The Tale of the Fat Tail

We have tough choices ahead, but are the tools adequate? Typical cost-benefit analysis of global warming scenarios only concerns itself with the likeliest outcomes, but possible extreme events still have a good chance of occurring, swamping the calculation

Melinda Kimble and Letha Tawney

Climate change policy is a minefield of tough choices. Making good decisions often comes down to the central question of cost. How affordable are greenhouse gas reduction commitments? How do we weigh the possible costs of a proposed law or regulation against the avoided costs of inaction, what we usually term benefits?

The global shift away from fossil fuels is often framed around an analysis of what the economy can bear. Cost issues drove inclusion of the emissions trading concept in the Kyoto Protocol, because the market approach explicitly values achieving controls through the lowest-cost alternative. Cost has not always been our central concern in managing environmental problems, however. Environmental legislation in the United States has run the gamut from laws that require regulators to act only on an assessment of risk (the Clean Air Act requires air quality standards at a level sufficient to protect the public health with an "adequate margin of safety") through technology-based requirements (the Clean Water Act's main regulatory standard is the best available technology that is economically feasible) to laws to that explicitly require cost-benefit balancing (the Toxic Substances Control Act and Safe Drinking Water Act).

Melinda Kimble (below) is a Senior Vice President at the UN Foundation and a former acting Assistant Secretary of State for Oceans, Environment and Science. **Letha Tawney** is a Climate Change Policy Analyst at Emerald Arc Consulting. Both are based in Washington, D.C.

For economists, assessing what the economy can bear is a way of asking how much we value emissions cuts. Underlying the cost-benefit approach to evaluating emissions cuts is an assumption that we should not pay more for future benefits than they are actually worth to us. There's a great deal of discussion in academic and policy circles about how to fine tune these estimates. Thus, the literature is full of intense debate about how much future benefits should be "discounted" in order to compare them accurately to costs today and other fine points. Other disputed issues include how to manage the uncertainty inherent in the climate science, assumptions about low-carbon innovations and technology spread, and many other details. As a result, the calculations can arrive at very different estimates of the cost of emissions cuts and the benefits of avoided damage.

But what if these debates are a distraction and cost-benefit analyses are simply not reliable guides when applied to an unprecedented environmental challenge? At least one prominent economist has issued exactly this warning. Martin Weitzman of Harvard University questions whether standard economic models can provide sound advice on emission permit prices or the timeline of greenhouse gas controls. He says the models artificially ignore the risk of a cataclysmic and irreversible outcomes. Once that risk is included in the analysis, it swamps the carefully estimated costs and benefits and argues for an approach that starts instead by valuing the risk of catastrophe.

There appears to be consensus within the Obama administration that putting a price on greenhouse

