October 25, 2010 version

CATASTROPHIC CLIMATE CHANGE

Mike Toman, World Bank Development Research Group

Draft, subject to revision; please do not cite or quote. Responsibility for content is the author's alone.

INTRODUCTION

The question of how to assess prospects of climate change catastrophes has been the focus of a great deal of recent research and debate. An example of the classic conundrum of low probability – high consequences events, a climate change catastrophe is a highly unlikely event, but if it did occur it would severely affect well-being across the world – though it would affect poor countries much more seriously than richer countries.[1] The larger geographical scale of climate change catastrophes distinguishes them from more localized extreme events. The consequences of catastrophes also are in varying degrees very costly, if not possible, to reverse.

Examples of global catastrophes include very large and relatively rapid increases in sea level from faster melting and collapse of ice sheets, slower changes in ocean currents that have insidious effects on weather patterns, and large scale destruction of forests and other ecosystems. fairly rapid loss of global forest cover. Unlike sudden disasters such as earthquakes, the onset of these events is measured in multiple decades or centuries; but once they occur it is impossible to reverse the impacts. Other permanent effects of climate change are anticipated to be increases in the frequency and severities of droughts, floods, and hurricanes, leading to corresponding destruction of crops, water supplies, and coastal infrastructure. While each of these individual events is a more localized disaster, the cumulative effect could be a global catastrophe created by the "cascading consequences" of more localized disasters occurring in relatively quick succession, each amplifying the effects of others.[2]

A key step in evaluating risks of climate change catastrophes is to assess not only the impacts on the physical climate system, but also the consequences in terms of human impacts. The most immediate implication is that while a physical "tipping point" may be reached at some unknown future date T^0, the human impacts will evolve more slowly, reaching an intensity viewed as catastrophic only at some date $T^1 > T^0$. This distinguishes climate change from, for example, the risk of catastrophe posed by a gigantic volcanic eruption, or nuclear war. While a gradual onset of impacts will not prevent a catastrophe if reversal is not possible, it can provide a window of time for major action to avert or adapt to the threat – if signals of the changes are detectable. More fundamentally, the assessment of what constitutes catastrophic human impacts involves not just climate change and earth system science, but also inherent value judgments about what magnitude and speed of consequences are deemed to be catastrophic. For example, the now-often-cited "scientific near-consensus" about the urgent need to hold warming to less than 2°C relative to pre-industrial times reflects more than a natural science evaluation of climate change impacts.

Climate change catastrophes pose a familiar challenge for assessing the impacts of low probability – high impact events: while exact quantification is not possible, the most extreme adverse impacts from climate change—say the worst 1% of scenarios—may account for a large portion of losses in expected value terms. This implies that focusing primarily on a trajectory of more likely anticipated climate change

[1] In terms of absolute numbers, losses are likely to be larger in richer nations. As a percentage of GDP, however, less developed countries are likely to face higher damages since most are more dependent on agriculture and less likely to have the resources to adopt measures that could reduce damages.

[2] This possibility appears to have received little systematic attention in reviews of climate change impacts by the IPCC and others, though it figures prominently in discourse about national security consequences of climate change.

damages may miss an important part of the problem.[3] Yet, these consequences of an unlikely but possible climate change catastrophe need to be weighed against a variety of other risks society faces.

Further complicating the problem is that climate change catastrophes may be better characterized by ignorance than uncertainty. That is, not only do we not know the probability of a particular mega-catastrophe occurring, we do not even know many of the possible outcomes. A catastrophe from climate change could stem from a cause or have impacts that currently receive little attention.[4] Some authors have suggested that this level of ignorance, coupled with the very low probability of an event and the possibility of extremely severe impacts, hamstrings the use of rational-choice based methods for analyzing response options. This in turn requires confronting the possibility that attitudes of the broader public about such events will not align very well with the results of a more systematic evaluation of the pros and cons of different response options, raising questions about what sets of preferences and beliefs should govern policy making.

