Comparing Alternative Policies Against Environmental Catastrophes

Timothy Besley, London School of Economics
and
Avinash Dixit, Princeton University

Preliminary version
(for comments only)

October 12, 2016

Abstract

We construct a model with three features important in the context of major environmental catastrophes: (1) the distribution of possible damage has a fat tail, (2) the probability of the catastrophic event increases as greenhouse gases accumulate, (3) a technological solution may emerge making conservation efforts unnecessary. We solve the model numerically for plausible values of the parameters, and evaluate the tradeoffs between alternative policies such as prevention, mitigation, and technological fixes.

1 Introduction

The weight of scientific opinion supports the proposition that a range of geophysical changes induced by emissions of green house gases place the world at risk of catastrophic changes which threaten human well-being. A key policy issue is how much current generations should be willing to sacrifice consumption to reduce the risk of a future catastrophe given that the probability of it occurring increases with delay. Debates about the optimal sacrifice to make can usefully be informed by calculating the discounted flow of benefits that accrue from any mitigation strategy. This should help to focus the minds of concerned citizens and policy-makers when deciding how much to invest in scientific research or to impose taxes on emissions.

Existing research has identified some important features of the dynamic interaction between economic activity and climate change. (1) There is much structural uncertainty, which is not satisfactorily captured in a normal distribution but has fat tails. Stern (2007, p. xiv) characterized the issue as "the economics of the management of very large risks"; see also Weitzman (2011, 2014), Barro (2015) and others. (2) There are dynamics with a stock effect: the probability of a rare environmental disaster increases as the stock of greenhouse gases accumulate in the atmosphere; see for example Brito and Intriligator (1987), Dietz (2011), and Kolstad (1996). (3) There is irreversibility on both sides of policy action: if we do nothing and the problem proves serious, the climate and economic activity will suffer long-term damage, but if we spend resources to mitigate or counter a problem that later proves to be less important or nonexistent, we will have sacrificed resources unnecessarily. Thus there are embedded options in both the decisions to act and to wait, and the two must be balanced to determine the optimal policy; see Kolstad (1996).

This paper constructs a model that includes all three features and uses it to examine some policy trade-offs. It is distinguished from the above-cited previous research in different ways. Much of the discussion of fat tails has focused on methodological issues of the appropriate choice of the discount rate and the specification of risk aversion in a utility function; for example Weitzman (2014), Barro (2015), Dietz (2011, esp. pp. 524–7), and Millner (2013). The work of Brito and Intriligator (1987) does not consider uncertainty while Kolstad (1996) has a two-period model that does not allow significant dynamics. Martin and Pindyck (2015) study the choice of which catastrophe(s) to avert when several threaten. And most work does not allow for the possibility that technological progress may allow the problem to be avoided or solved much more cheaply in the future.

The model will be used to explore general effects but the key contribution will be to offer a quantification of the benefits from mitigation how it depends on the underlying parameters of the problem. We do so without having to be specific about risk aversion nor do we need to take a strong stance on the discount rate. However, the approach can allow a reader to download the software and explore any set of parameter values that he or she regards as plausible. To get at policy-implications, we ask what would be the loss in real consumption that society should be willing to make in perpetuity in order to achieve a given discounted gain if a particular parameter of the problem could be changed in a favorable way. For example, how much sacrifice would be equivalent to reducing the risk of the arrival of a catastrophe by x%. We breath life into this by offering some concrete scenarios. This illustrates the heterogeneity in benefits according to which parameters are changed. For the specific cases that we consider, being willing to spend something of the order of 5% of real consumption to avert an environmental catastrophe appears plausible.

The remainder of the paper is organized as follows. In the next section, we lay out and solve our basic discrete time model. In section 3, we discuss the numerical solution. Then section 4 looks at some policy experiments while section 5 concludes.

2 The model

We model the idea of fat-tailed uncertainty of a rare disaster in the simplest and most extreme form, namely as a Poisson process. The stock effect is captured by making the arrival rate of this process an increasing function of a state variable that represents the accumulated stock of greenhouse gases in the atmosphere (henceforth called "pollution" for brevity). Denote this by x, and let t denote time. The state x_t evolves according to the equation:

$$x_{t+1} - x_t = \alpha\,, \tag{1}$$

where the parameter α will become a control variable when we come to consider policy (for example carbon capture and storage).[1]

Before a catastrophe strikes, there is a flow of benefits, for example real consumption, at rate y.[2] We specify a fairly general form of the consequences of a catastrophe: it causes a lump sum cost or wealth destruction K, and after that the economic activity continues with consumption flow and the lump sum costs of future catastrophes reduced to a fraction θ of its former level.[3] All benefits and costs are discounted at the rate r.

