



Environmental catastrophes and mitigation policies in a multiregion world

Timothy Besley^{a,1} and Avinash Dixit^{b,1,2}

^aDepartment of Economics, London School of Economics and Political Science, London WC2A 2AE, United Kingdom; and ^bDepartment of Economics, Princeton University, Princeton, NJ 08544-1021

Edited by Geoffrey M. Heal, Columbia University, New York, NY, and approved August 8, 2018 (received for review February 26, 2018)

In this paper we present a simple model for assessing the willingness to pay for reductions in the risk associated with catastrophic climate change. The model is extremely tractable and applies to a multiregion world but with global externalities and has five key features: (i) Neither the occurrence nor the costs of a catastrophic event in any one year are precisely predictable; (ii) the probability of a catastrophe occurring in any one year increases as the levels of greenhouse gases in the atmosphere increase; (iii) greenhouse gases are a worldwide public bad with emissions from any one country or region increasing the risks for all; (iv) there is two-sided irreversibility; if nothing is done and the problem proves serious, the climate, economic activity, and human life will suffer permanent damage, but if we spend large sums on countermeasures and the problem turns out to be minor or even nonexistent, we will have wasted resources unnecessarily; and (v) technological progress may yield partial or even complete solutions. The framework that we propose can give a sense of the quantitative significance of mitigation strategies. We illustrate these for a core set of parameter values.

catastrophic climate risk | global stock externality | climate change mitigation

It is a truth almost universally acknowledged that greenhouse gas (GHG) accumulation is contributing to climate change and increasing the risk of catastrophes such as cyclones, floods, droughts, and wildfires. Several aspects of this are important: (i) Although the broad mechanism by which climate change occurs is well established, neither the occurrence nor the costs of a catastrophic event in any one year are precisely predictable. There is much uncertainty, and the policy issue is mitigation of large risks. (ii) The probability of a catastrophe occurring in any one year increases as the levels of GHG in the atmosphere increase. (iii) GHGs are a worldwide public bad; emissions from any one country or region increase the risks for all. (iv) There is two-sided irreversibility of policies. If we do nothing and the problem proves serious, the climate, economic activity, and human life will suffer permanent damage, but if we spend large sums on countermeasures and the problem turns out to be minor or even nonexistent, we will have wasted resources unnecessarily. (v) Technological progress may yield partial or even complete solutions such as rapid and efficient carbon removal, injecting sulfur particles into the upper atmosphere, or some other form of geoengineering.

This paper presents a simple, user-friendly model that puts these elements together and which can be understood intuitively. Previous analytically tractable models have tended to have only one or two of the features highlighted above. More complex approaches, such as those embedded in integrated assessment models, do include all of these features but this comes at the cost of complexity, making it difficult to develop an intuitive understanding of findings.

Our approach is not meant to substitute for detailed modeling but rather is a way of building an intuitive understanding of some issues and creating a ready reckoner to inform debate. It permits calculation of the expected economic costs of such environmen-

tal catastrophes and yields upper bounds on the sums we should be willing to spend on countermeasures.

The tractability of the model makes it straightforward to apply to a multiregion world but with global externalities and hence to look at how willingness to pay in one region varies with mitigation measures taken in other regions. This is important since most examples of policies to combat climate change require international cooperation. For example, the United States pulling out of the Paris climate accord could affect the willingness of other countries to make sacrifices to combat climate change. We can also give a back-of-the-envelope sense of the quantitative significance of such effects.

We propose a simple way of looking at things by calculating the cost it would be worth paying to achieve specified target levels of mitigation. Even being quite conservative, we find numbers upward of 1% of gross domestic product (GDP). This can justify substantial expenditures in pursuit of those targets; for example, for the United States this amounts to spending about \$190 billion every year, far more than anyone has proposed for such policies.

Our numerical results, which can be performed using an Excel file ([Dataset S1](#)), show that the willingness of one region to spend resources to reduce GHG emissions is higher if other regions are also contributing their efforts; in game-theoretic terminology, different regions' expenditures are "strategic complements." This is borne out in reality by the fact that in international environmental forums countries often make "conditional commitments" rather than unconditional ones. The simplicity of our model makes the various effects and interactions directly interpretable.

The remainder of this paper is organized as follows. In *Related Literature*, we briefly discuss some of the existing literature. In *A Canonical Model*, we put forward a simple and general approach where the formulas for the compensating and equivalent variations from interventions are derived. We then extend the model to a multiregion setting where there are externalities across regions. *Numerical Solution* develops the model's implications which we implement and assess and results are in *Results*. Some suggestions are proposed in *Ideas for Future Work*. Some brief reflections on the value and limitations of the approach are in *Concluding Comments*.

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "Economics, Environment, and Sustainable Development," held January 17–18, 2018, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, CA. The complete program and video recordings of most presentations are available on the NAS website at www.nasonline.org/economics-environment-and.

Author contributions: T.B. and A.D. performed research and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Published under the [PNAS license](#).

¹T.B. and A.D. contributed equally to this work.

²To whom correspondence should be addressed. Email: dixitak@princeton.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1802864115/-DCSupplemental.

Published online September 25, 2018.

Related Literature

This paper is informed by the extensive body of existing research on the relationship between climate policy and the economy in dynamic settings. The interested reader will find an excellent overview of the voluminous literature in Hassler, Krusell, and Nycander (1).

