Chapter 2
The Value of Information in a Risk Management Approach to Climate Change

Carolyn Kousky and Roger M. Cooke

Abstract The standard economic approach to analyzing the climate change problem has been to search for efficient abatement policies. The massive uncertainties and the possibility for cataclysmic climate damages, however, suggest that a risk management approach is more appropriate. This shifts the policy question to how much risk of catastrophe society is willing to accept. Intuitively, this change in focus may shift our information needs, and the needs should be assessed through a value-of-information analysis. Such calculations should allow for examination of how improved information alters the probability of exceeding a given policy target, incorporate rigorous expert judgment for determining beliefs and quantifying uncertainties, and highlight what scientific information is most valuable for a policymaker attempting to keep the probability of catastrophic climate impacts below a set threshold. We discuss how Bayesian belief nets can be useful tools for this type of analysis.

Keywords Bayesian belief nets • Catastrophic climate change • Climate policy • Integrated assessment • Risk management • Value at risk • Value of information

2.1 Introduction

An early, and still predominant, economic approach to climate change has been to treat it as the challenge of pricing a large, global externality. The focus has been on economic efficiency and determining optimal emissions trajectories using integrated assessment models (IAMs) where the avoided damages of climate change could be

---

C. Kousky (✉)
Resources for the Future, Washington, DC, USA
e-mail: kousky@rff.org

R.M. Cooke
Resources for the Future, Department of Mathematics, TU Delft, Delft, The Netherlands
e-mail: cooke@rff.org

compared with abatement costs. Given the possibility of catastrophic impacts and the myriad uncertainties surrounding climate change, we, like several other authors, argue for a risk management approach, not an efficiency approach.

A risk management approach asks what policy should be, given the large range of possible outcomes from that choice. This is quite distinct from asking what the optimal policy is under different assumptions of our uncertain variables. Drawing an analogy to risk management in the insurance and financial sectors, society may wish to keep the probability of facing catastrophic damages to some determined low level. This change in focus to a risk management paradigm dramatically shifts our information needs. A risk management approach highlights the need for research on the possibility of climate catastrophes, their likelihood under various emissions scenarios, and whether we can detect impending catastrophes soon enough to avert them.

There are many places where we can improve our climate information to improve climate risk management, raising the question of where to spend scarce research dollars and when it is worth waiting for better information. Basic value-of-information models make clear that information will be valuable only when (1) the possible policy options perform quite differently in different states of the world; (2) our current beliefs would lead us to pick an option that is worse than what we would do with better information; and (3) we can undertake some measurement, the result of which would shift our belief substantially enough to change our preferred policy.

Bayesian belief nets (BBNs) can be used to take a distinctly risk management approach to the task of valuing improved climate information. BBNs are graphical models of the dependencies between multiple variables. They can be used to calculate how improved information on one or more parameters would change the estimated probability of meeting a given policy target, as well as how improved knowledge alters welfare estimates. BBNs can also incorporate expert judgment for determining beliefs and quantifying uncertainties and highlight what scientific information is most valuable for a policymaker taking a risk management approach, as opposed to an efficiency approach, to the climate problem.

The next section introduces what we consider to be a risk management approach to climate change. Section 2.3 offers an overview of a basic value-of-information framework. In Sect. 2.4 we move on to discussing how BBNs can be used to conduct a value-of-information analysis for the climate problem while incorporating a distinctly risk management flavor of analysis. Section 2.5 concludes.

2.2   A Risk Management Approach

As is known in finance, an increase in expected rewards usually carries with it an increase in risks. Prudent firms in the banking and insurance industries often manage the risk of insolvency using a value-at-risk (VaR) approach. A firm chooses a target solvency probability (or one is set for it through regulation) and then ensures that the risk of insolvency does not exceed this target, through, for example, building capital reserves or reducing exposure. So too with climate change, the benefits of increased economic growth from a carbon economy carry with them

risks of negative climate impacts, some of which could be quite catastrophic. When the uncertainties and nontrivial probability of catastrophic outcomes are recognized, it can change preferred policy choices; in these cases, some amount of abatement in the near term as a hedging strategy becomes optimal (see, e.g., Manne and Richels 1995; Lempert et al. 2000).

Following the VaR approach used in the private sector, society could choose to limit the risk of a climate-induced "insolvency." This would be some form of collapse in social welfare—a worst-case scenario whose probability should be kept beneath a defined tolerable level. The policy questions then become, first, what the worst-case outcome is we wish to avoid, and second, how much risk of such an outcome we are willing to tolerate. Regulations for the banking and insurance industries in the European Union dictate the solvency threshold for firms at around 1-in-200. We are currently taking much larger risks of large-scale climate damages than this.

In a risk management approach, then, fully assessing and clearly communicating the uncertainties become essential for policy. Too many studies conducted under an efficiency approach to the climate problem include the uncertainties as a caveat, and too many policymakers dismiss the uncertainties of modeling as fine print. In a world of climate risk management, the size and nature of these uncertainties and our attitudes toward risk determine the optimal amount of abatement today. This requires undertaking a complete uncertainty analysis with current climate models.

It is also the case that a risk management approach highlights different information needs. The correct discount rate becomes less important than an improved understanding of the nature of catastrophic consequences, their likelihood under differing emissions scenarios, our ability to detect tipping points before a catastrophe materializes, and the time frame for response should we pass such tipping points. Although we have some information on catastrophic impacts—for instance, numerous studies point to catastrophic consequences if global temperature exceeds 5 °C above preindustrial conditions, or even above 2.5 °C (Keller et al. 2008)—in general, we have a fairly poor understanding of the tail of the climate change damage distribution.

Satellite data are critical to this type of research. For instance, satellites can be used to document the trends that could be indicative of climate tipping points, such as melting of ice in Antarctica or the amount of methane in the atmosphere. They can also be used to look at effects as diverse as ocean acidification and desertification. This information, however, is under threat because the number of earth-observing satellites is declining, not increasing. Even rough model calculations of the value such information satellites provide in terms of detecting tipping points to avoid catastrophe could be useful for Congress when lawmakers consider appropriating more money to observation systems.

2.3   Value-of-Information Refresher

It is useful to recall a basic model of the value of information. Assume we can choose one of a set of available policy options, and that each option has a well-defined outcome with well-defined utility in each possible state of the world. If the

future state of the world were known, we would simply choose the option that would generate the highest utility. Unfortunately, the state is not known, and so we must quantify our uncertainty and then choose the option with the highest expected utility, given our beliefs about the state of the world.[1] Now, suppose we have the opportunity to perform an observation before choosing a policy, which will produce information to alter our beliefs about the likely state of the world. The observation may incline us to choose a different policy than we would have chosen without the observation. A simple result in decision theory states that it is never disadvantageous to perform a cost-free observation before choosing. That does not mean, however, that it is always worth spending money to obtain more information. The value of information quantifies the expected gain of performing this observation, relative to the given set of policy options.