Let $\lambda(x)$ denote the probability that a catastrophe occurs in the time period starting with state x; this is an increasing function with $\lambda(0) = 0$ and $\lim_{x\to\infty} \lambda(x) = 1$. Policies such as shifting housing to higher land or building levies to prevent flooding can lower the whole function; we consider this parametrically in the numerical solutions in Section 3.

There is another Poisson process with arrival rate μ, representing technical change such that if this event occurs, the risk of the catastrophe will disappear, i.e. $\lambda(x)$ will drop to zero. Geo-engineering solutions such as placing a layer of reflecting particles in the upper atmosphere might be an example of this even though it is somewhat speculative at present.

Let $V(x)$ denote the discounted present value of the net economic benefits starting from the level x of the state variable. This satisfies the Bellman-type recursion equation

$$\begin{aligned} V(x_t) \;\; = \;\; y + \frac{1}{1+r}\; \{ \; & \lambda(x_t)\,\mu \;\; [-K + \theta\, y\,(1+r)/r] \\ & +\lambda(x_t)\,(1-\mu) \;\; [-K + \theta\; V(x_{t+1})] \\ & +[1-\lambda(x_t)]\,\mu\; y\,(1+r)/r \\ & +[1-\lambda(x_t)]\,(1-\mu)\,V(x_{t+1}) \; \} \end{aligned} \tag{2}$$

On the right hand side, we have the immediate consumption flow y, and continuation payoffs in four contingencies with their appropriate probabilities. First, with probability $\lambda(x_t)\,\mu$, we have the event when a catastrophe occurs, but the magic technology that precludes any further catastrophes also occurs. We pay the lump sum cost K

[1] If geometric growth of the cumulative pollution level in absence of policy would be a better specification than linear growth, interpret x as the log of the level of pollution.

[2] GDP growth in the absence of a catastrophe, and its interaction with x allowing possibly simple normal uncertainty, can be included at the cost of only algebraic complexity.

[3] All these can be interpreted as money equivalents of utilities that embody risk aversion; we leave all that behind the picture to simplify the analysis and enable us to focus on the dynamics. Note that we are implicitly assuming that after a catastrophe, although the GDP flow is reduced as is the wealth destruction from any future catastrophe, the accumulation of pollution continues at the same rate α. The idea is that pollution is caused by inputs to production or consumption, for example burning fossil fuels or deforestation, and not by outputs. What the catastrophe does is to reduce the productivity of these inputs. The assumption can be altered to allow a different rate of increase of pollution after a catastrophe, although the mathematics gets much more complex.

of the first catastrophe; thereafter we have the capitalized value of the reduced consumption flow (reduced because one catastrophe has occurred) $\theta\, y$. The second line represents the event where a catastrophe occurs and the magic technology does not appear, so we pay the lump sum cost K and thereafter get only the fraction θ of the value function V evaluated at the new state x_{t+1}. The third line is where a catastrophe does not occur and the magic technology appears, so the continuation value is the full capitalization of consumption, $y(1+r)/r$. The fourth line neither the catastrophe nor the magic technology occur, so we have the continuation value $V(x_{t+1})$.

Rearranging terms on the right hand side gives

$$\begin{aligned} V(x_t) &= y + \frac{1}{1+r} \left\{ -\lambda(x_t)\, K + \mu\, y\, \frac{1+r}{r} \left[\theta\, \lambda(x_t) + 1 - \lambda(x_t)\right] \right\} \\ &\quad + \frac{1}{1+r}\, V(x_{t+1})\, (1-\mu)\, \left[\theta\, \lambda(x_t) + 1 - \lambda(x_t)\right] \\ &= y \left\{ 1 + \frac{\mu}{r} \left[\theta\, \lambda(x_t) + 1 - \lambda(x_t)\right] \right\} - \lambda(x_t)\, \frac{K}{1+r} \\ &\quad + \frac{1}{1+r}\, V(x_{t+1})\, (1-\mu)\, \left[\theta\, \lambda(x_t) + 1 - \lambda(x_t)\right] \end{aligned} \tag{3}$$

To simplify notation, define two functions

$$B(x) = y \left\{ 1 + \frac{\mu}{r} \left[1 - (1-\theta)\, \lambda(x) \right] \right\} - \lambda(x)\, \frac{K}{1+r}\,, \tag{4}$$

and

$$C(x) = \frac{1-\mu}{1+r} \left[1 - (1-\theta)\, \lambda(x) \right] . \tag{5}$$

The difference equation (3) can now be written as:

$$V(x_t) = B(x_t) + C(x_t)\; V(x_{t+1}) \tag{6}$$

or

$$V(x_{t+1}) = \frac{V(x_t) - B(x_t)}{C(x_t)}\,. \tag{7}$$

To solve this we need a way to fix $V(0)$. First we develop properties of all possible solution paths with different logically conceivable $V(0)$, and then show how the actual $V(0)$ can be determined.