CLIMATE CHANGE CATASTROPHES

The most widely discussed large-scale impact of climate change is **global sea level rise**. The collapse of the West Antarctic Ice Sheet (WAIS) or Greenland ice sheets could lead ultimately to a sea level rise of several meters, with consequences great enough to be considered a global catastrophe in the absence of massive and costly relocation because of the number of people living near the coasts. A key uncertainty is how rapidly this change in sea level might occur. Previously it had been thought that such large changes might require much longer than a century, but some recent studies suggest that substantial change could occur in this century. Anthoff et al. (2009) report figures for world losses (based on 1995 baseline conditions) that are relatively small – on the order of 0.5% of world GDP for a 5 m rise. Dasgupta et al. (2007) report figures for developing countries on the order of 6% of GDP, those these estimates do not take account of possibilities for ex ante efforts to mitigate risks. On the other hand, estimates based on historical baselines will tend to under-state the economic impacts of sea level rise by not taking account of likely future growth in the coming years in the share of GDP concentrated in coastal areas.[5]

A second important category of global catastrophe risk involves **disruptions of ocean circulation from climate change,** with potentially disastrous effects on regional weather patterns and long-term climate (Vellinga and Wood 2008). Such impacts are most commonly seen as developing over many hundreds of years. In contrast, **very large-scale ecosystem disruptions** could occur significantly sooner. Changes in ecosystems resulting from changes in temperature and rainfall incidence and increased climate variability have the potential to cause very significant loss of biodiversity—on the order of 20-30% extinction within a few decades. There is also the prospect of major changes in vegetation, in particular, irreversible

[3] For many classes of disasters and catastrophes, the most extreme small percent of the situations represent a significant proportion of the losses. We have witnessed this "fat tail" phenomenon recently with terrorist deaths and losses in a financial crisis. 9/11 and the 2008-09 financial meltdown caused more deaths and dollar losses respectively than all terrorist incidents and financial catastrophes in the post WWII era. With such phenomena, losses are better characterized by a power law than by a normal or even lognormal distribution. The debate about fat tails in relation to climate catastrophes has been a subject of lively recent debate among Weitzman, Pindyck, Nordhaus, and others.

[4] The history of the past 40 years is sobering with respect to the ability to identify catastrophe risks. In 1970, nuclear war would have been the leading contender for any world catastrophe, and looking forward few would have predicted the major looming threats of the current era, which would include not just climate change, but also global pandemics and terrorism.

[5] Using 1995 data, it has been estimated that around 400 million people would be impacted by a 5 m rise in sea level and that a WAIS collapse in 100 years could cause, at the peak, 350,000 forced migrations a year for a decade (Nicholls, Tol and Vafeidis 2008).

conversion of forest to grassland, desertification, and acidification of the ocean (Smith, Schneider, Oppenheimer et al. 2009). Another cause for significant concern is the possibility that positive feedback effects in the climate change process itself could occur (e.g., liberation of trapped methane from ice, rapid increases in CO_2 from vegetation dieback, or increased heat absorption as glaciers retreat), causing the abovementioned changes to occur more rapidly.

There also has been significant scientific research on how climate change can effect more localized disasters, such as heat waves, flooding, droughts, and changes in hurricane frequency or intensity. Less understood is how **a number of smaller disasters** all occurring over a relatively short time period could mutually reinforce each other in such as way that the **resulting "cascade of consequences" becomes a global catastrophe**. Extreme events can have secondary consequences that generate substantial amounts of additional damages; secondary consequences in turn can trigger tertiary consequences that further amplify the adverse consequences; and so on. One example would be if increased drought from climate change in different regions successively caused a series of local food shortages to occur in close proximity, leading to political instability, a breakdown of civil order, large-scale migration for survival, and regional conflicts. Another example could be a series of local fires occurring in climate-stressed forests and grasslands overly widely dispersed areas, adding up to a large-scale destruction of resources, ecosystem services, and livelihoods over a large area.

The compounding or amplifying effects of individual adverse impacts would be the result of exceeding the resilience of a number of local socioeconomic systems in rapid succession. More frail components of socioeconomic systems, such as marginal subsistence agriculture, represent potential places of vulnerability. Cascading-event catastrophes could occur much more rapidly than the slower-onset global impacts discussed above, especially as climate change accelerates and greater negative impacts occur at local scales. It is possible that more comprehensive local monitoring of disaster risks may facilitate the development of early warning indicators for cascading catastrophes. For example, if several years of historically unusual drought weakened agricultural systems in many vulnerable parts of the world, there would be a stronger basis for concern about cascading consequences than if agricultural failures were not occurring in such rapid succession. However, the time interval for action to avert the potential catastrophe could be short.

