



Treading Thin Air

By Geoff Mann

It's easy to forget that this isn't the first time the world has seemed to teeter on a precipice. 'To the question how shall we ever be able to extricate ourselves from the obvious insanity of this position, there is no answer,' Hannah Arendt wrote fifty years ago, reflecting on the nuclear arms race between the US and the Soviet Union. Then, as now, an apocalyptic mood had become quite normal. There is some solace in the fact that as bad as things seemed to Arendt, the world did not end in 1970, or in any of the ends-of-days foretold before and since. Is there anything that makes the present moment different? Although we don't lack for murderous leaders in charge of nuclear stockpiles, it isn't fingers on triggers that pose the greatest threat to our survival; indeed, nothing momentous need happen at all. If we simply keep doing what we are doing, then to the best of our knowledge, more and more of the planet will catch fire or be submerged under water; coastlines will wash away, glaciers collapse and rivers dry up; soils will desiccate and blow away; and millions will be on the move or dying of disease. To put it in the careful terminology of the Intergovernmental Panel on Climate Change's 'Summary for Policymakers', we have 'high confidence' that on a business-as-usual trajectory, widespread, irreversible social devastation eventually becomes 'very likely'. The question, of course, is when 'eventually' will be, and whether we can change course between now and then. One thing we do share with the frostiest years of the Cold War is how little we know about the way this disastrous future might unfold. It's true that those who modelled the dynamics of 'nuclear winter' had the advantage of knowing what nuclear destruction looked like, if only on a limited scale, and what the 'tipping point' would involve (the dropping of bombs, though it wasn't clear how many would be required to get there). Beyond that they could only speculate. In the event of nuclear war, how much of the earth's surface would burn, how high in the atmosphere would the smoke rise, and how fast and far would it spread? Would it matter if the war started in winter and not summer? How would crops respond to the reduction in sunlight? Government reports from the 1970s and 1980s are eerily similar to climate projections today: 'Nuclear smoke clouds generated by a relatively low level of explosions (e.g. a hundred megatons) could cool the earth's surface by 36° to 72° Fahrenheit.' That's a pretty imprecise estimate; it points one way (few thought the planet would get warmer), but covers quite a temperature range.



Today's climate models produce somewhat more tightly bounded results, but the uncertainties are just as overwhelming. We know that we are headed in the wrong direction – the planet is heating up – but it remains unclear how far we have gone and how fast we are moving, especially when it comes to the impacts on planetary life. Our increasingly sophisticated climate science modelling may give us a clearer picture of some aspects of future warming, but we can't predict how the world's human and non-human systems will react. We don't have any useful analogues, and we can't run experiments on a practice planet to study their effects. When it comes to human welfare, as the economist Martin Weitzman once put it, attempts to anticipate the impact of climate change, especially beyond the short-term, face an 'immense cascading of huge uncertainties' of 'truly stupendous' scale and scope: There exists here a very long chain of tenuous inferences fraught with huge uncertainties in every link beginning with unknown base-case GHG [greenhouse gas] emissions; then compounded by huge uncertainties about how available policies and policy levers transfer into actual GHG emissions; compounded by huge uncertainties about how GHG-flow emissions accumulate via the carbon cycle into GHG-stock concentrations; compounded by huge uncertainties about how and when GHG-stock concentrations translate into global mean temperature changes; compounded by huge uncertainties about how global mean temperature changes decompose into regional temperature and climate changes; compounded by huge uncertainties about how adaptations to, and mitigations of, climate-change damages are translated into utility changes – especially at a regional level; compounded by huge uncertainties about how future regional utility changes are aggregated – and then how they are discounted. The resources people have to manage this uncertainty remain largely the same as in the past: a mixture of information and doubt, faith and fatalism. What is different is the vast expansion in the range of what is now thought possible, which is no longer bound to the patterns of the past. Communities all over the planet are falling forwards into a future for which history is probably not a useful guide. If there are limits on the range of possible futures, they will become clear only after we, or a substantial proportion of us, are gone. As a result, the words we use to calibrate our reality seem less and less like accurate descriptions of the conditions they are supposed to name. If 'crisis' is so continuous a state as to be 'normal', what help is either term?