A simple example clarifies the basic properties of the value of information. Suppose we have to choose between three climate policies: (1) business as usual (BAU) with no abatement; (2) tempered abatement (a little now with the possibility of more later); and (3) maximal abatement now. Suppose for illustration that there are two possible states of the world: climate sensitivity (cs) = 1.5 and cs = 5. The value cs = 1.5 corresponds to the most sanguine value given in the Fourth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC 4), and 5 is a very pessimistic value "which cannot be excluded" (IPCC 2007). BAU produces high utility if cs = 1.5, as no money is wasted on unnecessary abatement. It is catastrophic if cs = 5. The opposite holds for maximal abatement: it avoids catastrophe if cs = 5 but is very wasteful if it turns out that cs = 1.5. Tempered abatement is intermediate. Since in this simple example Prob$\{cs = 1.5\}$ = 1 − Prob$\{cs = 5\}$, the expected utility of each policy is a linear function of Prob $\{cs = 1.5\}$, as shown in Fig. 2.1.

According to our assessed probability of the event $\{cs = 1.5\}$, one of these three options will be optimal. Figure 2.1 shows that for a belief point Prob$\{cs = 1.5\}$ = 0.66, tempered abatement is optimal, but if the probability were a little higher, say 0.7, the preference would shift to BAU. Of course, we would like to know the true cs before choosing. If we could simply observe this number, then we would obviously choose BAU if cs = 1.5 and choose maximal abatement if cs = 5. Without knowing the outcome of this hypothetical measurement, we can compute the expected value of observing cs before choosing by drawing the thin dotted line in Fig. 2.1. The difference between this line and the maximum expectation of our policy options at our belief point is called the *value of perfect information*, for this belief point and these policy options.

---

[1] Much ink is spilled over whether we should choose according to the principle of maximal expected utility. We assume for the present discussion that the decider is a rational agent, in the sense of Savage (1954). A rational agent's preferences can always be decomposed into a unique probability over states of the world and an affine unique utility over consequences such that preferences are modeled as expected utility.
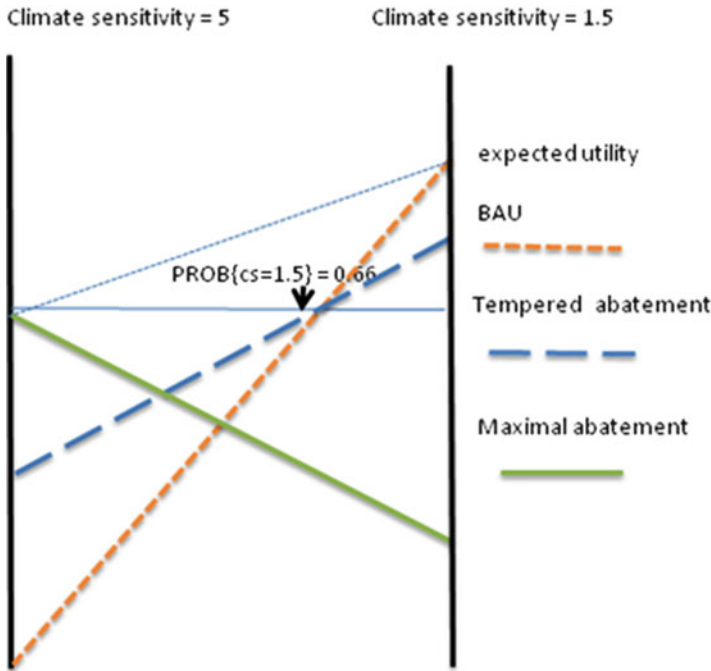
Fig. 2.1 Value of information for simple climate policy choice problem, with value of perfect information

Unfortunately, we are seldom afforded the possibility of performing a perfect observation. The best we can do in practice is find (costly) observations whose possible outcomes would alter our beliefs. Suppose scientists can undertake a study to get better but still imperfect information on the climate sensitivity. Keeping the example simple, suppose the possible outcomes of such a measurement are either HI or LO. Experts agree that if we observe LO, then the probability that $\{cs = 1.5\}$ $= 0.75$, whereas if we observe HI, then the probability of $\{cs = 1.5\} = 0.3$. It is easy to calculate that Prob$\{$outcome $=$ LO$\} = 0.8$.[2] If we observe LO, then we would choose BAU, whereas if we observe HI, then we would still choose tempered abatement, as shown in Fig. 2.2. The expected value of this observation is found by connecting the best choices for each possible outcome by the thin dotted line in Fig. 2.2. The value of this information in this problem is the difference between the thin dotted line and the value of the best option at Prob$\{cs = 1.5\} = 0.66$. This value is rather small because a HI value doesn't change our initial choice.

---

[2] The conditional probability Prob$(cs = 1.5 \mid$ LO$) = 0.75$, and similarly Prob$(cs = 1.5 \mid$ Hi$) = 0.3$. Solve Prob$(cs = 1.5) = 0.66 =$ Prob$(cs = 1.5 \mid$ LO$)$P(LO) + Prob$(cs = 1.5 \mid$ HI$)(1 - $P(LO)$)$ to find Prob(LO) $= 0.8$.
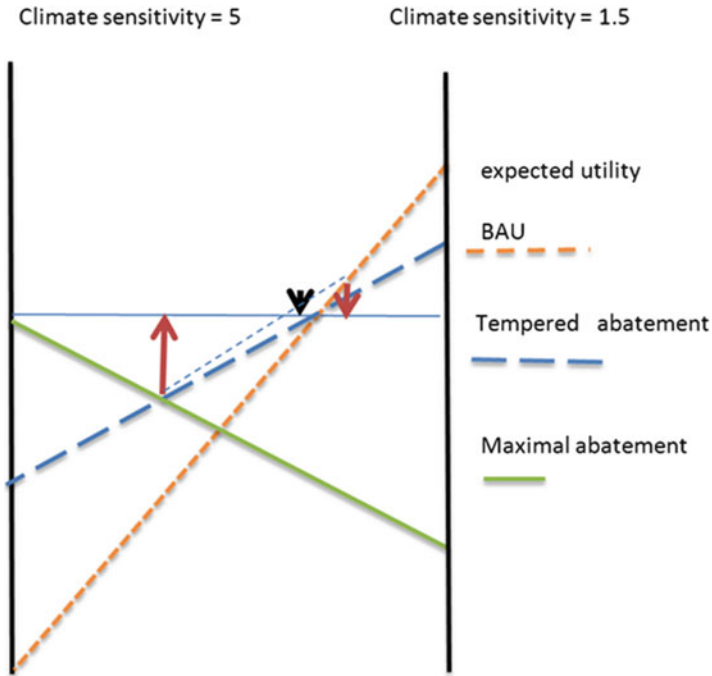
Fig. 2.2 Value of information simple climate policy choice problem with value of imperfect information

Our simple example demonstrates that for the value of information to be important, all of the following must obtain:

1. The set of available options is strongly concave in the sense that it consists of options that are very good in some states of the world and very bad in others, and options that are mediocre in all states of the world.
2. Our belief point leads us to choose an option that is much worse than what we would choose with perfect information.
3. There are observations whose possible outcomes would strongly influence our belief point.

2.4 A Risk Management Approach to the Value of Information

A risk management approach to climate change should translate through to value-of-information calculations. A risk management perspective suggests that value-of-information calculations should allow for examination of how improved information alters our estimate of the odds of meeting a given policy target. Furthermore, they should incorporate rigorous expert judgment for determining

beliefs and quantifying uncertainties and should highlight what scientific information is most valuable for a policymaker taking a risk management approach, as opposed to an efficiency approach, to the climate problem.