A key feature of climate change is the nature of the “stock externality” whose effects accumulate over time. This is a feature of the integrated assessment models of Manne and Richels (2) and Nordhaus (3). The two-sided uncertainty which creates embedded options in both the decisions to act and to wait also features and is a feature of our model and is studied in Kolstad (4).

There are many models which study issues related to those in this paper using dynamic general equilibrium models to assess the impact of policy alternatives. For example, Golosov et al. (5) study optimal fuel taxes and Gerlagh and Liski (6) study the role of commitment assumptions. Traeger (7) shares the ambition of this paper to develop a simple analytically tractable approach to climate change uncertainty, aiming to close the gap between complex numeric models used in policy advising and stylized models built for analytic insight. He posits a process by which temperature increases and uses a damage function based on the well-known Dynamic Integrated Climate-Economy (DICE) model’s specification. That paper builds on an intuitive understanding of the framework using computer simulations. Newbold et al. (8) provide a tool to help analysts and decision makers quickly explore the implications of various modeling assumptions with a focus on the social cost of carbon. Using a traditional growth model, they posit a damage function for GDP based on temperature increases due to accumulation of GHGs and conduct a sensitivity analysis to reflect parameter uncertainty, using a simulation model. Their framework does not have losses due to catastrophic events as such and also does not have the possibility of a savior technology.

Our paper follows in the footsteps of a range of papers that have modeled climate damages as stochastic catastrophes that depend on the stock of CO₂ and seek to provide simple analytical approaches to illuminate policy options. Gerlagh and Liski (6) model an event that causes damage as a Poisson process. However, unlike in our paper, once it occurs, the damage is a constant fraction of output and a constant subtracted from utility; we allow for multiple events that may or may not occur in successive years, such as hurricanes, droughts, floods, wildfires, etc. Lemoine and Traeger (9, 10) explore climate tipping points. They model these as a form of model discontinuity where the dynamics of the system shift abruptly, and global warming reduces the Earth’s ability to remove CO₂ from the atmosphere. The work of van den Bijgaart et al. (11) also explores the implications of a dynamic equilibrium model and simulates alternative scenarios to estimate the social cost of carbon. They do not have the possibility of a technological breakthrough to mitigate climate change effects. Van der Ploeg and de Zeeuw (12) model two effects of GHG accumulation: technical “regress” where the production function has a multiplicative factor which follows Golosov et al. (5). (A questionable feature of this specification is that the effect is biggest as it starts and flattens thereafter.) They also have a single “tipping” catastrophe at an unknown date whose hazard rate increases with GHG accumulation and which will cause a gradually increasing proportional drop in output. Tsur and Zemel (13) study a different source of uncertainty due to lack of knowledge. There is a critical level of the GHG stock which, if hit, will with certainty trigger catastrophic loss. The observation that the catastrophe has not yet occurred allows policy makers to update (truncate) their prior about this critical level.

We follow the literature in including the possibility of substantial technological progress reflecting advances in clean energy technology. This is endogenized in Acemoglu et al. (14) who model gradual endogenous technical change as is standard in economic growth models. We posit the possibility that technological progress could be a discontinuous “breakthrough” of the kind that is posited by climate optimists.

In common with Hassler and Krusell (15) our model has many regions. They develop a stochastic general equilibrium model with multiple regions and discuss the spillovers of policies such as carbon taxes. In their model, total factor productivity is a decreasing function of GHG stocks multiplied by a stochastic shock. Therefore, GHG accumulation actually reduces uncertainty in productivity equiproportionately with its level. This differs from our focus on increasing catastrophic risks that result from these accumulations. They also do not allow for the possibility that technological progress may allow the problem to be avoided or solved much more cheaply in the future.

A Canonical Model

We begin with a single-region model and then extend it to multiple regions or countries.

Core Single-Region Model. Let x_t denote the logarithm of the cumulated GHG level in the atmosphere in year t . This is our state variable; its dynamics are explained below. The expected GDP is denoted by y_t . This can be a decreasing function $y_t(x_t)$, interpreted as the certainty equivalent of some normal (non-catastrophic) uncertainty caused by GHG accumulation, for example some loss of efficiency of production processes.

The loss caused by a catastrophe event (conditional on one occurring in year t) is denoted by K_t , which can be an increasing function $K_t(x_t)$. It is interpreted as a comprehensive certainty-equivalent measure. For example, if the catastrophe lowers the path of GDP from its status quo for several years, K_t includes the discounted present value of the GDP gap. It is also intended to include the monetary equivalent of human costs such as loss of life and dislocation. Not surprisingly, this is a key parameter in our analysis.

The catastrophe is modeled as a Poisson process with arrival rate $\lambda(x_t)$, an increasing function. This is the simplest way to model fat-tailed risk that rises with GHG accumulation. In fact, with this formulation everything is in the tail. We interpret this risk as the addition to any environmental risk that may exist independently of climate change [Stott et al. (16) and van Oldenborgh et al. (17)].

The logarithm of the stock of GHGs, x , ranges from $-\infty$ to ∞ , and $\lambda(x_t)$ is bounded between 0 and 1. For GHG accumulations close to their current levels, which we normalize at 1 so the initial x is zero, we expect the catastrophic risk to be extremely small and to rise only slowly. Eventually it will rise more rapidly, but since the arrival probability is bounded above by 1, the function must eventually become concave. The obvious form for $\lambda(x_t)$ with these properties is a sigmoid; indeed that is the function used for many specifications of state-dependent increasing probabilities. [For example, Lin et al. (18) project that the arrival of storms of the magnitude of Hurricane Sandy are likely to increase by a factor as large as 17 in the period 2000–2100. This will severely threaten the flood defenses of New York City.] Below, we specify it parametrically for numerical calculations, but for the moment it is kept general.