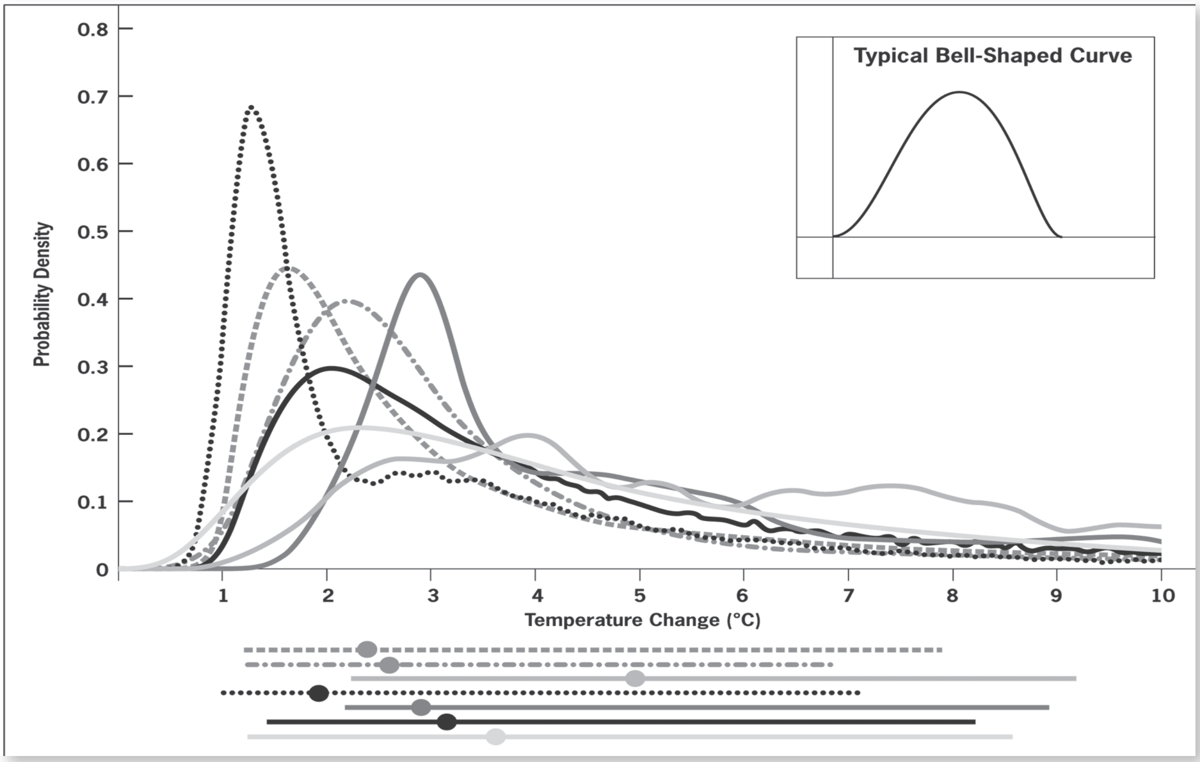

This chart is based on a similar one published by the Intergovernmental Panel on Climate Change in its latest assessment report, simplified to highlight its implications (Fig. 9.20, AR4 WG1). It shows the results of seven studies modeling the probability of various temperature shifts when the atmosphere contains 550 parts per million carbon equivalent. Below the curves are bars that give the probability range for each curve, lopping off the last 5 percent on each side as unlikely. Ordinary bell shaped curves (inset) fall off rapidly to zero probability on each side. The IPCC curves, however, extend well out to the right, where the most extreme temperature changes are, and their width makes them more probable. These "fat tails" mean that enormously consequential climate impacts are more likely than we assume in our usual decisionmaking.

gas pollution by means of a cap on carbon emissions and an emissions trading system is the cheapest and most politically feasible way to cut greenhouse gas emissions. But choosing a cap-and-trade system is only the first decision. Many details of the market design are hotly debated, including how to manage price volatility (using tools such as safety valves and banking and borrowing of emissions permits), auctioning versus giving away permits, and the length of time a permit would be valid. All of these debates flow from the basic assumption that it is possible to calculate a point of unreasonable or unnecessary stress to the economy.

This point of stress is derived from a systematic examination of how costs incurred today compare to benefits accrued in the future. Although policymakers may disagree about details, they all decide what the economy can bear today — the target permit price — based on this estimate. (A few economists have challenged the validity of using cost-benefit calculations for *any* environmental challenge, but that is not the subject of this examination.)

A very simplified version of how the cost-benefit analysis is typically applied to climate change can clarify. It starts from the assumption that there is no "safe" level of pollution, because climate change causes damage at every temperature change. Rather, there is an acceptable quantity of damage. The economist's approach to finding the right permit price calculates how much value society derives from cutting emissions. It values the emissions so that we know how much to pay for them — their economic utility. It assumes that society will not divert unlimited resources to halt climate change entirely. Therefore, some amount of damage is reasonable, affordable, even efficient compared to the cost of entirely stopping the pollution.

To establish the value of emissions cuts, standard economic models start with the costs we will all bear if we take no action. For example, the costs of failing to make emissions cuts today might include the need to build better flood control systems, improve disease surveillance, or rebuild a damaged community. Economists try to evaluate how much damage each change in temperature will cause by reviewing how the scientific community assesses impacts. They build a "damage curve" showing how, as temperature rises, the cost of the damage increases. Any of these damages that are avoided through emissions cuts are benefits of the regulation. The curve rises from left to right, as temperatures increase.