When 'five hundred year floods' happen every year, and 'human extinction' isn't ruled out as one plausible future, it might be time to acknowledge that what's lacking isn't simply – as researchers from the Cambridge Centre for the Study of Existential Risk recently suggested – a more precise 'lexicon of global calamity' for the 'climate endgame'. Of course we must rigorously assess where we are and what we know about the years to come. But instead of trying to fabricate precision out of desperation, or calculating just because we can, we need to take stock of what we do not or cannot know. Those whose business it is to understand climate change and its effects, such as oceanographers, climate scientists and economists, deal constantly with uncertainty. They build models, take measurements and produce 'scientific knowledge' about the real world which circulates in that world, perhaps helping the rest of us to manage uncertainties in our own lives. The sciences deal with uncertainty through the use of probability. The Marquis de Condorcet once called probability 'the kind of certainty that pertains to reality'. It takes account of the fact that for important quantities, from climate sensitivity (the way global temperature will respond to the doubling of atmospheric CO₂ concentrations) to the rate of population growth to the effect of taxation on energy demand, we cannot know the 'correct' value, but only posit a range of possible values. Taken together, these are expressed as a probability density function, or PDF: the x-axis presents the range of possible outcomes and the y-axis the frequency or likelihood of each outcome. Say, for example, you have a cardboard box filled with an equal number of red balls and blue balls. You pull a ball out at random, then put it back, then do the same over and over again, each time recording how many times you have pulled out red or blue balls consecutively. Over time you will likely find that the most common outcome is just one ball, red or blue; you will get two of the same colour in a row frequently, too. But it's rare that you'll pull out six or seven same-colour balls in a row, and you might not get twenty in a row even if you went on for days. The graph of these outcomes takes the shape of what is probably the best-known PDF, the bell-shaped curve of the 'normal' or 'Gaussian' distribution (after Carl Friedrich Gauss), with a single symmetrical peak at one blue or red ball, descending more or less steeply on each side to the 'tails' of the distribution, which cover the least likely outcomes. The 'confidence' with which we can expect an outcome within a defined range either side of the centre of the distribution is measured by the proportion of the total area under the curve that lies within that range: if 90 per cent of the PDF's 'density' lies between three consecutive red and three consecutive blue balls on the x-axis, then we say we have 90 per cent confidence that a random run will fall within that range – something like the



probability of a probability. When it comes to complex processes like climate change, however, the uncertainties are considerably harder to manage, and the PDFs less well anchored to what we can say we 'know' from theory or data. For many of the deep structural uncertainties that Weitzman listed, we aren't in a position to gather empirical data. Take climate sensitivity. We don't know what the global temperature response to doubled CO₂ concentrations will be. The current 'best estimate' is an increase of about 3°C, but many studies suggest it could be 4.5°C, and a few indicate a rise as high as 8°C. Others argue that because it will take hundreds or thousands of years for the climate to reach an equilibrium even after atmospheric CO₂ concentrations have stabilised, climate sensitivity is essentially unknowable and not worth trying to ascertain. Many modellers have turned their attention to 'transient climate response' – the real-time effects of increased GHG concentrations – instead. That's the physical science. Matters are at least as complicated on the social side. How will the population growth rate (itself a function of complex dynamics) change over the next fifty or a hundred years? Mass mortality events such as pandemics are almost impossible to anticipate, but we know they can happen. How will the price of renewable energy change over the same period: will it plummet and make fossil fuels redundant, or will it decline more slowly? We don't know what future population growth or energy prices will turn out to be. Yet this knowledge is fundamental to the way we model the future impacts of climate change. Models are essentially sets of equations, organised to simulate the way we think geophysical, biological, and social processes interact. Some are vast in their complexity, connecting and aggregating the impacts of many different dynamics at very fine spatial resolution, simulating, for example, the interaction of cloud cover, solar radiation, and atmospheric and ocean dynamics for every single cell with dimensions as small as 0.1° longitude by 0.1° latitude, in a grid covering the entire planet. Such models can only be run on expensive supercomputers, and may take a day or two to make the projections for, say, the next twenty years. Other models, of a much lower complexity, are calibrated to emulate the results of the big guns. Scientists can run these 'emulators' in a lab, using much less computing power and much more quickly. But increasing or decreasing the complexity of the models provides no answer to the fundamental question of how we can calculate anything at all when there are so many known unknowns, to say nothing of the 'unknown unknowns'. You can't solve an equation in which one of the key parameters is undefined: there is no meaningful solution to $2 \times \text{we're-not-sure}$. One way to deal with this uncertainty is to choose a value or 'point estimate' that accords with common wisdom or the modeller's best guess.