The dynamics of carbon accumulation have been found to follow multiple paths [Inman (19)]. We adopt the formulation in Hassler and Krusell (15). About 60% of the emissions dissipate very quickly, so we omit them from consideration. About 20%, i.e., half of the nontransient part, stay forever. The remaining

stock dissipates with a depreciation rate of around 2.8%/year. We use these numbers and base values for our model and calculations. Thus, if z_t denotes the emission flow in period t , and the fraction ϵ is permanent, the permanent stock P_t is

$$P_t = \sum_{\tau=0}^t \epsilon z_{\tau}.$$

Writing δ for the dissipation rate of the remaining fraction $(1 - \epsilon)$, the dissipating stock, S_t , grows as

$$S_t = (1 - \epsilon) z_t + (1 - \delta) S_{t-1}$$

with $S_0 = 0$. (Taking the stock at $t = 0$ to be zero is just a normalization, as this level gets incorporated into the parameters of the catastrophe hazard rate function $\lambda(x)$ defined below.) Then

$$x_t = \ln(P_t + S_t).$$

The emission flows z_t can in general have any specification; we expect this to be an increasing function of the GDP, i.e., y_t . In our numerical calculations we make specific assumptions; these will be stated at that point.

A second and independent Poisson process represents a technological solution to the whole climate change problem. [For example, Barrett (20) discusses the possibility of geoengineering solutions.] Its arrival rate is denoted by $\mu(x_t)$. This can be an increasing function of the GHG level—as the problem worsens, more resources are devoted to research and development (R&D)—or a decreasing function—as the problem worsens, more resources are needed to solve it, but because GDP falls as GHG accumulation lowers productivity, fewer resources are available for R&D. If the technological solution arrives in year t , thereafter no catastrophes will occur. We assume that GDP will then go on growing at rate g :

$$y_{t+\tau} = y_t (1 + g)^{\tau}.$$

So in the absence of a catastrophe, or after a miracle rescue technology has appeared, the value of the economy at period t will be

$$V^*(y_t) = \sum_{\tau=0}^{\infty} y_t \left[\frac{1+g}{1+r} \right]^{\tau} = y_t \left[\frac{1+r}{r-g} \right]$$

independently of x , and $r > g$ is the discount rate. [$r > g$ is the standard dynamic efficiency or convergence condition in growth models; see Dixit (ref. 21, pp. 59 and 109). For its empirical relevance, see Piketty (22).]

At $t = 0$, before the miracle rescue technology has appeared, the expected net present value (NPV) of the economy—NPV of the GDP minus the expected discounted costs of catastrophes—can be shown to be (see *SI Appendix, section 1* for the derivation)

$$V(x_0) = \sum_{t=0}^{\infty} D_t [b_t(x_t) y_t(x_t) - \lambda_t(x_t) K_t(x_t)], \quad [1]$$

where

$$D(0) = 1, \quad D_{T+1} = \prod_{i=0}^T \frac{1 - \mu(x_i)}{1 + r}$$

and

$$b_t(x_t) = \frac{r - g + (1 + g) \mu(x_t)}{r - g}.$$

Eq. 1 has an intuitive interpretation with $\lambda_t(x_t) K_t(x_t)$ being deducted from the maximum payoff in each period. [Note that only the expected loss $\lambda_t(x_t) K_t(x_t)$ from a catastrophe matters, not the probability and the loss separately. So we have some

freedom in what follows in specifying the x dependence of the two.]

We use the standard economic measures of willingness to pay for a change, namely compensating and equivalent variations. The former asks how much GDP a society would be willing to sacrifice to make it equally well off before and after a reduction in catastrophic risk. The latter asks the question in reverse: What increase in GDP would be needed without the reduction in catastrophic risk to make the value the same. The first measure is the willingness to pay for the change, and the second is the GDP increase that would be an acceptable alternative to the change.

To capture these ideas formally and compute them numerically, we suppose some parameters change with new values having a “ \sim ” above them; i.e., D_t is replaced by \widetilde{D}_t , the function $b_t(x_t)$ by $\widetilde{b}_t(x_t)$, etc., and the value $V(x_0)$ by $\widetilde{V}(x_0)$. For the compensating variation we ask what fractional decrease θ_{CV} in GDP at all times at the new parameters would yield the same value as before. That is, we want to find θ_{CV} such that

$$V(x_0) = \sum_{t=0}^{\infty} \widetilde{D}_t [\widetilde{b}_t(x_t) \{ (1 - \theta_{CV}) y_t(x_t) \} - \widetilde{\lambda}_t(x_t) K_t(x_t)],$$

where the tildes show that the right-hand side is evaluated at the new parameters. And for the equivalent variation we ask what fractional increase in GDP at the old parameters would yield the new value; i.e., we want to find θ_{EV} such that

$$\widetilde{V}(x_0) = \sum_{t=0}^{\infty} D_t [b_t(x_t) \{ (1 + \theta_{EV}) y_t(x_t) \} - \lambda_t(x_t) K_t(x_t)],$$

where the absence of tildes on the right-hand side means that it is evaluated at the old parameters. Simple algebra, detailed in *SI Appendix, section 3*, shows that

$$\theta_{CV} = \frac{\widetilde{V}(x_0) - V(x_0)}{\sum_{t=0}^{\infty} \widetilde{D}_t \widetilde{b}_t(x_t) y_t(x_t)} \quad [2]$$

and

$$\theta_{EV} = \frac{\widetilde{V}(x_0) - V(x_0)}{\sum_{t=0}^{\infty} D_t b_t(x_t) y_t(x_t)}. \quad [3]$$

Although the formulas for the two variations look very similar, they are not: The denominators on the right-hand sides are evaluated for different parameters, the new ones in Eq. 2 and the old ones in Eq. 3, as shown by the presence or absence of tildes. However, in the numerical example that we solve, these variations turn out to be equal.