Begin by comparing two solutions labeled by superscripts a and b. For any t, using (7) and $C(x_t) > 0$, we have that $V^a(x_t) > V^b(x_t)$ implies $V^a(x_{t+1}) > V^b(x_{t+1})$. Therefore if $V^a(0) > V^b(0)$, then $V^a(x_t) > V^b(x_t)$ for all t. That is, the family of solutions can be depicted graphically in (t, V) space (or equivalently, in (x, V) space; in fact by (1) t and x are effectively the same with just a change of scale) as a set of non-intersecting curves.

Next, observe that

$$0 < C(x) < 1 \quad \text{for all } x\,,$$

and

$$V(x_{t+1}) - V(x_t) = \frac{[1 - C(x_t)]\; V(x_t) - B(x_t)}{C(x_t)}\,, \tag{8}$$

so

$$V(x_{t+1}) > V(x_t) \quad \text{if and only if} \quad V(x_t) > \frac{B(x_t)}{1 - C(x_t)}$$

Define the function

$$Z(x) = \frac{B(x)}{1 - C(x)} \tag{9}$$

and observe from (4) and (5) that because $\lambda(x)$ is increasing, $B(x)$ and $C(x)$ are both decreasing, so $Z(x)$ is decreasing. Figure 1 shows this curve colored blue. (This figure uses numbers developed in Section 3, but think of it right now as merely depicting qualitative properties of the solution.)

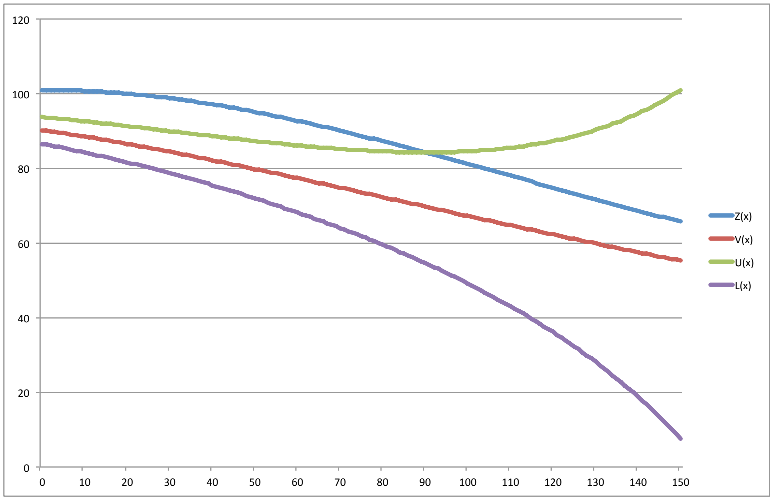

Figure 1: Solution paths

Next, we show that if $V(x_{t+1}) > V(x_t)$, then $V(x_{t+2}) - V(x_{t+1}) > V(x_{t+1}) - V(x_t)$. To prove this, begin by noting that $x_{t+1} > x_t$, and $B(x)$ and $C(x)$ are both decreasing functions. Inserting all this information in (8) yields the result.

Thus any solution path above the curve $Z(x)$ keeps rising at an increasing rate. However, there is an upper bound $\overline{V}$: $V(x)$ can never exceed what would happen if no catastrophe occurred. This is the discounted present value of the consumption stream y, i.e. $\overline{V} = y\,(1 + r)/r$. It is easy to check that this equals $Z(0)$. Any solution path that hits this bound then becomes infeasible; this is illustrated in Figure 1 by the green curve labeled $U(x)$.

Similarly, a lower bound $\underline{V}$ arises from the possibility that a catastrophe occurs every period until (if) the rescue technology arrives, i.e. as if $\lambda(x) \equiv 1$. Then the value is as if $B(x) = B(\infty)$ and $C(x) = C(\infty)$, i.e. (6) becomes

$$\underline{V} = B(\infty) + C(\infty)\ \underline{V}\,,$$

or

$$\underline{V} = \frac{B(\infty)}{1 - C(\infty)} = Z(\infty)\,.$$

In other words, the range of values of $Z(x)$ spans the range between the upper and lower bounds of values $V(x)$.

Another part of the family of solution paths can be worked backward from hitting the lower bound at different times. Figure 1 illustrates this by the purple curve labeled L(x). These curves all stay below the $Z(x)$ curve, so they are downward-sloping.

Separating all the curves that hit the upper or lower bounds and so become infeasible, there is just one that stays below the curve $Z(x)$ and above the lower bound, thus remaining feasible for all x. This is then the solution, and its $V(0)$ is the correct initial condition, completing the solution. Thus the solution has a familiar saddle-path structure. Figure 1 shows this as the red curve labeled V(x). The solution is downward-sloping throughout. This is intuitively reasonable: since x is a bad, a higher level of it entails a lower Bellman value $V(x)$.[4]

3 Numerical solution

We now find a numerical solution for a basic set of parameter values, and consider some parameter changes that represent possible policy interventions.