We argue that Bayesian belief nets are useful tools that meet all three criteria. A BBN is a graphical model representing variables and their conditional probabilities. It allows for quantification of uncertainty in complex models of multiple variables. A simple example based on the IAM of William Nordhaus, DICE, is used, with distributions on three uncertain parameters. We model temperature-induced damages, $\Omega(t)$, at time t as a function of global mean surface Temperature, $T(t)$, with uncertain parameter dx:

$$\Omega(t) = 1/[1 + \psi T(t)^{dx}]. \tag{2.1}$$

Temperature is a function of greenhouse gases (GHGs) and the uncertain climate sensitivity parameter, cs:

$$T(t) = cs \times \ln(GHG(t)/280)/\ln(2). \tag{2.2}$$

$\Omega(t)$ is a value between zero and one that scales down total output, Q, which is a function of abatement Λ, total factor productivity A (this is a parameter, evolving over time to capture technological change), capital stock K, and labor N, with uncertain Cobb-Douglas parameter gx:

$$Q(t) = \Omega(t)[1 - \Lambda(t)]A(t)K(t)^{gx}N(t)^{1-gx}. \tag{2.3}$$

Different policies are characterized by their GHG emissions: policy 1 involves the lowest emissions at highest abatement cost; policy 10 involves the highest emissions at lowest abatement cost. Greater abatement leads to reduced output.

This model as a BBN is shown in Fig. 2.3. The top three nodes represent uncertain variables in the model: the climate sensitivity cs, the exponent in our damage function dx, and the exponent in a Cobb-Douglas production function gx. We have assigned distributions to each of these variables.[3] One tenet of risk management is that these distributions should be assigned not in an ad hoc fashion by modelers (as we do here simply for purposes of illustration) but in a process of structured expert judgment. This involves transparently choosing a range of experts on the topic, familiarizing them with the study, allowing them to consider the problem and prepare a response, conducting a face-to-face interview, querying experts about measurable variables, querying experts about calibration variables,

---

[3] This model is meant simply to demonstrate the approach, and in that sense, the distributions chosen are somewhat arbitrary. We model the climate sensitivity as Beta distributed on [1, 15] with parameters (2, 24). The damage and production function exponent are both modeled as uniform variables, the first over [1, 3], and the second over [0.2, 0.4]. Part of the appeal of the BBN approach is that these distributions can be altered and the effects on output examined explicitly.
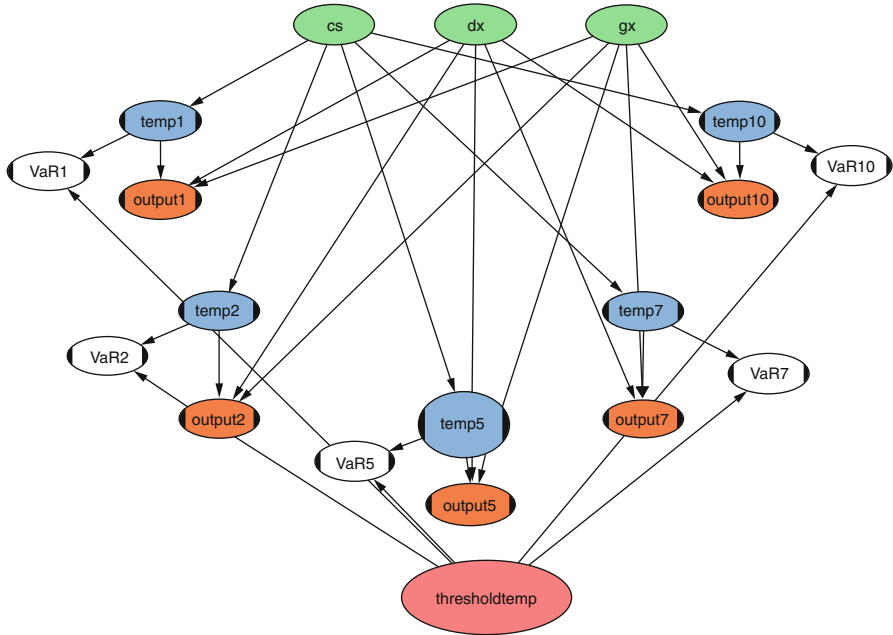
Fig. 2.3 Bayesian belief net for example climate model

and measuring performance on statistical accuracy and informativeness to aggre-gate judgments (Cooke and Kelly 2010). This process of expert judgment will allow for the best assessment of the uncertainties in the model.

The nodes labeled "output" in Fig. 2.3 represent output over the next 100 years under the five abatement policy options, which are shown in the Temp5 nodes. The arrows connecting the nodes represent defined relationships between those two variables. The thresholdtemp node at the base of the model allows for stipulating a threshold maximum temperature, and the model can then calculate the probability of exceeding this threshold for each policy option. When run, these probabilities would be shown in the White nodes, or the VaR nodes (see Fig. 2.4).

Figure 2.4 shows the result of running the model with temperature threshold set at 3 °C. The expected value of each variable ± the standard deviation is shown at the bottom of each node. The expected value of the VaR nodes is the probability of not exceeding the stipulated threshold. We see that the expected output and expected temperature increase as we move from option 1 to option 10, whereas the probability of staying below the stipulated temperature threshold drops. This reflects the fact, long obvious to investors, that increasing expected gain is coupled with greater risk. The first policy achieves our target 100% of the time, and the second policy achieves it 95.5%. By policy option 3, however, the target is met only about 39% of the time. If we defined 3° as our "collapse" point with a threshold of 5%, then only the first two policies would be deemed viable. We can see that output is higher under the second policy, as would be expected, since there are greater emissions.
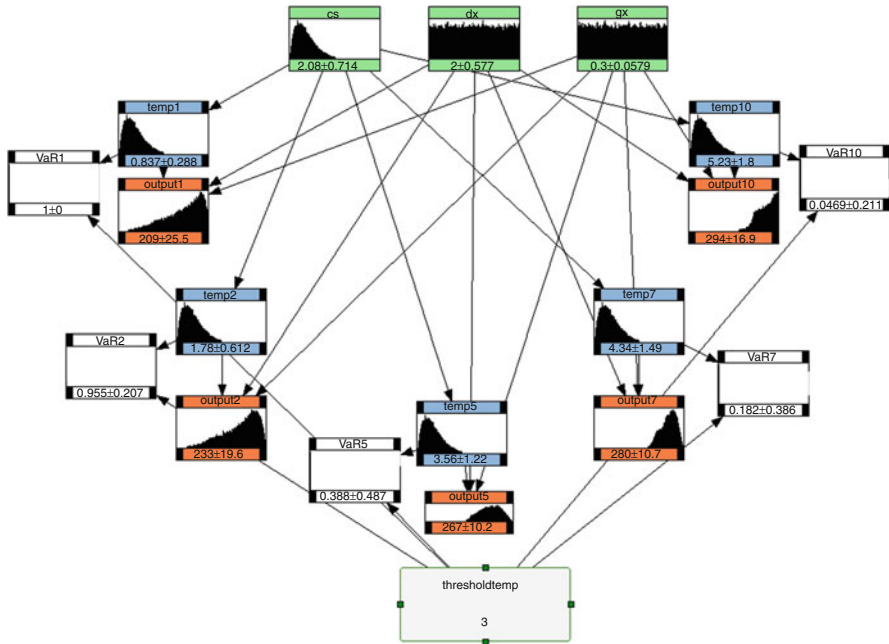
Fig. 2.4 Stipulating temperature increases to not exceed 3°

The BBN thus formalizes our uncertainty over particular parameters and, similar to the simple value-of-information model in the preceding section, will allow us to estimate the value of improved information on any of the uncertain variables. For instance, we can compare the distribution of output under the various policy options when climate sensitivity is modeled as an uncertain random variable and then compare this with the case when it is known with certainty or when its distribution narrows from improved information.