Next, the economic models estimate the expense of cutting each ton of emissions. This estimate creates an "abatement curve" that shows the costs of achieving particular pollution goals. This curve declines from left to right, as it is most expensive to eliminate pollution entirely (cut the very last ton and achieve the lowest concentration and lowest temperature change) and cheapest not to limit pollution at all.

The two curves have to be tied together using a common unit of measure, tons of greenhouse gas emissions, in order to balance the cost of cuts against the benefits they buy. To do this, the damage curve has to be converted from one that assesses damage at various temperature levels to one that shows damage per each additional ton of emissions. To make this conversion, temperature changes have to be connected to pollution concentrations, which in turn have to be connected to the emissions themselves. Economists look to climate scientists to help make this correlation. Connecting the curves is difficult because scientists are uncertain how much temperature change a particular concentration of pollution will create or how much pollution will lead to that concentration. The economists often simplify this uncertainty. They start by lopping off the low and high ends of the potential temperature change range, those with less than 5 percent chance of happening. After additional statistical work particular to each model, the correlation is complete, and the model can plot the damages expected for each additional ton of pollution.

Scientists are uncertain about how much change. Economists often simplify this uncertainty

In the cost-benefit analysis, economists find the optimal point where society spends no more on cutting the next ton of emissions than is gained in benefits. This is where the upward sloping damage curve crosses the downward sloping abatement curve on the newly combined analysis. Aggressive (and expensive) cuts in emissions produce more benefits by avoiding more damage. However, the assumption is that spending more on cutting emissions than we gain in benefits would be a waste of money. Spending less on emissions cuts and bearing more damage underestimates the cuts' actual value to society. The incremental emissions cuts should be worth the incremental damage they avoid in the future. This is "what the economy can bear."

Although this is the basic way of doing cost-benefit analysis, economists are not of one mind on various parts of the calculation, including the choice of discount rate (which determines the present value of future benefits), approaches to opportunities to change course, and assumptions about the wealth of future generations, among many other things. These decisions lead each analysis to a different estimate of damages. The models also vary in their estimates about the expense of emissions cuts, depending on assumptions about innovation, capital stock turnover, and other aspects of the shift to a low-carbon economy. A great deal of thought is put into how to estimate these curves properly, leading to a variety of conclusions. For example, *The Stern Review on the Economics of Climate Change* estimated that damages and risks of unchecked climate change cost 5 percent of annual global GDP currently and that costs could go as high as 20 percent. In comparison, in his 2008 book *A Matter of Balance,* Yale economist William Nordhaus estimates the damages of unchecked climate change may eventually reach 2.5 percent of global GDP by 2100. But the basic framework of analysis is the same.

Considering catastrophe

Harvard's Martin Weitzman has raised fundamental questions about this approach to cost-benefit analysis in the special case of climate change. Why is climate change different from most environmental issues? Weitzman's main concern is that the climate science points to a probability, very difficult to quantify but also too large to dismiss, of a very large change in temperature — an irreversible and very bad result at an unknowable threshold level of pollution. Standard analysis trims off the worst-case outcomes with less than a one in twenty chance of happening, as we have seen, but Weitzman tells us that seemingly remote possibility is exactly where we should be looking — because the costs are so high they overwhelm other elements of the cost-benefit analysis. He points attention to the potential for catastrophe and how that potential is calculated into a damage curve. Weitzman's argument undercuts any certainty in the standard calculations of what the economy can bear.

Weitzman's first paper on this issue, written in response to the *Stern Review,* pushed the economics community to "start posing and trying to answer tough questions about rare global-warming catastrophes." Without that challenge, Weitzman says, "We will not make real progress in dealing constructively with the nightmare scenarios."

Finding answers to these questions, Weitzman says, is not easy, because economics simply does not have "a commonly accepted usable economic framework for dealing with these kinds of *thick-tailed extreme disasters,* whose probability distributions are inherently difficult to estimate." (Emphasis added.)