The basic parameters are as follows: α is fixed at 1 which amounts to choosing the unit of x equal to the amount of greenhouse gas pollution that would be added in a unit of time (e.g. year). In this baseline, x and t become effectively the same thing, allowing easier calculation and interpretation. However, we can model reduced greenhouse gas emissions by lowering $\alpha < 1$. We set the discount rate r equal to 1% per year, i.e. $r = 0.01$. This is similar to the inflation-adjusted rate of interest that has prevailed for several decades, and is a compromise between those who advocate near-zero discounting (e.g. Stern 2007) and those who favor 5% or higher rates (e.g. Weitzman 2007).

The arrival rate function $\lambda(x)$ should be very low for small x, and must stay below 1, so an S-shape is most plausible. We specify the following functional form:

$$\lambda(x) = \frac{(\gamma\ x)^2}{1 + (\gamma\ x)^2} \tag{10}$$

and take $\gamma = 0.0025$ as our base value. Figure 2 graphs this function for the first 150 years; it has a point of inflection further out, after about 230 years, when the arrival rate reaches 0.25.

More intuitive and instructive is the probability $\Pi(x)$ that at least one catastrophe has occurred by the time x is reached, i.e. one minus the probability that no catastrophe has occurred until then. Figure 3 graphs this. A useful benchmark for comparison is the year in which this probability reaches 0.9; for our basic parameter values the answer is 104. But with $\gamma = 0.0015$ instead, this would be after 146 years.

Next, we set the arrival rate of the technology that rescues mankind from the risk of catastrophes at $\mu = 0.005$, i.e. $\frac{1}{2}$% per year. We set $y = 1$, which can be thought of

[4]It is somewhat easier to derive and understand in continuous time. However, the discrete-time model is better for numerical solutions and therefore for developing intuition for effects of different policy options; therefore we have used it in the main text. But to give readers a somewhat easier grasp of the theory, in Appendix A.1 we construct a continuous-time variant.

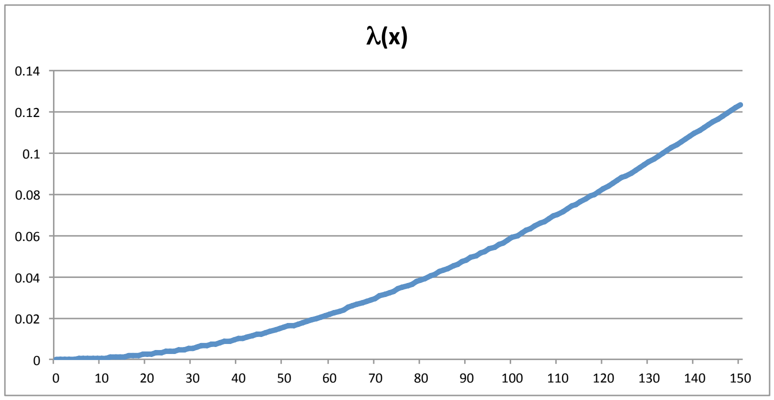

Figure 2: Catastrophe arrival rate function

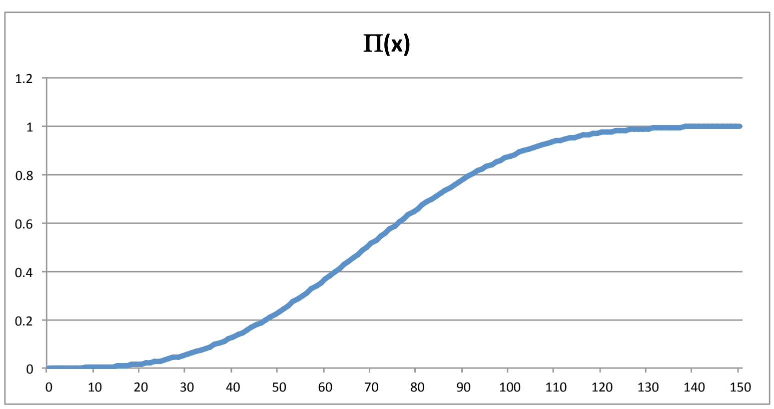

Figure 3: Cumulative catastrophe probability function

as a choice of the units of output in the absence of catastrophes; $K = 1$, which can be a large loss of capital, land etc., or a large disutility caused by the loss; and $\theta = 0.95$, the shrinkage factor of future possibilities that results from each catastrophe.

The calculation was done using an Excel spreadsheet.[5] The procedure was as follows. First we calculated $Z(x)$ for 300 years, and the lower bound $\underline{V}$. We set

$$V(300) = \tfrac{1}{2}\left[\, Z(300) + \underline{V}\,\right],$$

thereby choosing a path that remained feasible for 300 years without hitting either bound. Then we calculated $V(x)$ for earlier years backward using (6), and retained the first 150 years of this path as our solution. Because of the saddle-path instability property, our choice of $V(300)$ makes very little difference to the values from $V(0)$ to $V(150)$. The resulting solution, together with two unstable paths to either side of it, is what we have graphed for Figure 1 in the previous section.