Traditional responses to the risk of extreme events are of limited value in mitigating risks of a mega-catastrophe. The underlying changes in the climatic system could not be reversed over any time scale relevant for decision-makers. Traditional insurance mechanisms will not function effectively for this type of event, because the risks are "systemic" and cannot effectively be reallocated to diversify. Moreover, significant international transfers from richer to more vulnerable poorer countries are unlikely when a catastrophe affects broad swaths of the world.

EVALUATING CLIMATE CHANGE CATASTROPHE RISKS

The traditional economic model for decision making under uncertainty is expected utility theory, in which decision makers maximize the utility they receive from potential outcomes weighted by the probability the outcomes will occur. In the climate change economics literature, GHG abatement policies with the expected net benefits over time are identified using dynamic Integrated Assessment Models (IAMs) that compare the anticipated costs of abatement with avoided damages from climate change over time. By and large these models are deterministic and are used for scenario-based comparisons of policies under different assumptions about climate change damages and abatement costs. However, a literature has developed in which catastrophes are treated as (usually known) large-scale rapid-onset economic damages with an uncertain date of occurrence, the probability of which increases as atmospheric GHG concentrations rise.[6]

A common finding in these studies is that while the risk of such catastrophes increases the expected economic benefits of more rapid GHG mitigation, the effect is not that significant qualitatively unless the probability of nearer-term catastrophe is quite high, the size of the catastrophe is truly astronomical, or the discount rate used to value future catastrophic impacts is quite low. The scientific information on catastrophes summarized above indicates that catastrophes are extremely unlikely in any time frame short of several decades at the very least, and that while the ultimate effects may indeed be huge, the most severe impacts will develop only gradually. Until scientific understanding of climate change catastrophes leads to stronger findings on their proximity and severity, the choice of discount rate will be the most important determinant of the cost of future catastrophes in the expected-utility framework.

The discount rate issue in turn continues to be very hotly debated, and only a very brief summary of key points is offered here. Two strands of positive analysis has argued for applying a lower discount rate to longer-term climate change costs, including catastrophes, than might be inferred from research on consumer time preference or rates of return on investment. One is that individuals may discount the future hyperbolically, so rates of discount decline and ultimately plateau at a fairly low number as one goes out into the future. The other is that when one accounts for the higher marginal utility of income for the poor facing more adverse impacts from climate change, then under reasonable assumptions the effective time discount rate after adjusting for distributional differences is reduced. In addition, if climate change has the most severe effects on longer-term economic growth when growth itself is more likely to be weak, then policies to reduce the threat of catastrophe will have a lower effective discount rate because of their contribution to reducing intertemporal economic risk.[7]

Even with these considerations, however, the resulting implied discounting of future over current returns may not be small enough for catastrophes to carry major weight in evaluating the potential impacts of climate change. Unless the discount rate is under 1%, and perhaps even close to zero, severe future consequences that will not arrive for some time and are not world-threatening may still be too "telescoped." Stern and others have addressed the issue of discounting by using normative arguments to suggest a discount rate at or near zero is in fact appropriate. Two other arguments, not so dependent on normative precepts, may also add weight to the importance of catastrophe risks in evaluating climate change impacts.

Hypothesis 1: People are Not Expected Utility – Maximizers

There is a growing literature from behavioral economics and psychology which demonstrates that individuals do not consistently make decisions according to the expected utility paradigm.[8] If individuals are only boundedly rational, they have neither the time nor the capacity to fully assess the consequences of decisions. In that case, individuals adopt certain rules of thumb and mental shortcuts to make decisions. These so-called heuristics can lead to choices that depart from predictions of expected utility theory.

[6] References – Kverndokk et al, Pizer, Nordhaus. Earlier foreshadowing by Manne.