To illustrate, suppose we perform an imperfect observation on climate sensitivity, leading to the distribution shown in Fig. 2.5, with mean lowered from 2.08 to 1.2, and with narrower uncertainty. Now option 3 meets the risk management requirement of holding temperature below 3 °C with probability at least 0.95. The expected output of option 3 is 270. Without performing this observation, our best option meeting the risk management requirement was option 2 with expectation 233. The expected outputs in Fig. 2.5 are a bit higher than in Fig. 2.4, since the lower climate sensitivity leads to reduced damages for all options.

Once defined, the BBN can be sampled. Examining 1,000 such samples, displayed as a cobweb plot in Fig. 2.6, shows, as just one example, the relationship among climate sensitivity, temperature, and output, under the five policy options. The cobweb plot shows clearly that lower climate sensitivity values are associated with lower temperatures and higher output. Although this particular finding is obvious, it demonstrates the way in which a BBN can be used to explore the links among multiple variables. If a measurement could be taken to narrow the possible range for climate
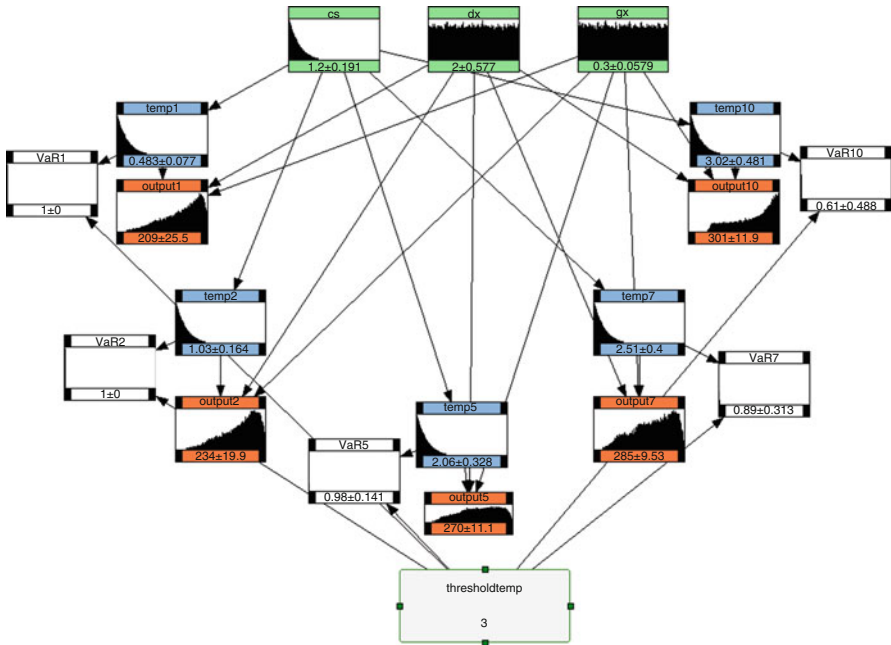
Fig. 2.5 Observation on climate sensitivity leading to shifted distribution

sensitivity, a value of that measurement can be determined by resampling our BBN with the narrowed range comparing output under the range of policies, as well as the probability that various policy options meet our threshold probability.

Finally, the BBN can help risk managers determine what type of information would be most useful and thus where best to direct scarce research dollars, or whether to invest in a particular research project. This can be done by comparing improved information on a variety of multiple uncertain variables. In this simple model, we only have three, but more complicated climate models would include a broader range of the uncertain variables. Our model would then let us uncover which types of information may be useful and which will not be. For instance, if catastrophic tipping points in the climate system are irreversible, detection is impossible, or detection would be too late for society to take action, Weitzman (2007) notes that the option value of waiting for more information would be zero. Thus, knowing when not to wait for more information and when not to invest in learning is just as important as knowing when to do so.

Note that to do this type of analysis effectively, we must clearly determine which uncertain variables are those over which we can undertake measurements to improve our knowledge and those where the uncertainty arises from other sources, such as differing value judgments. For instance, although there is uncertainty over the proper discount rate, this is at base a disagreement of values or opinion and cannot be resolved through better information.
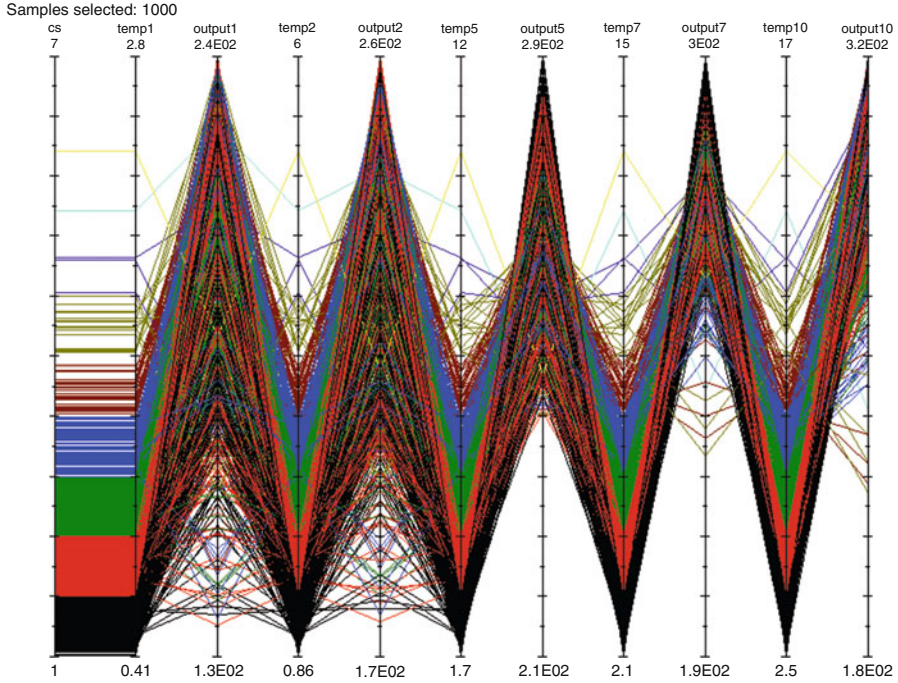
Fig. 2.6 Cobweb plot with 1,000 samples

The simple BBN used here represents just one climate model. It thus makes certain assumptions about the functional form of relationships among variables. From IAM modeling, however, we know that varying these assumptions can produce dramatically different outcomes. Fankhauser and Tol (2005), for example, observe that damages can affect capital depreciation, the utility function, the production function, and population growth. Which is chosen can create profound differences in predicted welfare for various policy choices. These differences can be addressed in one BBN by including different climate damage models.