Our solution has $Z(0) = 101$ and $V(0) = 90.23$. That is, if the numbers are interpreted as present discounted money-equivalent magnitudes, about 11% of the total productive potential of this economy is lost because of the risk of the catastrophes of the magnitude that we are contemplating. Below, we will see how sensitive this to parameter changes as a means of thinking about policy consequences of different strategies to mitigate this loss.

4 Policy experiments

The framework that we have put forward is useful for contemplating the value of policy interventions. Our "policy experiments" involve changing one parameter at a time, keeping all the others at their values specified above. Specifically, we will consider three main policy experiments in our framework: varying α, γ and μ. We then ask what reduction in y, which we will interpret as real consumption per head, we would be willing to accept[6] to secure a desirable parameter shift. We will not be specific about the forms that policies take. But broadly we have in mind those which increase investment in research and development or increase the cost of carbon-intensive forms of consumption.[7]

In our model, a catastrophe leads to falls in real consumption represented by K and θ.[8] There is a range of possible reasons why climate change will hit real consumption in future (see, for example, Stern, 2007). That said, providing reliable figures for the size of these damages has proved to be challenging given the uncertainties involved.

[5] The Excel file is available from the authors upon request, for readers to carry out their own experiments.

[6] In economics jargon this is the "compensating variation" for the parameter change.

[7] Nordhaus (2008) focused on policies which raise the price of carbon using an integrated model of climate dynamics and the economy. He advocatesd an optimal carbon tax in 2005 of $27 per metric ton in 2005 prices.

[8] More generally, though one should really think in terms of flows of utility rather than just consumption. Authors such as Barro (2015) have emphasized how risk and/or uncertainty aversion can increase the utility loss associated with a given profile of consumption. Many models of uncertainty aversion suggest looking at maximin criteria providing an even greater focus on worst case scenarios.

For example, future shocks to the climate are likely to result in flooding which have negative effects on productivity which result in lower consumption. But it is difficult to be precise about where flooding will occur and how effectively it can be mitigated. Another possibility is that climate change will necessitate greater spending on directly unproductive activities such as fighting territorial battles for land and resources or increased security costs due to mass migration. The costs of this will depend on how well political institutions function when put under stress and the efficacy of global cooperation.

We will compute the compensating change for three values of $\theta \in \{0.90, 0.95, 0.99\}$ and $K \in \{0.2, 1, 5\}$. These are broadly in line with standard IPCC estimates for output loss based on the status quo; they estimate that a 4^oC increase in global temperatures would lead to a $1\% - 5\%$ fall in future GDP[9] which is consistent with the range of θ that we consider. The case of $K = 0.2$ is similar to the kinds of consumption shocks considered in Barro (2015). While considering $K = 5$ may seem large, note that the precise way in which such a loss is distributed through time does not matter for the calculation, only the capital value of the loss. Over an infinite horizon with a 1% discount rate, $K = 5$ is equivalent to a per-period fall in consumption at *every* future date of only 5%.[10]

Changing the Level of Emissions Our first policy experiment supposes that α is reduced from from 1 to 0.7 (so a 30% reduction in the rate of greenhouse emissions). This is the order of magnitude of reductions in carbon emissions that have been contemplated in international agreements such as the Kyoto Protocol. A number of countries have responded to this with binding targets on carbon emissions backed up by regulations and tax incentives. Increased use of renewables and nuclear power in electricity generation form a key part of this strategy.

This change in α increases the year in which the probability of a catastrophe occurring hits 90% from 104 to 132 and $V(0)$ goes from 90.23 to 94.22. To bring $V(0)$ back to 90.22, y should be reduced from 1 to 0.958; hence, we should be willing to sacrifice 4.2% of annual consumption to bring about this reduction in emissions. Table 1 shows how this depends on θ and K.

The Table illustrates how this depends on the magnitude of the catastrophe being contemplated. In the most extreme case where $K = 5$ and $\theta = 0.9$, then the willingness to pay is 7.6% of real consumption which is quite substantial.

Varying Catastrophic Risk We now consider an increase in the risk of a catastrophe where we lower γ from 0.0025 to 0.002. This postpones the year in which the probability of at least one catastrophe hits 90% by 17 years, from 104 years to 121

[9]See Nordhaus (2008) for discussion of this.

[10]One feature of the model which merits mention although is not a feature of the tables below is that there can be a range of θ close to one where the loss is not monotonic in θ. This feature of the model is due to the possibility of having $K > 0$. Mathematically, this is a result of the fact that to maintain stationarity in the recursive relation that generates the equation for $V(x)$, it is necessary to have the same factor of proportionality θ for the decline of the size of catastrophes with respect to the capital loss K and the GDP reduction in y. Then the relative importance of the two as θ changes depends on the discount rate.