[7] [add references] Strictly speaking, the second and third arguments are not about the actual rate of time preference, but rather about how factors related to distributional impacts and risk that enter the maximand of the intertemporal utility calculation affect the implied discounting of future over current returns.

[8] This discussion is taken from Kousky et al (2009), which contains references to the relevant behavioral economics literature.

When thinking about possible disasters, it has been found that people tend to be over-optimistic, thinking negative outcomes are less likely to happen to them. When a risk is highly emotional, however, people can disregard probabilities altogether, treating all outcomes as equal ("probability neglect"). Individuals also seem to place an added value on certainty, preferring to reduce a small risk to zero by more than they value reducing a larger risk by a greater amount. Errors of commission are viewed as worse than errors of omission. This can lead to a tilt to the side of inaction.

Experimental also has found that context matters, often significantly, a when making decisions. For instance, when probabilities are unknown and must be estimated, individuals have been found to assess an event as more likely when examples come to mind more easily (the "availability heuristic"). People can disproportionately prefer to maintain the status quo in their choices, even if conditions or options change. Individuals sometimes "anchor" their preferences on an available piece of information, and fail to update their assessments adequately in the face of new information. Individual choices are also strongly affected by the way that information is presented. Thus, individuals may make different choices for the same decision if it is merely phrased differently ("framing effects"). Choices depend upon the extent to which a risk evokes feelings of dread. Personal utility also is sensitive to individuals' perceptions of equity and fairness.

These various behavioral attributes can imply higher or lower values attached to catastrophe risks than would be implied by expected utility theory. The former would follow from dread or the evaluation of all catastrophes as roughly equal in likelihood. The latter would follow from optimism bias, or a preference for reducing small and familiar risks to zero over reducing more substantially an unfamiliar risk – of which climate change catastrophe certainly is an example. While the direction of bias has to be assessed empirically, the existence of these various "non-rational" attitudes raises an important but not new question for evaluating climate change catastrophe risks in setting public policy: if decision makers believe they have better information than the general public and that they are less subject to emotional biases, to what extent should their valuation of alternatives supersede those of members of the general public?

Hypothesis 2: People are Non-Egoistic Expected Utility – Maximizers

A second approach that has been taken in the literature for addressing long-term threats posed by climate change is to see individuals today, imperfect information and all, as interested in more than maximizing the discounted present value of their lifetime expected utility streams. One can broadly define this as altruistic preferences, but this label can cover several different forms of preferences.

A traditional approach to altruistic preferences is to include some measure of next-generation or other future utility in the preferences of members of the current generation. In this setting, individuals will weigh the potential costs of a climate change catastrophe in terms of its anticipated impacts on future welfare, as well as the possibly slight impact on current individuals' egoistic well-being. Consequently, individuals will derive utility in part from the "bequest they leave to the future in terms of a lowered (endogenous) risk of a climate change catastrophe. However, there are both theoretical and empirical reasons to expect individuals to discount the welfare of future generations relative to their own egoistic welfare. This takes us back to the question previously mentioned in the context of time preference, as to how powerful an influence this form of altruism might be in the current generation's assessment of risks of climate change catastrophes.[9]

A second approach is to depart from a purely utilitarian framework by supposing that individuals see themselves (or should do) as having a moral obligation to future generations. This mixing of obligations and conventional utilitarian motivations implies some degree of lexicography in individuals' preferences – or, critics of utilitarianism might say, an innate failure of the standard economic model to describe what really motivates people. In this view, if a potential future catastrophe threatens to impose a morally inacceptable burden on the future, people will be (or at least can be) motivated to endure potentially extra-ordinary sacrifices to reduce the threat. The expression of that moral sentiment by individuals as citizens and stewards, versus utilitarian consumers, would be found through public choice exercises like voting for tough restrictions on future GHG emissions.

This conception is both stimulating and frustrating, since it does not offer any straightforward way of assessing how economically significant is the threat of a future climate change catastrophe. Aside from uncertainty about what the triggering level of threat to the future might be, does one regard current almost universal reticence to support tough GHG restrictions as due to (correctable) moral failing? Lack of information? Lack of leadership? The result of rational leadership, because the threat of climate change is seen as less significant than other threats or because international collective action problems have not been solved?