What is a "thick tail?" Weitzman tells us in his paper to take a good look at the right side of a probability distribution of temperature rise published in the 2007 report of the Intergovernmental Panel on Climate Change. (See figure on page 25.) The IPCC graph illustrates the chance of each temperature change that 550 parts per million of carbon dioxide equivalent (CO_2e) in the atmosphere might cause when the climate system finally comes to equilibrium. The horizontal axis shows global average temperature change, ranging from no change to a very large change, such as +10°C (+18°F). The vertical axis is a surrogate for probability. The probability distribution climbs, peaks at the most likely temperature change and then falls to less likely temperatures. However, the curve does not fall directly to zero probability. Instead, there is a "thick tail," a space below the curve that represents the small chance that very large temperature changes might happen at 550 ppm CO_2e.

This graph is the sort of input economists use to correlate temperature and concentrations of pollution into their own models. As discussed, standard economic analysis to determine the damage caused by each concentration of pollution would at most examine only the portion of the curve directly over

Wake-up Call for Economics

Harvard economist Martin Weitzman published an important analysis last year in which he explained why conventional economic analyses of climate change are "arbitrarily inaccurate." Weitzman's bottom line: If you don't factor in plausible extreme-impact scenarios — and the vast majority of economic analyses don't — your analysis is worse than useless.

The extreme or fat tail of the damage function represents what Weitzman calls "rare climate disasters." My one disagreement with Weitzman is that the science says they aren't rare at all, they are near certain with business-as-usual emissions. For Weitzman, disaster is a temperature change of greater than 6°C (11°F) in a century, "a terra incognita biosphere" with "mass species extinctions, radical alterations of natural environments, and other extreme outdoor consequences."

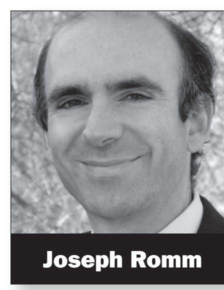
Joseph Romm

Weitzman says the Fourth Assessment of the Intergovernmental Panel on Climate Change gives the probability of such an "extreme" temperature change as 3 percent, and that "to ignore or suppress the significance of rare tail disasters is to ignore or suppress what economic theory is telling us loudly and clearly is potentially the most important part of the analysis."

What is especially important about Weitzman's analysis is that we know now there is far greater chance than 3 percent that total warming will exceed 6°C if we don't reverse emissions trends soon. Indeed, there is a 50 percent chance we will see warming above 5.5°C this century on our current emissions path, according to the United Kingdom's prestigious Hadley Center — and a 10 percent change we'll beat 7°C. MIT's Joint Program on the Science and Policy of Climate Change comes to a similar conclusion. A close reading of the 2007 IPCC report makes the same point.

The failure to explain that business-as-usual greenhouse gas emissions leads to 5°C or higher total warming and thus catastrophe — indeed, most likely an irreversible 1,000-year catastrophe — may be the greatest single messaging failure of the scientific community (and science media).

What exactly is the cost of sea levels in 2100 being 5 feet higher and rising thereafter 8 to 20 inches or more a decade until the planet is ice free in several centuries and sea levels are 250 feet higher? That is now business as usual. Same for much of the planet becoming a permanent Dust Bowl, most species going extinct, and large parts of the ocean becoming hot, acidic dead zones.

That is why essentially every cost-benefit analysis on climate in the literature is wrong and useless and hence very dangerous if taken serious by policymakers.

Weitzman concludes that conventional economic models are "especially and unusually misleading." He notes that "in rare situations like climate change" where extreme outcomes are a plausible outcome, "we may be deluding ourselves and others with misplaced concreteness."

Yes, that is what mainstream economists are doing by peddling their standard cost-benefit analyses to the public, the media, and policymakers — deluding themselves and others.

Joe Romm is a Senior Fellow at the Center for American Progress and editor of climateprogress.org.

the bars at the bottom of the graph, the 5–95 percent range of each probability distribution shown as a full curve above. Weitzman is instead interested in implications of the temperatures that, in the IPCC's graphic, have a 5 percent chance of happening. In the IPCC graphic, these are where the curves extend beyond the bars on either side. Weitzman used this graph and supporting studies to estimate the risk of catastrophic temperature change, changes greater than +6°C (+10.8°F), as at least 5 percent at 550 ppm CO_2e. He accurately described this scale of temperature change as a "terra incognita," beyond human experience.