Multiple Regions. Now consider a world with many regions indexed by superscript i , where the externality from emissions and proneness to catastrophe is global. We can think of regions as either countries or groups of countries.

Writing z_t^i for the emission flows in region i , the permanent and dissipating components P_t and S_t of the global stock now follow

$$P_t = \sum_{\tau=0}^t \sum_i \epsilon z_{\tau}^i$$

and

$$S_t = (1 - \epsilon) \sum_i z_t^i + (1 - \delta) S_{t-1}$$

with $S_0 = 0$, and then the state variable, namely log aggregate log-GHG accumulation X_t , is

$$X_t = \ln(P_t + S_t).$$

Thus, emissions are a global public bad. To reflect this, region i 's GDP is denoted by $y^i(X)$ and the cost of a catastrophe in region i is $K^i(X)$. The growth rate of these, g^i , can also be region specific. The arrival rate of the catastrophe process is $\lambda^i(X)$ for region i ; it can differ across the regions because although they are all affected by the worldwide X , their probabilities and costs can depend on whether they are in a hurricane-prone area or a flood zone, etc. The technological solution, if it materializes, is highly likely to be global. This implies that the arrival rate function for that process should be the same for all regions i and hence denoted by $\mu(X)$. However, our framework can cope straightforwardly with the possibility of more local solutions, such as levees or better rain capture to cope with droughts, in which case there would be separate functions $\mu^i(X)$.

For the core analysis, it makes sense to have the discount rate r being common to all regions if capital markets are functioning well. But it could also differ in a more general setting to capture the possibility of region-specific capital market imperfections which mean that the returns to capital are not equalized across regions.

Putting this together, we can compute the value in any region i using the recursion relation specified above, yielding a solution very similar to Eq. 1 for the one-region or whole-world case,

$$V^i(X_0) = \sum_{t=0}^{\infty} D_t^i \left[b_t^i(X_t) y_t^i(X_t) - \lambda_t^i(X_t) K_t^i(X_t) \right], \quad [4]$$

where

$$D^i(0) = 1, \quad D_{T+1}^i = \prod_{t=0}^T \frac{1 - \mu^i(X_t)}{1 + r^i}$$

and

$$b_t^i(X_t) = \frac{r^i - g^i + (1 + g^i) \mu^i(X_t)}{r^i - g^i}.$$

The crucial difference is that X_t rather than the region-specific x_t^i enters, reflecting the global interdependence. Hence the willingness to pay for reductions in emissions will be interdependent and depend on the time path of emissions in other countries.

Numerical Solution

We study the model's implications using a numerical solution. To implement this, we make a few specific assumptions and solve the model with these in place. These assumptions imply that the equivalent and compensating variations are the same. After giving the solution, we choose specific parameter values to provide a quantitative assessment of the willingness to pay to avoid the risks associated with climate change. These are informed by a range of considerations including the losses that were experienced following Hurricane Katrina in the United States. Interested readers can download the spreadsheet (Dataset S1) and perform their own calculations with different parameter values.

Parameterization. To simplify notation, we temporarily revert to the case of a single region and hence drop the i superscript. We assume that, under status quo policies, the GDP (although the model is couched in terms of GDP, in principle these could be utility payoffs rather than income) and the cost of a catastrophe keep on growing at a fixed rate g so that

$$y_t = y_0 (1 + g)^t, \quad K_t = K_0 (1 + g)^t.$$

We also assume that arrival rate for the savior technology is constant at μ , which captures a rough balance of the two forces

mentioned above. Then Eq. 1 simplifies to (the derivation is in *SI Appendix, section 2*)

$$V(x_0) = y_0 \frac{1 + r}{r - g} - K_0 \Lambda, \quad [5]$$

where

$$\Lambda = \sum_{t=0}^{\infty} \left[\frac{(1 + g)(1 - \mu)}{1 + r} \right]^t \lambda(x_t).$$

This can be thought of as an expected present value operator that captures the influence of the key parameters embedded in $\lambda(x_t)$ acting through the growth parameters in x_t and μ . Moreover the solution in Eq. 5 has a nice interpretation: It is the full discounted present value of GDP absent any catastrophes, minus the expected discounted cost of catastrophes as captured in Λ and scaled by the initial condition K_0 . In our base-case numerical calculations we assume that emission flows z_t grow at a rate α which is equal to the GDP rate of growth g . We examine the effects of various policies which lower α . For example, policies which effect a Kyoto-style reduction would lower α by 30%.

We specify the arrival rate of the catastrophe Poisson process as the usual logistic function

$$\lambda(x) = e^{\gamma x} / [J + e^{\gamma x}]$$

with two parameters J and γ . (The logistic is a natural functional form in this context. The state variable, the natural log of accumulated GHGs, goes from $-\infty$ to ∞ ; the probability goes from 0 to 1. We should expect the function to be initially convex while it flattens out eventually.)

