Commission on the Precautionary Principle*. COM(2000)-1, Commission of the European Union.
- Cheve, M. and Congar, R. (2002). Managing environmental risks under scientific uncertainty and controversy. In *International Conference on Risk and Uncertainty in Environmental and Resource Economics*, van Ierland, E.C., Weikard, H.P. and Wesseler, J. (eds), Wageningen, The Netherlands.
- DeRobertis, L. and Hartigan, J. (1981). Bayesian inference using intervals of measures. *The Annals of Statistics*, **9**(2), 235–244.
- Forest, C., Stone, P., Sokolov, A., Allen, M. and Webster, M. (2002). Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science*, **295**(5552), 113–117.
- IPCC (2001). *Intergovernmental Panel on Climate Change, Third Assessment Report of Working Group I: The Science of Climate Change*, Cambridge University Press, Cambridge, UK.

- Kann, A. and Weyant, J. (2000). Approaches for performing uncertainty analysis in large-scale energy/economic policy models. *Environmental Modeling & Assessment*, **5**(1), 29–46.
- Katz, R. (2002). Techniques for estimating uncertainty in climate change scenarios and impact studies. *Climate Research*, **20**(2), 167–185.
- Keith, D. (1996). When is it appropriate to combine expert judgments? *Climatic Change*, **33**(2), 139–144.
- Kriegler, E. and Held, H. (2003). Climate projections for the 21st century using random sets. In *Third International Symposium on Imprecise Probabilities and Their Applications*, Bernard, J.M., Seidenfeld, T. and Zaffalon, M. (eds), Carleton Scientific, Lugano, Switzerland, pp. 345–360.
- Lavine, M. (1991). An approach to robust Bayesian analysis for multidimensional parameter spaces. *Journal of the American Statistical Association*, **86**(414), 400–403.
- Maddison, D. (1995). A cost-benefit analysis of slowing climate change. *Energy Policy*, **23**(4-5), 337–346.
- Morgan, M. and Keith, D. (1995). Subjective judgments by climate experts. *Environmental Science & Technology*, **29**(10), A468–A479.
- New, M. and Hulme, M. (2000). Representing uncertainty in climate change scenarios: a Monte-Carlo approach. *Integrated Assessment*, **1**(3), 203–213.
- Pittock, A., Jones, R. and Mitchell, C. (2001). Probabilities will help us plan for climate change. *Nature*, **413**(6853), 249.
- Raiffa, H. and Schlaifer, R. (1968). *Applied Statistical Decision Theory*, John Wiley & Sons, New York.
- Reckhow, K. (1994). Importance of scientific uncertainty in decision-making. *Environmental Management*, **18**(2), 161–166.
- Reichert, P. (1997). On the necessity of using imprecise probabilities for modelling environmental systems. *Water Science and Technology*, **36**(5), 149–156.
- Reichert, P. and Borsuk, M. (2005). Does high forecast uncertainty preclude effective decision support? *Environmental Modelling & Software*, **20**, 991–1001.
- Tol, R. (2003). Is the uncertainty about climate change too large for expected cost-benefit analysis? *Climatic Change*, **56**(3), 265–289.
- Wasserman, L. and Kadane, J. (1992). Computing bounds on expectations. *Journal of the American Statistical Association*, **87**(418), 516–522.
- Webster, M., Forest, C., Reilly, J., Babiker, M., Kicklighter, D., Mayer, M., Prinn, R., Sarofilm, M., Sokolov, A., Stone, P. and Wang, C. (2003). Uncertainty analysis of climate change and policy response. *Climatic Change*, **61**(3), 295–320.
- Webster, M. and Sokolov, A. (2000). A methodology for quantifying uncertainty in climate projections. *Climatic Change*, **46**(4), 417–446.