		K		
		0.2	1	5
Θ	0.99	1.3%	2.1%	6.3%
	0.95	3.6%	4.2%	6.9%
	0.90	5.2%	5.6%	7.6%

Table 1: Reduced Emissions

years. With $K = 1$ and $\theta = 0.95$ (our middle case), this increases $V(0)$ from 90.23 to 92.89. Lowering y to 0.972 has a similar impact on $V(0)$. Hence, we should be willing to spend $1 - 0.972 = 0.028$, or 2.8% of annual real consumption to achieve this reduction in risk. Table 2 shows how this number varies for a range of values of the loss parameters. It shows that this behaves predictably with the compensating variation increasing in θ and K.

		K		
		0.2	1	5
Θ	0.99	0.9%	1.4%	4.3%
	0.95	2.5%	2.8%	4.5%
	0.90	3.5%	3.7%	5.2%

Table 2: Reduced Catastrophe Risk

It is clear from this Table how the size of the catastrophe affects the willingness to sacrifice consumption to reduce risk. Even in the least pessimistic case, devoting a little less than 1% of real consumption is justified while this increases to more than 5% for larger catastrophes.

Increasing the Probability of a Technological Solution We now suppose now that μ is doubled to 0.01, i.e. there is now a 1% chance of finding a solution that eliminates the problem. Then in the case where $K = 1$ and $\theta = 0.95$, $V(0)$ rises to

95.08. Lowering y to compensate, the $V(0)$ stays above the original value of 90.23 so long as $y > 0.949$ suggesting a willingness to invest which reduces real consumption by 5.1% if it could double the probability of finding a technology which would fix the problem. Table 3 illustrates how this calculation depends on the loss parameters. The dependence on θ and K is once again predictable. Here, the compensating variation rises to 9.1% for the worst case that we consider.

		K		
		0.2	1	5
ϴ	0.99	1.6%	2.7%	8.1%
	0.95	4.4%	5.1%	8.3%
	0.90	6.2%	6.7%	9.1%

Table 3: Increased Probability of a Technological Solution

Summary These results are purely illustrative and the spread sheet which produced them is available on request. Overall, the results illustrate that focusing on tail risk can motivate a substantial "willingness to pay" to reduce risk even with a fairly standard discount rate and with economic loss parameters that are quite modest. More generally, our approach provides a quantitative "ready-reckoner" to explore even more extreme threats to future living standards.

5 Concluding comments

This analysis offers a useful starting point for studying the effects of rare large disasters that may result from climate change, and the policies for deterring and countering these effects. The model can be augmented, extended, and altered in various ways to include related effects and further details; these are topics for our future research. Here we outline some of these.

We have captured the idea of rare but large disasters not by fattening the tails of an ordinary normal distribution, for example by changing it into another continuous distribution like the Pareto-Levy or Cauchy distribution, but as a Poisson process, which is nothing but tail. This can be combined with some component of ordinary normal risk. For example, our y can be replaced by a function $y(x)$, which captures the decreasing certainty equivalent of an increasing normal risk as greenhouse gases accumulate.

We have not considered optimal policies. This is what kept the recursion equation (6) linear. An optimized right hand side of (2) would have been a non-linear function of $V(x_{t+1})$, and characterization and numerical solution of the resulting Bellman equation would have been much harder. The compensating advantage we gained was that we could compute policy tradeoffs (what fraction of annual consumption we should be willing to sacrifice to obtain various reductions in the probabilities of rare disasters) in a simple and intuitive way. But a more ambitious analysis of optimal policies remains an interesting topic for future research.

We have considered the "demand side" of policies, namely how much current consumption we should be willing to invest in order to reduce the catastrophic risks by a specified amount. This can be combined with the "supply side," namely the marginal productivity of various investments in deterrence and countering of these risks, to complete the analysis of which investments should be undertaken.

Finally, we have considered the "demand side" as if there were one central decision-maker whose objective was our function $V(x)$. In reality the issue of climate change involves the interaction of many countries that stand to lose or gain from climate change to different degrees, and have partly conflicting beliefs and objectives about the effects. Even within one country, there are political differences about the issue. The analysis has to be extended to encompass the political economy, not just the technological and aggregate economic aspects, if it is to be realistic.

References

Barro, Robert J. 2015. Environmental protection, rare disasters and discount rates. *Economica* 82(1): 1–23.

Brito, Dagobert L. and Michael D. Intriligator. 1987. Stock externalities, Pigovian taxation and dynamic stability. *Journal of Public Economics* 33(1): 59–72.

Dietz, Simon. 2011. High impact, low probability? An empirical analysis of risk in the economics of climate change. *Climate Change* 108: 519–541.