But many dynamics have a range of possible outcomes, and the uncertainties are often so extreme that this practice is difficult to justify. Take, for example, the rate of economic growth. Setting aside the fact that different regions will continue to experience growth at very different rates, ‘expert’ estimates of the average annual global growth rate range between 0.5 per cent and 3.5 per cent across the next 75 years. You could simply pick, say, a figure of 2 per cent and run the model, but that will drastically narrow the usefulness of the results. An alternative might be to run it multiple times using a range of values then average the results, but that’s to assume all the values are equally probable, which is highly unlikely. The struggle is to find a sweet spot between illusory exactitude and unhelpful handwaving. Even the simplest models might involve fifteen or twenty uncertain parameters, so this is not a small problem. The standard approach to dealing with multiple uncertain parameters is the Monte Carlo simulation. The model is programmed to run many, many times over, each time randomly selecting a value from each of the parameters’ PDFs. In a random draw, we can assume that the more likely a value is, the more often it will be drawn, so the results will reflect the relative likelihood of the parameter values. If population growth, for example, is distributed between 1 and 4 per cent, but the tails (1 and 4 per cent) are considered unlikely while something closer to 3 per cent is suggested by current estimates, then more of the model’s runs will reflect the 3 per cent than the tails. The trouble here is that it isn’t only the range of parameter values that is uncertain: the shapes of many of the PDFs are pretty much unknown, even unknowable. We could just assume everything is normally distributed, make an informed estimate of each PDF’s centre (the mean) and standard deviation (the degree to which the distribution is spread out), and hope for the best. This is in fact standard practice when modellers have nothing more to go on. But it’s problematic when we have strong evidence, or an intuition, that a parameter is not normally distributed. If population growth is expected to fall between 1 and 4 per cent, but 3 per cent is considered most probable, then the PDF is not normally distributed: it will be ‘skewed’ towards the higher end, with values above the middle of the range more likely. There is an array of PDF ‘families’ modellers can use to ‘fit’ ranges of parameter values that deviate from the normal distribution; some peak near the low end, some near the high end; some decline towards the tails very steeply, others slowly descend over a spread of extreme values. This technical detail matters a lot, because the shape of the PDF determines the probability that a randomly drawn parameter value is in the tails – in other words, in the extremes of the distribution, like climate sensitivity of 8°C.



If you run a model hundreds or thousands of times using PDFs that make it very unlikely extreme values will register, then the results will probably suggest we can be less worried about catastrophic futures. In fact, the choice of PDF can go some way towards predetermining a model's outcomes. If you select a PDF because it makes your calculations easier, or you simply guess, then the model might not be merely unreliable, but unreliable in ways that can't be gauged. Then we'll 'march on', as the MIT economist Robert Pindyck puts it, acting as if we have knowledge we don't really have. This is sometimes called the problem of 'fat tails', that is, possibility that PDFs underestimate the probability of extreme 'tail events'. Models with thin tails concentrate probability away from the extremes and can effectively assume away disaster, especially if, as is not uncommon, the PDF defines extremes as having such low probability that we can ignore them completely (to increase computing efficiency, perhaps). If, on the other hand, our parameter selections told us that extreme outcomes were far more likely, and the probability of ecological catastrophe wasn't necessarily remote, then the models would tell us to take them much more seriously. Related problems beset the methods used to estimate economic losses. Virtually all attempts to determine an appropriate carbon price use models incorporating a 'damage function', which calculates the costs incurred by global temperature changes. But the damage function is arbitrary - 'made up out of thin air,' in Pindyck's words - largely because we don't have much to base it on beyond intuition: no theory, no data, no nothing. Some of Pindyck's colleagues think he overstates the difficulties, but even they admit that 'the best we can do' is use historical data and hope the trends remain stable in the years ahead. Pindyck's response amounts to 'I rest my case': The basic problem is that we know as little about the correct probability distributions as we do about the damage function to which they are being applied. What can we possibly learn from assigning arbitrary probability distributions to the parameters of an arbitrary function and running Monte Carlo simulations? I would argue that the answer is nothing. It gets worse. The damage function is usually mapped onto 'utility' - the principal measure of human welfare in economics - to determine how economic changes will affect the 'social welfare function'. The welfare function is enormously important, but tough to define. One of the things it is supposed to tell us is how the impact on welfare changes as damages accrue. This is another quantity determined by the modeller, and there is no reason to expect it to be constant: the welfare losses associated with losing one litre of water for agriculture will be far greater if it is the last litre than if it is one of billions.