Five percent sounds quite small. But translating this "fat tail" into a gamble, the odds of extreme catastrophe are 1 in 20. (For comparison, a "thin" tail might have odds of .1 percent or one in one thousand. Much of modern environmental law concerns itself with even smaller risks.) We might live within a short walk of a Superfund site whose hazardous wastes are calculated to produce one excess cancer in 100,000 residents. But none of us would board an airplane that had a 1 in 20 chance of crashing. We would never buckle our children into car seats with a 1 in 20 rate of failure. While these are small odds at first glance, the downside is disastrous, making the chance unacceptably large. The same applies to the safety of the biosphere. As a result, Weitzman argues these events should be explicitly accounted for in the cost-benefit analysis, not simply lopped off.

Significant emissions cuts will be necessary even to level off at the 550 ppm CO_2e example from the IPCC graph. The IPCC estimated that in 2005, CO_2e was about 375 ppm. The current trajectory of pollution growth means we will easily exceed 550 ppm CO_2e in the coming decades if we don't take rapid action. New research from MIT gives some idea of the risk of catastrophe on our current emissions pathway. The MIT Joint Program on the Science and Policy of Global Change uses 1990 temperature as their baseline in contrast with the IPCC, which uses pre-industrial temperature. However, technical adjustments to add the 20th century warming (about +.7°C or +1.26°F) allow a rough comparison of risk between business as usual and an ambitious policy. (Equilibrium climate sensitivity and temperature change by 2100 are not identical, but serve for comparison.) MIT reports 50–50 odds of +5.1°C (+9.2°F) or larger change compared to 1990 temperatures from business as usual emissions. This is roughly the equivalent of +5.8°C (+10.4°F) compared to a pre-industrial baseline and very close to Weitzman's approximate definition of catastrophic. That is, on the current pathway without dramatic policy changes, there is a 50–50 chance of "terra incognita" by 2100. By holding our emissions down to 550 ppm CO_2e we can lower that risk to about 5 percent, which still may be unacceptably large.

Weitzman accurately described this scale of change as a "terra incognita," beyond human experience

All of this discussion of temperature probability distributions seems perhaps unrelated to economic modeling, and this is exactly Weitzman's complaint. It is intuitive to use the "most likely" temperature to connect pollution concentration to temperature. For example, Nordhaus assumes a change of +3.1°C (+5.58°F) in 2100 relative to 1900 temperatures and a pollution concentration of 685 ppm CO_2 in his damage estimates. He sees the risk of catastrophic temperature change as too small to alter the model's outcome. Weitzman is less interested in the "likely" temperature and instead focuses the damage estimate for 550 ppm CO_2e on the 1 in 20 risk of catastrophe because the downside is so disastrous.

Weitzman's second paper on the topic, written in the fall of 2007 after long reflection on the climate science, ends with an even stronger conclusion than the first, one that should reinforce concerns about the limitations of traditional cost-benefit estimates for deciding the value of emissions cuts.

"The economics of fat-tailed catastrophes raises difficult conceptual issues which cause the analysis to appear less scientifically conclusive and look more contentiously subjective than what comes out of an empirical [cost-benefit analysis] of more usual thin-tailed situations. But if this is the way things are with fat tails, then this is the way things are, and it is an inconvenient truth to be lived with rather than a fact to be evaded just because it looks less scientifically objective in cost-benefit applications. . . . Perhaps in the end the climate-change economist can help most by *not* presenting a cost-benefit estimate for what is inherently a fat-tailed situation with potentially unlimited downside exposure as if it is accurate and objective and perhaps not even presenting the analysis as if it is an approximation to something that is accurate and objective…" (Emphasis in original.)

Risk takes center stage

This critique is far more than a dispute among economists. It goes to the heart of how policy is set — the target permit price and how we value emissions cuts. Weitzman is saying that a cost-benefit analysis that acknowledges the risk of catastrophe arrives at a more subjective, but realistic value for emissions cuts. In the article, Weitzman details several significant reasons for why risk overwhelms the standard models and produces markedly different results than a typical environmental cost-benefit analysis.