2.5 Conclusion

Among the most challenging aspects of addressing climate change are the uncertainties and the possibility of truly catastrophic damages should we fail to abate sufficiently. Rather than neglecting these features of the problem, we suggest that a risk management policy approach be pursued, which would aim to keep the probability of reaching catastrophic damages below some tolerable threshold. Within this framework, improved information on some aspects of the climate problem will be more useful than other aspects. A simple model of the value of

information suggests when improved information will be helpful—namely, when we have policy options that produce very different outcomes in different states of the world, when our current beliefs lead us to choose a policy we would not choose if we had better information, and when it is possible to learn information that would alter our beliefs substantially. Although these heuristics are useful, more sophisticated analyses of the value of improved climate information based on detailed climate models would help policymakers make improved decisions about where to invest in information, how much to invest, and when more research is even worthwhile.

Such calculations can be performed using BBNs. Translating climate models into this framework creates a visually intuitive model in which it is easy to stipulate risk management thresholds and observe the consequences of improved learning within such targets. We have presented a simple illustration here, but a true analysis would, of course, require a much more detailed model. It would also require the use of expert judgment to adequately characterize the uncertainties, as well as discussion with scientists to discover what uncertainties could be reduced through various investments in research.

2.   Commentary: Valuing Information, Ascertaining Risk, and Setting the Target

Timothy J. Brennan[4]

In their contribution to this volume, Kousky and Cooke (this volume) offer a method, Bayesian belief nets (BBNs), as a way to acquire the relevant stochastic information under a risk management approach to climate policy. Under such an approach, the goal is to come up with a policy to meet a given target—for example, that the probability of a given level of warming is less than a tolerable maximum. KC present this in a unified framework, but I want to suggest that the three aspects of the discussion—valuing information, ascertaining risk, and setting the target— are separable issues, in the sense that the merits of the approaches to any of these aspects can be assessed independently of how we regard the others. To support this, I will look at these three aspects in turn, and then conclude with some observations on the use of aggregate expected utility in setting climate policy, whether under a cost-benefit or risk management framework.

2.C.1.   The Value of Information: Some General Issues

That information is valuable is obvious; more interesting is why that requires investigation in a way that the value of hamburgers does not. Part of the problem is that markets may not work as well for information as they do for hamburgers. Three aspects of information bring this out. A first is the circular regress in the purchase of information. Applying to information the general principle that one needs to know about a product before one can formulate a willingness to pay for it implies that before one would buy information, one would need to be informed about what the information is (Arrow 1959, 10). Were that true, why would one need it?[5]

---

[4] I thank Molly Macauley for a long series of extensive discussions on the value of information; she should be held blameless for mistakes I persist in making. I am also grateful to Roger Cooke and Carolyn Kousky for discussions of issues raised in their chapter; they too are without responsibility for my errors.

[5] A related issue is the distinction between data, the raw material an instrument gathers, and information, after the data have been turned into useful knowledge through some theoretical interpretation in the chain between initial the gathering of those data and their potential use.

T.J. Brennan (✉)
Public Policy and Economics, University of Maryland–Baltimore County, Baltimore, MD, USA

Resources for the Future, Washington, DC, USA
e-mail: Brennan@rff.org

Two other aspects are intriguing because they are somewhat contradictory. One is that information, apart from any physical medium in which it may be embodied (book, DVD) or communicated (theater, computer), meets the economist's definition of a "public good." In that definition, the consumption or possession of information is "nonrivalrous," meaning that one person's having it does not disable others from having or it as well. Consequently, information should be available to all who place a positive value on it, which generally requires that the price be zero. This is the somewhat technical argument behind the aphorism "information wants to be free" (Brand 1985, 49).

However, if information is free, one has the problem of covering the cost of discovering and providing it. Until recently, information providers recovered these costs through a combination of embedding the information in tangible goods that were relatively costly to duplicate and intellectual property protections against unauthorized copying and resale. The ability to convert information into digital formats detached from tangible goods and easily copied on computers and transmitted through broadband networks has blown a substantial hole in these protections. The good news is that a lot of information is free, but the bad news is that the business models used to support information provision have become unsustainable.

In contrast, although information meets the economic definition of a public good, the value of information often depends on exclusivity—that it is not made available as a public good in practice, even if it could be in principle.[6] One obvious context is business strategy, where the value of information depends on the competitive advantage it conveys over relatively ignorant rivals. Information can be valuable in interactions with buyers or sellers, although in those cases one runs the risk that asymmetry in information between buyers and sellers could cause adverse selection, leading to a collapse of markets to the detriment of all. In such cases, the value of information can be negative: market participants as a whole would be better off without it.[7] But business is not the only context where information is valuable to the extent it is private. Academic researchers strive to limit access to information to protect the priority of discovery on which reputations depend. The second person to write "$E = mc^2$" on a piece of paper is likely unknown, probably forever.

---

[6] The full quote from Brand (1985, 49) is instructive: "On the one hand, information wants to be expensive, because it is so valuable. The right information in the right place just changes your life. On the other hand, information wants to be free, because the cost of getting it out is getting lower and lower all the time. So you have these two fighting against each other." A quarter of a century later, this is still true.

[7] A virtue of having health insurance purchased through businesses or the government rather than individually is that it creates an information firewall that keeps individual knowledge of health status from creating an adverse selection problem in health insurance markets, where relatively healthy people choose not to purchase insurance, increasing the expected costs of ensuring those who purchase insurance, perhaps to the point where the insurance market disappears altogether. When employees can choose their health care provider, that firewall comes down, and the costs of insurance to all can increase (Cutler and Reber 1998).

2.C.2.   *Modeling the Value of Information*

Fortunately and understandably, KC focus not on the strategic value of climate information, but on its public good aspect. Moreover, they get around the paradox by invoking the assumption that the procurer of information knows that the information will resolve a question or reduce uncertainty regarding it, but does not know how that uncertainty will be resolved. They present this graphically, but it is useful to look at the value of information algebraically to see that their formulation applies not just to expected utility maximization but to risk management as well.[8] It also allows us to see that BBNs used to refine uncertainties could in principle be equally useful in either setting.

We face uncertainty about a parameter w reflecting sensitivity of climate to emissions. KC have w equal either to 1.5 or 5, but for notational convenience, we allow w to vary over a connected range W. Before gathering new information I, the likelihood that w takes a particular value is $f(w)$, which as a probability distribution satisfies

$$\int_W f(w)dw = 1.$$

Absent I, a policymaker charged with maximizing expected utility would have to choose a policy x^* to maximize

$$\int_W U(x, w)f(w)dw,$$

where $U(x, w)$ is the utility from policy x if w is realized.

Following KC, assume first that I allows the policymaker to know the value of w before choosing x. If so, she would choose $x(w)$ to maximize $U(x, w)$ given w. Thus, prior to acquiring I, the expected utility the policymaker would achieve with I is

$$\int_W U(x(w), w)f(w)dw.$$

The value of information V at the time one decides to acquire it is the difference between the expected utility with the information and the expected utility without it:[9]

$$V = \int_W [U(x(w), w) - U(x^*, w)]f(w)dw.$$

---

[8] None of the formal constructions in this section are novel.

[9] Finkel (this volume) points out that one can view this as the expected avoided cost of error or "regret" from choosing x^* instead of $x(w)$.

Figure 2.1 in KC illustrates this formulation in the case where w can take one of only two values.