One of the more common ways of getting at expected welfare losses is the use of 'willingness to pay' (WTP) measures, which register how much 'we' would be willing to give up to avoid a given level of damage. WTP has been central to environmental economics for decades, but it has serious drawbacks. For one thing, there's potentially a very big difference between being 'willing to pay' and 'having to pay'. And, even more fundamental, it isn't clear that people grasp what an environmental 'amenity' (a curiously restrained term) such as clean water is 'really' worth. This is not to suggest that we might be making a big deal out of nothing. The concerns about fat tails focus on the upper tails of the probability distributions: only lunatics and shills for the fossil fuel industry would suggest that the lower tails might be fat and that we're worrying about outcomes we are unlikely to experience. The point of highlighting the vertiginous degree of uncertainty is that we might not be making nearly as big a deal of climate change as we should. We are, as a result, tragically under-prepared for the possibility of really bad outcomes, yet at the same time far too confident in our level of preparation. If we took this possibility more seriously, our models would indicate potential losses in productivity, consumption and welfare that dwarf what currently counts as professional common sense – and that's to say nothing of mortality, a topic most damage estimates are keen to avoid. In making probability the main tool we use to manage uncertainty about climate change, we have perhaps taken Condorcet's remark too literally. If probability is the 'kind of certainty that pertains to reality', that doesn't mean we should treat it as certainty. Yet those making key decisions treat probability as if, paradoxically, embracing uncertainty meant taming it. This gives it a special political utility. 'Probability cannot dictate values,' Ian Hacking wrote in *The Taming of Chance* (1990), 'but it now lies at the basis of all reasonable choices made by officials. No public decision, no risk analysis, no environmental impact, no military strategy can be conducted without decision theory couched in terms of probabilities.' Hacking went on to say that by 'covering opinion with a veneer of objectivity, we replace judgment with computation.' He wasn't talking about the economics of climate change, but he might as well have been. His point was that computation, which is what modelling is all about, suggests precision, even when it shouldn't. Precision has political value because it appears to be above politics; the assumption is that technical complexity gets us closer to the 'truth', which is (or was, until recently) a useful thing to have on one's side in an argument. In light of all this unknowability, there are people, experts among them, who question the 'scientific legitimacy' of the knowledge that climate policymakers rely on.