A third possible approach that has received less attention is that individuals have preferences that include some notion of "planetary health" as a global public good. Rather than seek to describe concern about risks of catastrophe from climate change as deriving only from more fundamental concerns for intergenerational altruism or fairness, one could posit that individuals derive some direct benefit from having greater confidence in the ability of planetary systems to remain undisrupted, without the need to unpack the rationales in terms of future human well-being, satisfaction of moral sentiment, or a pure existence benefit. This approach allows one to sidestep some of the difficulties encountered in either the altruistic utilitarian or moral-obligations conceptions. In particular, the normative approach to setting discount rates can be embedded in a framework of preferences without having to be an ad hoc add-on.[10] However, this does not get around the huge empirical problems in assessing the value that members of the current generation might place on reducing risks of future climate change catastrophes. [11]

CATASTROPHE RISKS AND RATIONAL CHOICE APPROACHES TO POLICY

While it is certainly possible to debate the capacity of expected –utility types of analyses to adequately capture the social opportunity cost of climate change catastrophe threats, it is in cases like this that a disciplined application of rational-choice based analysis more broadly defined can prove most useful. A

---

[9] Current individuals also could believe, as Schelling for example has suggested, that other kinds of bequests to the future would have higher value; or they could further discount bequests of a less risky climate out of concern that unless the "chain of obligation" is maintained, something impossible to assure, the sacrifice made today would be wasted in the future.

[10] A fundamental criticism of conventional expected-utility analysis for assessing future climate change risks is that it combines conventional time-preference considerations in assessing the opportunity cost of reducing threats with the explicitly ethical question of how much the current generation will feel willing or bound to do in protecting the future.

[11] Ideas like this arise often in literature on environmental stewardship, but I am not aware of many treatments of the idea in economic terms. One example is the paper by Kopp and Portney [ref to add], who describe a thought experiment in which individuals value "well being of the future," and the willingness-to-pay for that value can be discerned through a stated preference valuation effort. While one can debate the merits of the valuation approach even in a thought experiment, the concept is very similar to what I am trying to describe here. Unfortunately, the question of how one would ascertain such valuation remains a barrier to empirical implementation of the concept.

thoughtful, systematic, and transparent weighing of benefits and costs, broadly defined, is at the heart of such an approach. The presence of “deep” uncertainty or ignorance about the types and likelihoods of potential catastrophes means that we must include, in addition to sensitivity analysis on these characteristics, focused analysis of the robustness and flexibility of options in addition to the benefits and costs. With respect to what seem to be behavioral biases in the assessment of catastrophe risks by individuals, decision makers must make (and then defend) informed judgments on behalf of those they serve as to when the seeming biases reflect a high degree of economic risk aversion, or dread, and when the biases reflect other factors (framing effects, optimism bias, and the like) that can be viewed as inaccurate comprehension of the tradeoffs involved.

Posner (2005) argues that uncertainty over benefits and costs should not prevent using the basic structure of cost-benefit analysis for evaluating and comparing options, but that this should be framed in a “tolerable-windows” approach. This involves using a range of plausible risk estimates to help identify levels of spending on reducing risk for which benefits clearly exceed the costs, for which costs clearly exceed benefits. Policies then can be designed with the goal to remain in this window.[12] This approach does not provide or depend on “a number” for how to evaluate the impacts of potential future climate change catastrophes. In particular, it does not treat them as largely irrelevant economically given their low probabilities and long time frames to be realized. Instead it provides flexibility as to how different considerations about climate change catastrophes are brought into the assessment, including risk aversion and concerns about future sustainability as well as costs of risk mitigation, while insisting on transparency and a persuasive argument for how these considerations are to be addressed.

[12] This idea is akin to value-of-information approaches. If one has some confidence in the evaluation of costs of different policies but great uncertainty about the potential benefits, one could investigate how large the potential benefits might have to be to make a case for the selection of one set of options over another in a portfolio. Similarly, if the benefits are reasonably well understood conditional on a catastrophe occurring, but there is uncertainty about the probability of a catastrophe, then one can ask how large the probability would have to be to justify a particular portfolio of actions.