If V exceeds the expected cost of gathering the information—for example, launching an earth observation satellite, carrying out malaria detection tests, or advertising product characteristics—the investment in information gathering is worth the cost; if V is smaller than that expected cost, the information is not worth obtaining. V will tend to be larger as the difference between the utility of choice knowing the information, $U(x(w), w)$ exceeds $U(x^*, w)$, over ranges of w that are relatively likely—that is, where $f(w)$ is relatively large. When $U(x(w), w)$ is not much different from $U(x^*, w)$, either because decisionmakers cannot choose $x(w)$ to be much different from x^* or because the choice does not affect U that much, or when the set of w *values* where the difference matters has a low *ex ante* likelihood of occurring, where $f(w)$ is small, the value of information will not be great and investments to procure it are less likely to be justified on cost-benefit grounds.

It is important to recognize that the value of information here is based on $f(w)$, the likelihood the information will be valuable *before the information is gathered*. Consequently, a study showing that information happened to be valuable, given a particular realization of w after the information was gathered, does not prove that the information was worth gathering in the first place. If the *ex ante* chance of observing that value of w, $f(w)$, was particularly small, the information investment may not have been worth making, even if it happened to turn out to be valuable. Similarly, an information investment could have been worth the cost even if the particular piece of information revealed turned out to have little effect on the decisionmaker's choice or utility, if outcomes that could have been important were sufficiently likely at the time the investment in information gathering was made.[10]

As KC point out in their second example, information may not enable a decisionmaker to identify w; it may give her only better information regarding the underlying distribution of w. To describe this, let the preinformation $f(w)$ be given as the weighted average of a set of probability distributions $p(w, z)$, where the preinformation probability of observing distribution $p(w, \mathrm{z})$ is $g(z)$ taken over a domain Z of distributions.

$$f(w) = \int_Z p(w,z)g(z)dz,$$

---

[10] This distinction has implications for whether case studies on the value of particular pieces of information really can tell us very much about whether it was valuable to invest in the ability to acquire that information. Such case studies tend to be valuable to the degree that the benefit of the information, $U(x(w), w) - U(x^*, w)$, is relatively constant over the set of w that one is likely to observe. For example, if (a) one knows that a disease might break out in one out of N areas but one does not know which area will be the one, (b) stopping the disease would have the same benefit regardless of where it broke out, and (c) preventive measures would not be effective absent knowing where the breakout is, then a study showing the benefits of detection will be informative regarding the value of the determining which area will have the outbreak. Whether assumptions along these lines are generally valid could be a useful question to study.

where

$$\int_Z g(z)dz = 1.$$

and for any z,

$$\int_W p(w,z)dw = 1.^{11}$$

Prior to learning z, the policymaker would choose x^* that maximizes expected utility

$$\int_W U(x,w) \int_Z p(w,z)g(z)dzdw.$$

After learning, the policymaker can choose $x(z)$ to maximize expected utility with the distribution $p(z, w)$ that describes the postinformation uncertainty about w:

$$\int_W U(x(z),w)p(w,z)dw.$$

The value of information is, as above, the expected improvement in being able to choose after learning, in this case the distribution $g(z)$ rather than w directly:

$$V = \int_Z \int_W [U(x(z),w) - U(x^*,w)]p(w,z)g(z)dzdw.$$

KC's Fig. 2.2 illustrates this where there are just two potential distributions.
    So far we have two uncertainties, the distribution of the climate variable w given the parameter z, and the distribution of the parameter z. Of course, there are other uncertainties as well, associated with how x will affect utility given w, such as how

---

[11] It may be that $p(w, z)$, the probability of w conditional on z, is learned through a process of Bayesian updating based on $g(z)$, prior beliefs regarding z, and $h(z, w)$, the likelihood that evidence z would be observed were a particular climate observation w valid:

$$p(w,z) = \frac{h(z,w)g(z)}{\int_Z h(z,w)g(z)dz}.$$

well the policy will work and what it will cost. One uncertainty is technological change that may reduce the cost of substituting away from fossil fuel use; another surrounds the ecological costs associated with extreme-case geoengineering. Let θ be a potentially multidimensional parameter reflecting these uncertainties over a range Θ distributed by $h(\theta, w)$, so that the "utility" $U(x, w)$ for any policy under any realization of climate sensitivity is really an expected utility

$$U(x, w) = \int_{\Theta} U(x, w, \theta) h(\theta, w) d\theta$$

where, for any w,

$$\int_{\Theta} h(\theta, w) d\theta = 1$$

Substituting this into the expression for the value of information in the general case where information narrows but does not eliminate uncertainty regarding the climate parameter gives

$$V = \int_{Z} \int_{W} \int_{\Theta} [U(x(z), w, \theta) - U(x^*, w, \theta)] p(w, z) g(z) h(\theta, w) d\theta dw dz.$$

We could increase the formal complexity by looking at the value of information about θ that changes the degree of uncertainty regarding the costs and benefits of climate policy. However, this formulation is sufficient to illustrate that the methods for constructing the value of information, the use of Bayesian belief nets, and the choice between expected utility and risk management frameworks, can be regarded as independent issues.

2.C.3.   Separating Value of Information, Risk Management, and BBNs

KC propose that climate policy be considered using a risk management framework, as distinguished, presumably, from an expected utility framework. KC (2012, page number tk) characterize the distinction in the following way:

> A risk management approach asks what policy should be, given the large range of possible outcomes from that choice. This is quite distinct from asking what the optimal policy is under different assumptions of our uncertain variables.

Perhaps I'm wearing blinders acquired from decades of being an economist, but the source of "quite distinct" is not obvious. Both seem to be doing the same thing.

The essential difference may be that under risk management, rather than maximizing expected utility, the task is "to keep the probability of facing catastrophic damages to some determined low level," or as said later, a "defined tolerable level." If so, the utility or benefits of the outcome itself are essentially out of the calculation. Maximizing expected utility thus becomes a matter of minimizing the cost of meeting the risk management probability target. In the formulation above, then, we can substitute for expected utility (which was net of cost) the negative of cost.

To put it another way,

$$U(x, w, \theta) = \overline{U} - C(x, w, \theta),$$

where $\overline{U}$ is the utility achieved at the "defined tolerable level" and $C(x, w, \theta)$ is the cost of implementing policy x with climate sensitivity y under circumstances θ. Substituting this for the value of information V above gives that value as the difference between the cost we have to pay for acting before getting the information and the cost we would expect to have to pay after we learned more about the distribution of climate conditions.

$$V = \int_Z \int_W \int_\Theta [C(x^*, w, \theta) - C(x(z), w, \theta)]p(w, z)g(z)h(\theta, w)d\theta dw dz$$

where x^* is the choice of policy that minimizes expected costs prior to acquiring more information about the distribution of climate sensitivity.

In my interpretation of KC, the value of information shows up in the context of a cost-effectiveness test rather than as expected utility maximization. However, as a matter of formalism and, importantly, the underlying uncertainties, the factors determining the value of information are essentially the same. It may be that some values of θ would affect utility but not cost, so the relevant distribution function $h(\theta, w)$ could be less difficult to ascertain. But we still need to know $h(\theta, w)$, $p(w, z)$, and $g(z)$ to determine whether an effort to acquire information to determine z is worth the cost and whether we are operating under an expected utility framework or a risk management framework.