A convenient feature of this special case is that the compensating and equivalent variations are the same, making it unnecessary to differentiate between them in what follows. We derive the formula for them in *SI Appendix, section 3* and show that this is equal to

$$\theta = \frac{r - g}{1 + r} \frac{K_0}{y_0} (\Lambda - \tilde{\Lambda}) \quad [6]$$

when there is a parameter shift which moves from Λ to $\tilde{\Lambda}$. A convenient feature of Eq. 6 is that the "loss ratio," K_0/y_0 , simply multiplies the expression for the compensating variation, scaling it up or down.

The generalization of this formula to allow for regional differences is straightforward. To do so, we suppose that ϵ and δ pertain to global carbon dynamics so that they are the same for all regions. This also implies that the state variable X is global. The GDP levels y_0^i ; the costs of catastrophes K_0^i ; the functional form of the catastrophe hazard functions $\lambda^i(X)$; and the parameters g^i , r^i , and μ^i are all allowed to be region specific. Then Eq. 6 in a multiregion world becomes

$$\theta^i = \frac{r^i - g^i}{1 + r^i} \left(\frac{K_0^i}{y_0^i} \right)^i (\Lambda^i - \tilde{\Lambda}^i), \quad [7]$$

where

$$\Lambda^i = \sum_{t=0}^{\infty} \left[\frac{(1 + g^i)(1 - \mu^i)}{1 + r^i} \right]^t \lambda^i(X_t).$$

This has the neat feature that all of the interdependence is captured entirely through $\lambda^i(X_t)$. In this formula, all strategic interaction between regions is mediated by the common state variable X_t . So whether there are strategic complements or strategic substitutes depends on how the interventions of other countries change X_t .

Choice of Parameter Values. We specify baseline values for the parameters and consider a range around them in numerical solutions. Since we err toward assuming somewhat optimistic parameter values, any conclusions below are quite conservative. Advocating significant expenditures based on the chosen parameter values would hold a fortiori for the specifications and parameter values favored by those with more alarmist views of climate change problems.

Hence our baseline sets

$$\begin{aligned} r = 0.05, \quad g = 0.03, \quad \alpha = 0.03, \quad \delta = 0.03, \\ \epsilon = 0.5, \quad \mu = 0.01, \quad y_0 = 1, \quad K_0 = 2. \end{aligned} \quad [8]$$

The discount rate r is much higher than the rate advocated by those who favor strong policies to counter climate change, for example Stern (23), and close to the 5% or more that was advocated by critics of the Stern review, for example Weitzman (24). This is a specific example of our desire to make conservative assumptions since the lower the discount rate, the more future damage weighs in the calculation and the greater will be the justification for taking countermeasures. The 3% status quo growth rate is again quite optimistic. We have set $\alpha = g$, so under the status quo GHG emissions would keep step with economic growth. The values of the permanent component of emissions ϵ and the dissipation rate δ of the rest are in broad agreement with Hasler and Krusell (15) and Inman (19). Very little is known about the likelihood of a total technological solution, but the choice $\mu = 0.01$ implies that the probability of such a solution having arrived rises to 50% in 70 years, which also seems optimistic. Setting $y_0 = 1$ is just a normalization and we discuss the justification of K_0 below.

In the logistic specification of the function $\lambda(x)$ our base values are $\gamma = 1.5$ and $J = 20,000$. With these parameter choices, the probability of at least one catastrophe occurring by time T rises to 50% in $T = 56$ years and to 90% in 81 years. These again are fairly optimistic numbers. It should be clear to the reader that all of these magnitudes could be varied and an Excel spreadsheet (Dataset S1) allows the interested reader to do so.

In our Excel file (Dataset S1), we carry out the sum defining Λ in Eq. 5 to 1,000 years, when the terms generally become of the order of 10^{-12} . Again readers can easily alter the file as they wish.

Specifying the Expected Cost of a Catastrophe. As emphasized in reviews of integrated assessment models such as Metcalf (25), there is a considerable uncertainty about the right assumptions to make about the likely damages from higher carbon emissions. We anchor our estimates around the US experience of Hurricane Katrina, which hit New Orleans and other parts of the southeastern United States in late August 2005 and had many of the features that are expected to figure in future environmental catastrophes—flooding, wind and water damage to structures, loss of life, dislocation of populations and disruption of economic activity, and so on. That storm cannot be attributed directly to climate change, but a rough quantification of its effects gives us a useful starting point for considering costs of catastrophes. Although this is a specific case, it would be straightforward to assess the sensitivity of the results to alternative scenarios.

We begin with the loss of GDP. Before Katrina, New Orleans' GDP was growing fast, from \$52.38 billion in 2001 to \$72.91 billion in 2005, which is an annual growth rate of 8.6%. (These numbers come from an article on the Atlanta Federal Reserve website, "New Orleans, 10 Years after Katrina," <https://www.frbatlanta.org/economy-matters/2015/08/20/new-orleans-10-years-after-katrina>, accessed May 10, 2017.) To assess the shortfall of GDP below what it would have been without Katrina, let us take a highly conservative approach by assuming

that the GDP would have grown for the next three years at the slower rate experienced by the United States as a whole. Table 1 shows the calculations. The cumulative shortfall for the three years from 2006 to 2008 is then \$24.8 billion, which is 34% of the 2005 GDP level. Effects of Katrina continued for much longer than three years, but after 2009 the calculation gets trickier because of the need to separate the effects of Katrina from those of the Great Recession that began in 2008. We again take a conservative approach by omitting any GDP losses beyond three years.