I think this goes too far. In many ways, climate scientists' and economists' explicit, careful use of probability to treat uncertainty is exactly what we should expect of science. Every climate scientist I have encountered wants to be honest about what they know and what they don't, and is aware of the ways their knowledge might be applied in the 'real world'. The science has made it possible to ask questions we couldn't otherwise formulate and to look again at our assumptions. But what will matter in the end is the politics: however much we reduce our carbon emissions, we are going to have to live with uncertainties that have existential implications. The way we respond to that condition is above all a political question about the kind of uncertainty that pertains to our reality, which is not a question probability can answer. Another way of putting it would be to say that the problem isn't probability, but the faith we have put in probability as a way of dealing with uncertainty. This is about more than the natural scepticism of the non-expert, who might demand to know, for example, how we can be sure our PDFs are even in the right place on the x-axis given that certain crucial parameters are at best educated guesses? It's possible that if one could just sit down with the experts they might be able to allay these fears, but the technical literature certainly doesn't. Beyond this, though, I think there is something about probability that fundamentally misconstrues reality. When the models run hundreds, even thousands of times, each time using slightly different parameters, it is as if there were dimensions parallel to our own, all with of more or less plausible futures, from which our own world is assembled. Modelling supposedly gives us access to the full range of these futures, so we consider them without having to actually experience them. But in our collective reality, we don't get thousands of runs at life and future. We just get one. In *The Emergence of Probability* (1975), Hacking tells a story about a naval base during the Second World War. The soda machine on the base had always charged 5¢ for a bottle of Coke. Then one day the price went up to 6¢. But the machine could only take nickels, so the operators decided that every sixth purchase would be a dud, and the buyer would be out of luck. The residents on the base accepted the new arrangement. Sometimes they would lose out, but most of the time they would get the Coke they'd paid for. The casual visitor to the base, however, only had one shot, and there was a one-in-six chance he'd get nothing. This is the sense in which multi-simulation probabilistic models miss the visceral content of the reality they are intended to describe. Whatever 'mid-course corrections' we make along the way, all but one of our planetary futures will inexorably fall away as time passes. We have put ourselves in the position of Hacking's 'casual visitor'.



When Martin Weitzman, who along with Robert Pindyck and Nicholas Stern has done more than anyone to get mainstream economics to take catastrophe seriously, said that all his work points towards a 'generalised precautionary principle', he was saying that we, as a species, cannot afford to use the Coke machine because we have just one coin to put in the slot. It is for these reasons that the philosopher C.S. Peirce argued that 'in reference to a single case considered in itself, probability can have no meaning.' He explains with a thought experiment: If a man had to choose between drawing a card from a pack containing 25 red cards and a black one, or from a pack containing 25 black cards and a red one, and if the drawing of a red card were destined to transport him to eternal felicity, and that of a black one to consign him to everlasting woe, it would be folly to deny that he ought to prefer the pack containing the larger proportion of red cards, although, from the nature of the risk, it could not be repeated. We would, of course, all choose the red pack, even though we don't get repeated draws. Peirce, remarkably, looked on this faith in probability as proof that 'logic is rooted in the social principle': our interests 'must not stop at our own fate, but must embrace the whole community'. This doesn't hold for climate change, however, because picking the black card affects everyone; there is no second or third or fourth chance, no other people or planets to benefit from the probability that a decent outcome will eventually come to pass. At this point, it's time to stop playing the game. This is where climate policy, despite policymakers' best efforts, cannot help but become politics. We have reached a stage of global warming at which every decision is critical: we don't know when our last chance will have been. So when, for example, we base the vast part of our climate policy on offset markets and carbon taxes, as we are doing, and proceed to calculate the social cost of carbon to determine an 'optimal' carbon tax that 'efficiently' manages the 'trade-offs' between the costs and benefits of emitting GHGs, we are doing something much more dangerous than is usually acknowledged. A precise calculation of the 'optimal' carbon tax is nothing more than a claim that the best way forward is to perch the gargantuan machine of contemporary capitalism as close as possible to the precipice without tipping us all over the edge. That is neither efficient nor optimal. It is a myopic and recklessly arrogant approach to the unknown fate of life on earth. What we need is a much more honest assessment of what we do not or cannot know, which is, among other important things, where the edge is. We might, in fact, be past it already, trading thin air like Wile E. Coyote before the fall. Today's politicians don't like uncertainty: it introduces doubt.



Yet we are in desperate need of a politics that looks catastrophic uncertainty square in the face. That would mean taking much bigger and more transformative steps: all but eliminating fossil fuels, for a start, and prioritising democratic institutions over markets. The burden of this effort must fall almost entirely on the richest people and richest parts of the world, because it is they who continue to gamble with everyone else's fate.

-Geoff Mann

Geoff Mann is at work on a book about uncertainty. He teaches at Simon Fraser University

This essay first appeared in [The London Review of Books](#).

©LRB, 2023