Kolstad, Charles D. 1996. Fundamental irreversibilities in stock externalities. *Journal of Public Economics* 60(2): 221–233.

Martin, Ian W.R. and Robert S. Pindyck. 2015. Averting catastrophes: The strange economics of Scylla and Charybdis. *American Economic Review* 105(10): 2947–2985.

Millner, Antony. 2013. On welfare frameworks and catastrophic climate risks. *Journal of Environmental Economics and Management* 65: 310–325.

Nordhaus, William. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven, CT: Yale University Press.

Stern, Nicholas H. 2007. *The Economics of Climate Change: the Stern Review.* Cambridge, UK: Cambridge University Press.

Weitzman, Martin L. 2007. A review of the 'Stern Review on the Economics of Climate Change'. *Journal of Economic Literature* 45(3): 703–724.

——— . 2011. Fat-Tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy* 5(2): 275–292.

——— . 2014. Fat tails and the social cost of carbon. *American Economic Review: Papers and Proceedings* 104(5): 544–546.

Appendix

A.1 Continuous-time Model

Here we develop a continuous-time model similar to the discrete-time model in the text, adhering to the same notation. The dynamics of the state variable x are now given by

$$\frac{dx}{dt} = \alpha\,. \tag{A.1}$$

A catastrophe can strike as a Poisson process with arrival rate $\lambda(x)$, an increasing function, so the probability that a catastrophe occurs in an infinitesimal time interval dt at state x is $\lambda(x)\,dt$. Before a catastrophe strikes, there is a flow of benefits, for example real consumption, at rate y. The catastrophe causes a lump sum cost or wealth destruction K, and after that the economic activity continues with consumption flow and the lump sum costs of any future catastrophes reduced to a fraction θ of its former level. All benefits and costs are discounted at the rate r.

There is another Poisson process with arrival rate μ representing technical change such that if this event occurs, the risk of the catastrophe will disappear, i.e. $\lambda(x)$ will drop to zero.

Various policies to cope with this risk can be contemplated. Some examples: (a) Mitigation: by paying a flow cost m, which may be related to real consumption, the arrival rate of the catastrophe can be reduced to $\lambda(x,m)$, where the function is decreasing in m. For example, in a linear specification $\lambda(x,m) = x/m$, an increase in the parameter m shifts the arrival rate function downward. (b) Innovation: by paying a flow cost n, the arrival rate of the technical change process can be increased to $\mu(n)$. (c) Sudden stop: by paying a lump sum cost L, the catastrophe process can be stopped, i.e. $\lambda(x)$ reduced to zero immediately. Other examples can also be imagined and analyzed using similar methods.

Let $V(x)$ denote the Bellman value function starting at x. This has the recursion relation:

$$\begin{aligned}
V(x) &= y\,dt + \Big\{ \lambda(x)\,dt\;[-K + \theta\,V(x+\alpha\,dt)\,] + \mu\,dt\;(y/r) \\
&\qquad\qquad + [\,1 - (\lambda(x)+\mu)\,dt\,]\;V(x+\alpha\,dt) \Big\}\, e^{-r dt} \\
&= y\,dt + \Big\{ \lambda(x)\,dt\,[-K + \theta\,V(x) + \theta\,V'(x)\,\alpha\,dt] + \mu\,dt\;(y/r) \\
&\qquad\qquad + [\,1 - (\lambda(x)+\mu)\,dt\,]\;[V(x) + V'(x)\;\alpha\,dt] \Big\}\;(1 - r\,dt) \\
&= y\,dt + \lambda(x)\,dt\,[-K + \theta\,V(x)\,] + \mu\,dt\;\frac{y}{r} \\
&\qquad\qquad + V(x) - r\,V(x)\,dt - [\lambda(x)+\mu]\;V(x)\,dt + V'(x)\;\alpha\,dt \\
&= V(x) + \Big\{ y - K\,\lambda(x) + \theta\,\lambda(x)V(x) + \mu\,\frac{y}{r} - r\,V(x) - [\lambda(x)+\mu]\;V(x) + \alpha\,V'(x) \Big\}\;dt \\
&= V(x) + \Big\{ y\,\frac{\mu + r}{r} - K\,\lambda(x) + \alpha\,V'(x) - [r + (1-\theta)\,\lambda(x) + \mu]\;V(x) \Big\}\;dt\,.
\end{aligned}$$

In the right-hand side of the first line, the first term is the consumption flow, and the rest are the various continuation values starting time dt later and hence discounted by the factor e^{-rdt}. Of these three continuation values, each with its appropriate probability, the first captures the catastrophe event and its consequences, namely the fixed cost and the reduced continuation value $\theta\, V$ starting from the slightly increased state $x + \alpha\, dt$; the second captures the technical change event (after which the flow y continues for ever with the continuation capitalized value y/r);[11] and the third is the case where neither of the two Poisson events occurs in this infinitesimal time interval, so the continuation is given by the same value function V evaluated at the slightly increased $x + \alpha\, dt$. The subsequent lines are successive simplifications, at each step neglecting terms of order $(dt)^2$ and higher. Note that in continuous time the probability of joint arrival of a catastrophe, and of the technology that prevents subsequent catastrophes, in the same time interval dt is of order $(dt)^2$; that is why this derivation has only three eventualities instead of the four in the corresponding discrete time equation (2) in the text.