Risk takes center stage because there are no reliable ways to estimate the cost of that extremely large temperature change. Without experience to draw on, trustworthy estimates of the cost of such a disaster are impossible, though we can suppose it is very expensive. Weitzman argues that extrapolating from the damage caused by small temperature changes is purely an academic exercise. We know the cost of the crop losses during the Dust Bowl in the North American prairies during the 1930s, but we don't know the cost of simultaneous droughts across Africa, North America, and Australia. The inundation of New Orleans during Hurricane Katrina provides an imperfect analogue to the potential costs of rapid sea level rise throughout the United States, much less the possibility of millions of environmental refugees across the globe.

Additionally, temperature changes are likely to be virtually irreversible in any short time frame. The climate system has tremendous momentum and does not allow for a sudden change of course at some future time. Carbon dioxide and other greenhouse gasses continue to accrue damage for decades or centuries, so emissions today commit us to long-run future damage.

Finally, there is much debate about when tipping points might wrest control of the warming process from us entirely, rocketing the planet to much larger temperature changes than we've modeled to date. Despite intense research, science's ability to predict that moment remains clouded. James Hansen at NASA's Goddard Institute asserts that the current figure of 385 ppm of CO_2 already has too large a risk of irreversible tipping points based on his research into CO_2 in the paleo-climate. Others argue that 450 ppm CO_2 or 550 ppm CO_2 will not push us over the edge, but there is deep uncertainty that will not be easily dispelled. This uncertainty is troubling enough, but Weitzman points out the limited potential of further learning to inform our course of action in a timely way because of the strong momentum in the climate system. Getting the price wrong for most environmental challenges is reparable. If monitoring reveals more acid rain falling than optimal for example, the price on sulfur dioxide pollution can be adjusted up and further damage quickly ceases. The feedback loop between regulation, pollution, and damage is assumed to be relatively short, allowing tuning of the cost-benefit analysis as we learn. In climate change, the feedback loop is very long and seeing a catastrophic change before we've tripped it will continue to be difficult or impossible.

Weitzman's concerns drive to the crux of the issue. If a stalwart of environmental decisionmaking tools — cost-benefit analysis and all the traditional assumptions about damage that feed into it — is uncertain about how to calculate the value of emissions cuts, what can policymakers rely on? How can we acknowledge Weitzman's version of the "inconvenient truth," the disconnect between the risk of irreversible catastrophe detailed by the science and the simplified damage estimate assumed in the economic models?

We are not arguing that cost not be considered in formulating climate remedies. Society must make some assessment of how much of its limited resources to spend on emissions cuts. We are not even arguing to abandon the cost-benefit analysis of climate change mitigation. Rather, we are questioning how the estimates of acceptable cost, the value of emissions cuts today, best incorporate the unique nature of the climate problem.

We might value emissions cuts based on how we value lowering the risk of irreversible catastrophe

Weitzman suggests we might value emissions cuts not by the estimated future damages they'll avoid but based on how we value lowering the risk of irreversible catastrophe. A transparent and realistic discussion of the odds of disaster might lead to an emissions cap that made a more appropriate trade-off between cost today and how society values the reduced risk. Faced with the example of an airplane with a one in twenty chance of crashing, a significantly more expensive train ticket might look more attractive. Faced with climate catastrophe odds of 50-50 or even 1 in 20, the costs of slimming down the fat tail look more acceptable. A firm emissions target based on achieving an acceptable level of risk of catastrophe would more clearly acknowledge the limitations of our current scientific and economic tools.

It would be a relief if there was a precise way to measure the trade-off between cost now and damage in the future, as we've done with other pollution problems, but Weitzman's analysis shows it is at best a very rough estimate that misses the significant risk of catastrophe. We would hope Congress and the administration make a strong economic and scientific argument for considering the risk of catastrophe and align the politics with this sobering reality, designing an environmentally robust, ambitious policy that puts that fat tail on a strict diet. •