Hurricane Katrina caused considerable loss of, and damage to, capital. For the whole region affected by Katrina, which comprises the states of Louisiana, Florida, and Mississippi, private insured and uninsured losses are estimated at \$108 billion, of which about half occurred in New Orleans alone. In addition, restoration of the damaged levees and coastal restoration and urban water management projects in New Orleans required about \$29 billion. These costs—\$54 billion + \$29 billion = \$83 billion—amount to over 113% of the city's 2005 GDP (from <https://www.thebalance.com/hurricane-katrina-facts-damage-and-economic-effects-3306023>, accessed May 10, 2017).

Thus, the capital costs (113%) and GDP losses (34%) taken together come to 147% of one year's GDP in New Orleans. It should also be noted that these calculations do not include human costs—1,836 lives lost, over 100,000 people displaced and their lives disrupted, trauma suffered by nearly the whole population (over 400,000) of that city, and much more. Our conclusion from this exercise is that a reasonable baseline (conservative) case is $K_0 = 2 y_0$. But to recognize the wide possible range of estimates, we consider variations in K_0/y_0 between 1 and 3.

Two other aspects of our specification go in opposite directions. First, we are extrapolating from Katrina, which was a local event, to the whole United States (or other comparable regions). At the present time it seems implausible that such an event could hit the whole country in any one year. However, multiple events of this kind are already hitting different areas of the country—major hurricanes, nor'easters, California's sequence of drought—wildfires—rainstorms—mudslides, and so on. Therefore, in the time frame where the probability of catastrophe in our model rises substantially—50–60 years—such losses on a national scale do not seem so implausible. However, readers can easily alter our results to suit their preferred estimates, since our formula for θ is proportional to K_0/y_0 . Second, we assume that after the loss from a catastrophic event, growth is restored to the old level g . However, catastrophes such as a shift in the pattern of a seasonal monsoon can lower the growth rate permanently, magnifying the formula for θ . Once again interested readers can easily modify the Excel file (Dataset S1) to include such effects.

Results

Single Country. We begin with the single-country case. This will help to get a feel for the quantitative magnitudes which come out of the model. In the first instance, we use the base parameters above. But we will also see how our conclusions vary with some of these.

Table 1. GDP loss due to Katrina in New Orleans

Year	Actual GDP, \$ billion	US growth rate, % per year	Hypothetical GDP, \$ billion	Shortfall, \$ billion
2005	72.91			
2006	71.18	5.8	77.14	5.96
2007	70.93	4.5	80.61	9.68
2008	72.82	1.7	81.98	9.16

First, consider the Kyoto reduction in GHG emissions, lowering α by 30%, i.e., from 0.03 to 0.021. This has a variation (compensating or equivalent) of 0.033. That is, we should be willing to pay a cost of 3.3% of GDP each year to bring about the Kyoto reduction. This is a large number; for the United States it amounts to about \$500 billion/year (and growing at 3% with GDP growth). However, to put it in perspective, it is equivalent only to permanently sacrificing one good year of economic growth.

This number is quite sensitive to the outlook for growth and if the growth projection were 1% instead of 3%, which is more in line with recent growth pessimism, then this number also falls by around one-third to 1.1% of current GDP. Perhaps not surprisingly, the willingness to pay for the Kyoto reduction is increased significantly by having a smaller discount rate with an increase in the willingness to pay to 4.8% if the discount rate is $r = 0.04$. One could also argue that the Hurricane Katrina output loss is too conservative to capture the kind of catastrophic change that could be envisaged. Suppose that K_0/y_0 is equal to 3 instead of 2; then the willingness to pay would increase to 5% of GDP. The bottom line in all cases is that the plausible willingness to pay for Kyoto-style reductions is in the range of 1 – 5%. While one would seek to design policies which do this both fairly and efficiently, the sizes of the sacrifices in consumption that are needed are small in comparison with historic increases in material living standards.

Next consider increasing μ from 0.01 to 0.015. This would raise the probability that a solution has been found in 25 years' time from 22% to 31%. This might be feasible with the kind of investment in science that has been seen in the past in pursuit of military ends or space travel. But how much does our model suggest would be a reasonable commitment of aggregate resources to achieve this end? The willingness to pay for this in our baseline case is 0.033; i.e., we should be willing to invest 3.3% of GDP, or \$500 billion/y for the United States (and growing at 3%), to raise the probability of a complete technological solution by 50%. To put this in perspective, note that this is less than the roughly 4% of GDP that the US Government spends on national defense although considerably in excess of the (around) \$20 billion/y spent on NASA and the total NSF budget of around \$6 billion.

Multiple Countries. We consider four regions which we call China, the United States, Europe, and the rest of the world (RoW). [Hassler and Krusell (15) specify Africa as the fourth region in their approach.] Their GDP shares and share of CO₂ emissions are as follows:

Region	GDP share, %	CO ₂ share, %
China	15	30
United States	15	15
Europe	20	15
RoW	50	40

The world GDP and initial emission level are both normalized to 1. And in the baseline we assume that the parameter values in Eq. 8 are maintained.