Now we can cancel $V(x)$ from both sides to get

$$y\, \frac{\mu + r}{r} - K\, \lambda(x) + \alpha\, V'(x) - [r + (1-\theta)\, \lambda(x) + \mu]\ V(x) = 0\,,$$

which yields the following differential equation for the dynamics of $V(x)$:

$$\alpha\, V'(x) = [r + \mu + (1-\theta)\, \lambda(x)]\ V(x) - \left[y\, \frac{\mu + r}{r} - K\, \lambda(x) \right] \tag{A.2}$$

We can also establish upper and lower bounds for $V(x)$. The upper bound occurs if the catastrophe never happens; it is the capitalized value of the consumption flow, namely y/r. The lower bound occurs if catastrophes happen in rapid succession with only infinitesimal consumption flows in between; the outcome is

$$-K - \theta\, K - \theta^2\, K - \ldots = -K/(1-\theta)\,.$$

Neither of these bounds is actually attained so long as $0 < \lambda(x) < \infty$; therefore

$$y/r > V(x) > -K/(1-\theta)\,. \tag{A.3}$$

The differential equation can be solved in closed form, but the solution is better understood graphically. In (x, V) space, we can show the solution curves corresponding to different potential values of $V(0)$, and then pick the right initial condition. Since x increases linearly with time, the trajectories can also be interpreted as the time paths of $V(x(t))$.

Begin by drawing the locus defined by $V'(x) = 0$, i.e.

$$V(x) = \frac{y\,(\mu + r)/r - K\, \lambda(x)}{r + \mu + (1-\theta)\, \lambda(x)}\ . \tag{A.4}$$

Since (A.2) shows that $V'(x)$ is increasing in $V(x)$ for given x, we have $V'(x) > 0$ above this locus and < 0 below it. The solution trajectories that cross this locus do so horizontally.

[11] If in the absence of a catastrophe GDP would grow exponentially at rate g, then this capitalized value would be $y/(r-g)$.

Under the reasonable assumption that $\lambda(0) = 0$, this locus has the intercept (y/r) on the vertical axis. Next, since $\lambda(x)$ increases as x increases, on the right hand side of (A.4) the numerator decreases and the denominator increases; therefore the curve is decreasing. If $\lambda(x) \to \infty$ as $x \to \infty$, this locus asymptotes to $-K/(1-\theta)$. Thus the locus exactly spans the bounds of the solution. Figure 1 shows the locus, the blue curve labelled Z(x).

Now consider the curvature of the solutions of (A.2) starting from different initial $V(0)$. Differentiating the equation gives

$$\alpha\, V''(x) = [\, r + \mu + (1-\theta)\, \lambda(x) \,]\; V'(x) + [\, (1-\theta)\, V(x) + K \,]\; \lambda'(x)\,. \tag{A.5}$$

From (A.3) we have $(1-\theta)\, V(x) + K > 0$. Therefore $V'(x) > 0$ implies $V''(x) > 0$, i.e. the solution trajectories above the curve Z(x) are convex. Below that locus but close to it, $V'(x)$ is negative but small in magnitude. Therefore the second term on the right hand side of (A.5), which is positive, dominates, and $V''(x)$ is positive. On the other hand, close to the lower bound, the second term is close to zero, and the horizontal line at $V = -K/(1-\theta)$ is below the locus Z(x) so $V'(x) < 0$ there. Then the first term on the right hand side of (A.5) dominates and $V''(x) < 0$, so the solution trajectories are concave.

Combining all this information we have the following picture. From any initial $V(0)$ in the range $(-K/(1-\theta)\,,\, y/r)$ the trajectory is initially downward-sloping. If $V(0)$ is close to the lower bound, the trajectory will stay downward-sloping, eventually turn concave, and hit the lower bound for finite x, becoming infeasible. If $V(0)$ is close to the upper bound, the trajectory will be initially downward-sloping, but it will turn convex, cross the locus Z(x), and then increase and remain convex, hitting the upper bound for a finite x and becoming infeasible. Figure 1 shows one such trajectory of each kind.

Between these two types of infeasible trajectories, there is just one, the red curve in Figure 1 labeled V(x), which slopes downward and remains squeezed between the locus Z(x) and the horizontal line at the lower bound. This "saddle-path" is the only one that can remain feasible for all time. Its $V(0)$ is then the correct initial condition.