Similar considerations show that the Bayesian belief nets KC advocate would be relevant in an expected utility context. The purpose of a BBN is essentially to use the knowledge of a set of experts to learn the probabilities relevant to a particular policy. For expected utility, one could interpret this as using those experts to learn the distribution of effects on utility instead of costs ($h(\theta, w)$). The expert information can also be used to determine which distribution of climate sensitivity $p(w, z)$ we have, essentially by finding z and eliminating the stochastic step added by $g(z)$. This suggests that the virtues of BBN in improving information regarding risk, and whether it is worth the cost to set one up, are at least qualitatively just as plausible in an expected utility framework as in a risk management framework. The expert survey methods KC propose are equally applicable in both circumstances, so their value is largely separate from whether one maximizes utility or minimizes the cost of reaching a maximum tolerable probability of a climate catastrophe.

2.C.4.   *Critiques Nonetheless*

Establishing that methods for calculating the value of information and Bayesian belief nets could apply equally to expected utility and risk management doesn't mean that risk management and expected utility are equally (in)valid ways of determining policy responses. Neither does applicability to expected utility theory insulate BBNs from critiques of risk management. The point of showing that the concepts are largely independent in principle allows their merits to be assessed largely independently. Some brief observations on risk management, BBNs, and their interrelationship follow. I conclude with a critique of economic approaches to utility maximization in the climate context.

2.C.4.1.   Risk Management

One can think of three justifications for using risk management over an expected utility approach. The first two are conventional arguments that justify cost-effectiveness over wealth maximization approaches generally. First, if the benefits of a policy are too difficult to quantify, one might simply evaluate policies in terms of how well they achieve a predetermined policy target. Quantification may be inherently difficult because the data on valuation are highly noisy. For the instructive case of the value of reducing mortality risk, the "statistical value of life," the underlying data on willingness to pay for incremental safety benefits are notoriously variable—but at least there are some markets or behavior trails from which a willingness to pay may be inferred. Since the effects of climate change are global, nonexcludable (and thus outside markets), and future, the present willingness to pay to mitigate it may be impossible to measure with any real confidence. One might be better off making a considered judgment regarding acceptable risk.

The second conventional argument for taking risk management rather than expected utility approach is that ethical considerations as well as economic factors determine the appropriate target, the instant example being the maximum tolerable probability of a sufficiently large-scale climate effect. I conclude below with some observations on the limits of the economic approach and, derivatively, the unavoidability of ethical considerations in the specific context of climate policy. In general, however, when lives or major changes to the social or physical environment hang in the balance, some may argue that policy responses ought not be determined solely by how much people might be willing to pay for them, even if that willingness could be measured accurately. One could set the level of tolerable risk using qualitative assessment, normative judgment, and communal deliberation (Ackerman and Heinzerling 2002) and then manage that risk by gathering information to seek out the most cost-effective programs.

A more recent argument for risk management and against expected utility, cited by KC, comes from Weitzman (2007). As I understand it, Weitzman's argument rests on two premises. The first is that the distributions for climate events that one could statistically infer from the data, the $f(w)$ in the models above, have fatter tails— more weight toward high climate sensitivity—than the underlying distributions

might entail. He contrasts the "t-distribution" from the normal distribution in that regard. The second premise is that the appropriate form for modeling the utility of wealth, where wealth is affected by climate, entails assuming constant relative risk aversion (CRRA). These two premises together imply that the expected (dis)utility from climate change is $-\infty$ and that any finite effort to alleviate it is justified.

I lack the expertise to address Weitzman's statistical premise, but the CRRA assumption is unsupported by theory and contradicted by experience. With regard to theory, CRRA is based on a quadratic Taylor series approximation to a utility function to model the willingness to pay to avoid variance in wealth. As such, it is constant only within a small distance around a target wealth level. Nothing suggests that one could extrapolate that approximation far outside such small variations, particularly to catastrophes. Were this so, individuals would regard the loss of life as sufficiently harmful to warrant arbitrarily large expenditures to limit mortality risk. That is, the observed value of a statistical life would be infinite— a prediction violated by almost everyone's behavior almost every day. KC's analysis does not rely on going this far, but it does undercut using Weitzman's argument to justify a risk management rather than expected utility approach.

I was surprised to see KC tout the advantages of "value-at-risk" (VaR) models for risk management. VaR has been taking a beating in the press, where its widespread adoption in assessing derivative portfolio risk has been blamed for the financial meltdown of 2008:

> Given the calamity that has since occurred, there has been a great deal of talk, even in quant circles, that this widespread institutional reliance on VaR was a terrible mistake. At the very least, the risks that VaR measured did not include the biggest risk of all: the possibility of a financial meltdown. "Risk modeling didn't help as much as it should have," says Aaron Brown, a former risk manager at Morgan Stanley who now works at AQR, a big quant-oriented hedge fund. A risk consultant named Marc Groz says, "VaR is a very limited tool." David Einhorn, who founded Greenlight Capital, a prominent hedge fund, wrote not long ago that VaR was "relatively useless as a risk-management tool and potentially catastrophic when its use creates a false sense of security among senior managers and watchdogs. This is like an air bag that works all the time, except when you have a car accident." Nassim Nicholas Taleb, the best-selling author of "The Black Swan," has crusaded against VaR for more than a decade. He calls it, flatly, "a fraud." (Nocera 2009)
>
> VaR, by the way, is a horrible way to measure risk, as has been said again and again by economists, because it calculates the risk for only 99% of the time. As [Simon Johnson, a professor at Sloan School of Management at MIT] says, "VaR misses everything that matters when it matters." Indeed, the VaR metrics obviously missed what led to what now has been dubbed the Great Recession. (Sorkin 2009)

I am not an expert in assessing financial risk, but I find one potential flaw of VaR, a shortcoming it shares with risk management generally: once the level of acceptable risk is determined, factors that might make the costs of unlikely events outside the range of acceptability become irrelevant to decisions, rather than ratcheting down acceptable probabilities as would happen with an expected utility approach. However, a reasonable response may be that the financial meltdown was the fault of not VaR but its application. To paraphrase the National Rifle Association, one could say, "Models don't kill economies, bankers kill economies."

2.C.4.2.   Bayesian Belief Nets

My initial reaction to KC's BBN proposal was skepticism. My epistemological predisposition is that knowledge is something an individual acquires by examining the evidence and analyzing relevant theories, not by what appears from taking a poll. Of course, as KC go to some length to point out, BBN is more than mere poll taking. Underlying it is a process, not described in detail in their chapter, for treating expert opinions themselves as data amenable to Bayesian updating and maximum likelihood testing to reduce the uncertainty associated with a particular potential phenomenon, such as climate sensitivity or the costs of various policies.

Along with recognizing the statistical aspect of BBNs is the point that knowledge is routinely combined through institutional mechanisms to arrive at better estimates. Markets, certainly since von Hayek (1945), can be seen as information media in which disaggregated estimates of costs and value are combined to provide prices, the best estimates possible of marginal benefits and marginal costs. Asset markets—stocks and bonds, commodity futures and options, derivatives—similarly combine information regarding expectations of those benefits and costs to estimate their present values. The efficient market hypothesis is that those markets cannot be systematically beaten without unique information (Malkiel 2003), a view currently under fire since the 2008 financial market meltdown (Wighton 2009). Variations on this theme involve prediction markets (Iowa Electronic Markets, Intrade) in which assets are created with payoffs based on the outcome of elections, legislation, wars, or other events. Such markets have been proposed, albeit controversially, to predict the likelihood of catastrophes or terrorist attacks (Hanson 2007).