What makes the multicountry case interesting is how the willingness to pay in one country is affected by actions taken elsewhere. Due to the global nature of the externality, the willingness of the United States or China to take a Kyoto-style cut would depend on the path of emissions taken elsewhere. Associated with any proposal, therefore, would be an associated vector $\{\theta^1, \theta^2, \theta^3, \theta^4\}$ denoting the willingness to pay in each region. One of the critical issues in negotiations over emissions reduction is how the benefits and costs are shared, which creates potential for free riding, particularly when there are difficulties of enforce-

ment. Aldy et al. (26) discuss the complex issues that are involved in aligning this. Our ready-reckoner approach will be useful in giving an insight into how the heterogeneity in willingness to pay depends upon underlying differences in economic prospects for the regions of the world. If the willingness to pay is similar, then it should be easier to achieve consensus. However, there are still issues of how to enforce agreements, which we do not discuss here.

To provide a benchmark, we consider an optimistic case in which all countries follow the Kyoto benchmark with a cut in emissions such that α^i falls from 0.03 to 0.021. In this case, the willingness to pay is equal to about 3.3% in all regions of the world just as in the single-country model. Now consider what happens if one region decides to opt out of the deal and free ride. Then how big would the loss from that be to the participating countries? Suppose, for the sake of illustration, that it is the United States that opts out. Now the willingness of the other three regions to pay for Kyoto falls to around 2.7%. But the United States still gets a benefit of this amount from what other countries are doing without paying any cost. Thus, around 80% of the benefit from emissions reductions is available to a free rider conditional on full compliance elsewhere. Thus, more relevant for calculating whether a country is willing to participate is not the total gain but rather the marginal willingness to pay, assuming that other countries go along with a Kyoto-style cut, and this is only around 0.6% of GDP, much lower than the 3.3% of the one-country model. That said, it still amounts to around \$90 billion in the case of the United States or a little under \$300 per US citizen.

Another interesting question to ask is, What should be China's or the United States' willingness to pay for unilateral action, assuming that no other country participates? For this we suppose that the rest of the world maintains $\alpha^i = 0.03$ and that either China or the United States cuts α^i to 0.21. This yields a willingness to pay for the country that is cutting only 0.4% of GDP since many of the gains accrue to other countries. Note that this is less than the willingness to pay conditional on full participation by others and vividly illustrates how strategic interdependence can change willingness to pay to reduce emissions. The fact that unilateral action is more valuable when more countries participate illustrates how emissions reductions are strategic complements with greater action in one part of the world when more countries also participate in emissions reduction. So the cost of the United States leaving an agreement is also detrimental to the willingness of other countries to participate quite apart from any moral sense of shared burden.

The complementarity result does not, however, appear to be a general feature of the model. It arises because with the parameterization that we use, the world is operating in a region of the $\lambda(X)$ function whose curvature is such that, when a country makes a contribution to reducing X , then this has a bigger effect on $\lambda(X)$ when other countries also contribute, thereby increasing that country's willingness to pay.

Given the baseline parameter values, the proportionate gains and losses are similar in all regions. We now illustrate two interesting sources of heterogeneity. The first one is where countries face different growth prospects. And the second one is where the costs of catastrophes differ by region.

Recent discussions around the potential for secular stagnation [for example, Summers (27)] suggest that advanced economies such as those of the United States and Europe face a weaker growth outlook. One influential version of this view due to Gordon (28) centers on a slowdown in the innovation process, afflicting mainly advanced economies. This is possible even if there is maintenance of strong growth in China and the RoW as the process of catchup and convergence continues. We now show that views about this have a material effect on the willingness-to-pay calculations in our model.

To capture this, we assume that growth will be only around 1% in the future for the United States and Europe while remaining at 3% for China and the RoW. In our first calculation, we suppose that emissions' growth in the United States and Europe is not affected by slow growth so $\alpha = 0.03$ without emissions reduction and 0.021 in a Kyoto-style cut. Now the benefit from collective action for the United States and Europe falls to around 1% of GDP. More strikingly still, the value of unilateral action is only 0.01% of GDP or just \$18 billion in the case of the United States. So unilateral action under a pessimistic scenario for growth is highly unrewarding when there is growth pessimism for the United States.

The second source of heterogeneity that we consider concerns what happens if the losses from catastrophes are unevenly distributed. Suppose that the United States and Europe have reasons to be sanguine about the cost of catastrophes and their willingness to pay is based on $K_0/y_0 = 1$ while in China and the RoW, the losses are larger with $K_0/y_0 = 3$. Then how do we think that this will affect the geopolitics of reaching an agreement? First, consider a multilateral agreement to $\alpha = 0.021$. The willingness to pay in China and the RoW now increases by a factor of 1.5 to around 5% of GDP while that in the United States and Europe would fall to around 1.7%. The willingness to pay with unilateral action will similarly fall.

Together these results illustrate how regional interdependence matters, more so when the outlook is heterogeneous across regions. Moreover, it is clear that the magnitude of the willingness to pay is influenced significantly by views about the outlook for growth and the potential sizes of the catastrophes that might be faced. Of course, all of these numbers are only illustrative but they show that the perceptions around the distribution of damages due to climate change affect the potential for self-interest to motivate action and our framework allows us to think about the magnitudes involved and sensitivity to parameter values in a highly transparent way. The very simple model structure that we use also means that it is easy to gain an intuitive understanding of what is going on.

Ideas for Future Work

The framework that we have proposed suggests various directions for future work. One interesting issue is to explore incentives for joining coalitions and to explore a formal analysis of stable coalitions along the lines of Nordhaus (29). The model as it stands is sufficiently linear that the value function for a coalition is simply the sum of those for any partitions, but introducing gains from trade, or Nordhaus-style enforcement mechanisms by climate clubs, could create superadditivity and allow one to look for the core of the game for a given X .

Future work could also consider alternative ways in which growth is affected by climate change. In the framework presented here, catastrophes have transitory effects and eventually growth is restored. However, there is the possibility that growth could be reduced permanently and this would be an interesting issue to explore in the future.

Concluding Comments

This paper has put forward a model to evaluate the risks of climate catastrophes in a multiregion world. We have developed a simple formulation of the costs of catastrophic risk and the willingness to pay for mitigation. The model is simple and transparent and can be solved on a spread sheet, thereby giving a simple way of considering the kinds of sacrifice that a society might make to mitigate these risks. In our baseline, the numbers turn out to be quite large (typically in excess of 1% of GDP). Of course, ensuring that reductions in consumption brought about by taxation are actually spent wisely and effectively to bring about emissions reduction and/or investments in technology is by no means easy. And there are complicated issues in policy design that we have not tackled here.

ACKNOWLEDGMENTS. The authors are grateful to James Mirrlees, Simon Dietz, Gunnar Eskeland, Ian Martin, Frikk Nesje, Victor Norman, Robert Pindyck, Nicholas Stern, and participants in a workshop at the Norwegian School of Economics for helpful feedback on an earlier version of this paper. We are grateful to Takafumi Kawakubo for research assistance.

- Hassler J, Krusell P, Nycander J (2016) Climate policy. *Econ Pol* 31:503–558.
- Manne AS, Richels RG (1992) *Buying Greenhouse Insurance: The Economic Costs of CO₂ Emission Limits* (MIT Press, Cambridge, MA).
- Nordhaus W (1994) *Managing the Global Commons: The Economics of Climate Change* (MIT Press, Cambridge, MA).
- Kolstad CD (1996) Fundamental irreversibilities in stock externalities. *J Public Econ* 60:221–233.
- Golosov M, Hassler J, Krusell P, Tsyvinski A (2014) Optimal taxes on fossil fuel in general equilibrium. *Econometrica* 82:41–88.
- Gerlagh R, Liski M (2016) Carbon prices for the next hundred years. *Econ J* 128:728–757.
- Traeger C (2015) Analytic integrated assessment and uncertainty. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2667972. Accessed September 19, 2018.
- Newbold SC, Griffiths C, Moore C, Wolverton A, Kopits E (2013) A rapid assessment model for understanding the social cost of carbon. *Clim Change Econ* 4:1350001.
- Derek L, Traeger C (2014) Watch your step: Optimal policy in a tipping climate. *Am Econ J Econ Policy* 6:137–66.
- Derek L, Traeger C (2016) The economics of tipping the climate dominoes. *Nat Clim Change* 6:514–519.
- Inge van den B, Gerlagh R, Liski M (2016) A simple formula for the social cost of carbon. *J Environ Econ Manag* 77:75–94.
- van der Ploeg F, Zeeuw A (October 25, 2017) Climate tipping and economic growth: Precautionary capital and the price of carbon. *J Eur Econ Assoc*, 10.1093/jeaa/jvx036.
- Tsur Y, Zemel A (1996) Accounting for global warming risks: Resource management under event uncertainty. *J Econ Dyn Control* 20:1289–1305.
- Acemoglu D, Aghion P, Bursztyn L, Hemous D (2012) The environment and directed technical change. *Am Econ Rev* 102:131–166.
- Hassler J, Krusell P (2012) Economics and climate change: Integrated assessment in a multi-region world. *J Eur Econ Assoc* 10:974–1000.
- Stott PA, Stone DA, Allen MR (2004) Human contribution to the European heatwave of 2003. *Nature* 432:610–613.
- van Oldenborgh GJ, et al. (2017) Attribution of extreme rainfall from Hurricane Harvey. *Environ Res Lett* 12 124009.
- Lin N, Kopp RE, Hortond BP, Donnelly JP (2016) Hurricane Sandy's flood frequency increasing from year 1800 to 2100. *Proc Natl Acad Sci USA* 113:12071–12075.
- Inman M (2008) Carbon is forever. *Nat Climate Change*. Available at www.nature.com/climate/2008/0812/full/climate.2008.122.html. Accessed September 12, 2018.
- Barrett S (2008) The incredible economics of geoengineering. *Environ Resour Econ* 39:45–54.
- Dixit A (1976) *The Theory of Equilibrium Growth* (Oxford Univ Press, Oxford).
- Piketty T (2014) *Capital in the Twenty-First Century* (Harvard Univ Press, Cambridge, MA).
- Stern NH (2007) *The Economics of Climate Change: The Stern Review* (Cambridge Univ Press, Cambridge, UK).
- Weitzman ML (2007) A review of the 'Stern Review on the Economics of Climate Change'. *J Econ Lit* 45:703–724.
- Metcalfe G, Stock J (2015) The role of integrated assessment models in climate policy: A user's guide and assessment. Available at belfercenter.ksg.harvard.edu/files/dp68_metcalfe-stock.pdf. Accessed September 12, 2018.
- Aldy JE, Krupnick AJ, Newell RG, Parry IWH, Pizer WA (2010) Designing climate mitigation policy. *J Econ Lit* 48:903–934.
- Summers L (2016) The age of secular stagnation: What it is and what to do about it. *Foreign Aff* 95:2–6.
- Gordon R (2016) *The Rise and Fall of American Growth* (Princeton Univ Press, Princeton).
- Nordhaus W (2015) Climate clubs: Overcoming free-riding in international climate policy. *Am Econ Rev* 105:1339–1370.