We, or at least I, should also keep in mind a couple of things about the individual character of knowledge. Academically, material does not become accepted, nor is a lot of research funded, without peer review, itself a kind of collective expert assessment. More broadly, what all of us believe we know is far, far greater than what we have individually found out. From childhood to the present, we depend on teachers, books, colleagues, journals, libraries, and maybe even Wikipedia, to tell us what we believe we "know." Each of those sources is fundamentally a sort of "belief net." That they lack the statistical foundation of BBNs is in BBNs' favor; the crucial point is that resistance to BBNs on the basis of their being collective rather than individual may be misplaced.

It does remain the case that a BBN need not be restricted to risk management settings. As noted above, a BBN can reduce uncertainties to improve policy choices under expected utility maximization as well. We do need to be careful to keep in mind that the value of a BBN is something that needs to be known *ex ante*. KC's portrayal gives the impression that the value of a BBN is realized *ex post*. We need the assorted probability distributions described above to determine whether the costs of a BBN are worth undertaking. More on how one would make this *ex ante* assessment would be useful. One could imagine a staggered set of BBNs, where one undertakes a relatively low cost survey of a small number of experts to determine whether a full-blown BBN would be worth the costs.

2.C.4.3.   Putting Them Together

The theme here is that value-of-information calculations, risk management versus expected utility, and the role of BBNs can all be assessed independently. Although that seems largely true, there is one sense in which BBNs combined with risk management could be problematic. As KC say, the goal under risk management is to find the best way to deal with a "defined tolerable level" of risk. This leaves open the question of how "tolerable level" comes to be defined. Under risk management, this is specifically not a matter of ascertaining individual willingness to pay, whatever that may be, and defining tolerable level as the point at which the revealed willingness to pay for further reductions in probability no longer exceeds the cost of those reductions.

If not, then how do we define the tolerable level? If the preferences of the general public cannot be ascertained or relied on, then the question of who gets to choose becomes unavoidable. If so, the BBN framework raises the possibility that the experts may interpret questions to be not about their best guess of a relevant value or probability distribution, but about their interpretation of "tolerable level." This may happen unintentionally but could be problematic nonetheless.

For example, one might find a climate expert who is an avid bicyclist and vegan who finds air-conditioning oppressive. Another expert may love his Hummer and steaks and believe that air-conditioning is great, especially when one keeps the windows open so the house doesn't smell musty. Those experts' assessments regarding relevant probabilities of climate sensitivity and thus justification for incurring costs of policies may become difficult to disentangle from their own judgments about what costs are worth incurring in a society. Perhaps this is just a predictable observation from a paternalism-averse economist, but we probably need to be careful in framing BBN surveys to minimize the degree to which they become a forum for elites' desires trumping the preferences of the public. Even if one has few qualms about letting elites make policy decisions, the political legitimacy of BBNs as an advisory tool may depend on limiting them to "just the facts."

2.C.5.   Expected Utility, Risk Management, and Climate Policy

We can conclude with an observation (Brennan 2010) that could cut in favor of the risk management approach, and less obviously but possibly Bayesian belief nets as well. When economists employ utility maximization models for policy, they—we—typically interpret "utility" as in consumer surplus terms—that is, aggregate willingness to pay. The standard normative critique of this approach is that in aggregating surplus across everyone in an economy, it treats each dollar of net benefit as equal, whether the recipient is homeless or a billionaire. The justification for ignoring the normative consequences of the distribution of net benefits is that the winners could in principle compensate the losers, creating a situation where

everyone is better off—a generally uncontroversial improvement. Someone other than the policy evaluator can decide explicitly or implicitly whether some other distribution of benefits is preferable.

However, there may be no compensation-in-principle in the case of climate change. If the benefits of climate policy are realized only by generations far in the future, this compensation would require that those in the future pay us for the sacrifices we make on their behalf. Since future output cannot be put into a time machine for delivery to us, such compensation may be impossible. If that is the case, the standard economic approach does not suffice; explicit moral assessment of our obligations regarding the welfare of future generations and environmental protections becomes paramount.

The inevitability of an explicitly normative dimension beyond efficiency to climate mitigation policy suggests we ought to focus not on maximizing expected utility but on determining an ethically tolerable level of risk of severe climate change to be met at least cost—the risk management approach that KC advocate. This would entail gathering information to reduce the relevant uncertainties associated with the costs of various means of meeting that objective. KC provide a framework to accomplish this, making their contribution to the climate policy community important.

References

Ackerman, F., & Heinzerling, L. (2002). Pricing the priceless: Cost-benefit analysis of environmental protection. *University of Pennsylvania Law Review, 150*, 1553–1584.

Arrow, K. (1959). *Economic welfare and the allocation of resources for innovation* (P-1856-RC). Santa Monica: RAND Corporation.

Brand, S. (1985, May). Keep designing: Discussions from the Hackers' conference, November 1984. *Whole Earth Review, 46*, 44–55.

Brennan, T. (2010). The challenges of climate policy. *The Australian Economic Review, 43*, 225–239.

Cooke, R. M., & Kelly, G. N. (2010). *Climate change uncertainty quantification: Lessons learned from the joint EU-USNRC project on uncertainty analysis of probabilistic accident consequence code* (Resources for the Future Discussion Paper 10-29). Washington, DC: Resources for the Future.

Cutler, D., & Reber, S. (1998). Paying for health insurance: The trade-off between competition and adverse selection. *Quarterly Journal of Economics, 113*, 433–466.

Fankhauser, S., & Tol, R. S. J. (2005). On climate change and economic growth. *Resource and Energy Economics, 27*(1), 1–17.

Hanson, R. (2007). The policy analysis market: A thwarted experiment in the use of prediction markets for public policy. *Innovations, 2*(Summer), 73–88.

IPCC. (2007). *Climate change 2007: Synthesis report*. Contribution of Working Groups I, II and III to the fourth assessment report of the Intergovernmental Panel on Climate Change. Geneva: Intergovernmental Panel on Climate Change.

Keller, K., Yohe, G., & Schlesinger, M. (2008). Managing the risks of climate thresholds: Uncertainties and information needs. *Climatic Change, 91*, 5–10.

Lempert, R. J., Schlesinger, M. E., Bankes, S. C., & Andronova, N. G. (2000). The impacts of climate variability on near-term policy choices and the value of information. *Climatic Change, 45*, 129–161.

Malkiel, B. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives, 17*, 59–82.

Manne, A., & Richels, R. (1995). The greenhouse debate: Economic efficiency, burden sharing and hedging strategies. *Energy Journal, 16*(4), 1.

Nocera, J. (2009, January 2). Risk mismanagement. *New York Times Magazine*. http://www.nytimes.com/2009/01/04/magazine/04risk-t.html. Accessed 1 July 2011.

Savage, L. J. (1954). *Foundations of statistics*. New York: Wiley.

Sorkin, A. R. (2009, September 16). Taking a chance on risk, again. *New York Times*. http://dealbook.nytimes.com/2009/09/17/taking-a-chance-on-risk-again/. Accessed 1 July 2011.

von Hayek, F. (1945). The use of knowledge in society. *American Economic Review, 35*, 519–530.

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

Wighton, D. (2009, January 29). Efficient market hypothesis is dead—for now: Business editor's Davos commentary. *The (London) Sunday Times*.