

Averting Catastrophes: The Strange Economics of Scylla and Charybdis[†]

By IAN W. R. MARTIN AND ROBERT S. PINDYCK*

Faced with numerous potential catastrophes—nuclear and bioterrorism, mega-viruses, climate change, and others—which should society attempt to avert? A policy to avert one catastrophe considered in isolation might be evaluated in cost-benefit terms. But because society faces multiple catastrophes, simple cost-benefit analysis fails: even if the benefit of averting each one exceeds the cost, we should not necessarily avert them all. We explore the policy interdependence of catastrophic events, and develop a rule for determining which catastrophes should be averted and which should not. (JEL D61, Q51, Q54)

“Is there no way,” said I, ‘of escaping Charybdis, and at the same time keeping Scylla off when she is trying to harm my men?’

“You dare-devil,” replied the goddess, ‘you are always wanting to fight somebody or something; you will not let yourself be beaten even by the immortals.’”

—Homer, *Odyssey*¹

Like any good sailor, Odysseus sought to avoid every potential catastrophe that might harm him and his crew. But, as the goddess Circe made clear, although he could avoid the six-headed sea monster Scylla or the “sucking whirlpool” of Charybdis, he could not avoid both. Circe explained that the greatest expected loss would come from an encounter with Charybdis, which should therefore be avoided, even at the cost of an encounter with Scylla.

We modern mortals likewise face myriad potential catastrophes, some more daunting than those faced by Odysseus. Nuclear or bioterrorism, an uncontrolled viral epidemic on the scale of the 1918 Spanish flu, or a climate change catastrophe

*Martin: London School of Economics, Houghton Street, London WC2A 2AE, UK (e-mail: i.w.martin@lse.ac.uk); Pindyck: Sloan School of Management, Massachusetts Institute of Technology, 100 Main Street, Cambridge, MA 02139 (e-mail: rpindyck@mit.edu). Our thanks to Thomas Albert, Robert Barro, Simon Dietz, Christian Gollier, Derek Lemoine, Bob Litterman, Deborah Lucas, Antony Millner, Gita Rao, Edward Schlee, V. Kerry Smith, Nicholas Stern, Martin Weitzman, three anonymous referees, and seminar participants at LSE, Oxford, the University of Amsterdam, Tel-Aviv University, Universidad de Chile, Harvard, MIT, the University of Arizona, Arizona State University, and the Toulouse School of Economics for helpful comments and suggestions. Ian Martin is supported by ERC Starting Grant 639744. Both authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/aer.20140806> to visit the article page for additional materials and author disclosure statement(s).

¹*Odyssey*, Book XII, translated by Samuel Butler (1900).

are examples. Naturally, we would like to avoid all such catastrophes. But even if it were feasible, is that goal advisable? Should we instead avoid some catastrophes and accept the inevitability of others? If so, which ones should we avoid? Unlike Odysseus, we cannot turn to the gods for advice. We must turn instead to economics, the truly dismal science.

Those readers hoping that economics will provide simple advice, such as “avert a catastrophe if the benefits of doing so exceed the cost,” will be disappointed. We will see that deciding which catastrophes to avert is a much more difficult problem than it might first appear, and a simple cost-benefit rule doesn’t work. Suppose, for example, that society faces five major potential catastrophes. If the benefit of averting each one exceeds the cost, straightforward cost-benefit analysis would say we should avert all five.² We show, however, that it may be optimal to avert only (say) three of the five, and not necessarily the three with the highest benefit/cost ratios. This result might at first seem “strange” (hence the title of the paper), but we will see that it follows from basic economic principles.

Our results highlight a fundamental flaw in the way economists usually approach potential catastrophes. Consider the possibility of a climate change catastrophe—a climate outcome so severe in terms of higher temperatures and rising sea levels that it would sharply reduce economic output and consumption (broadly understood). A number of studies have tried to evaluate greenhouse gas (GHG) abatement policies by combining GHG abatement cost estimates with estimates of the expected benefits to society (in terms of reduced future damages) from avoiding or reducing the likelihood of a bad outcome.³ To our knowledge, however, all such studies look at climate change in isolation. We show that this is misleading.

A climate catastrophe is only one of a number of catastrophes that might occur and cause major damage on a global scale. Other catastrophic events may be as likely or more likely to occur, could occur much sooner, and could have an even worse impact on economic output and even mortality. One might estimate the benefits to society from averting each of these other catastrophes, again taking each in isolation, and then, given estimates of the cost of averting the event, come up with a policy recommendation. But applying cost-benefit analysis to each event in isolation can lead to a policy that is far from optimal.

Conventional cost-benefit analysis can be applied directly to “marginal” projects, i.e., projects whose costs and benefits have no significant impact on the overall economy. But policies or projects to avert major catastrophes are not marginal; their costs and benefits can alter society’s aggregate consumption, and that is why they cannot be studied in isolation.

Like many other studies, we measure benefits in terms of willingness to pay (WTP), i.e., the maximum fraction of consumption society would be willing to

²Although we will often talk of “averting” or “eliminating” catastrophes, our framework allows for the possibility of only partially alleviating one or more catastrophes, as we show in Section IVA.

³Most of these studies develop integrated assessment models (IAMs) and use them for policy evaluation. The literature is vast, but Nordhaus (2008) and Stern (2007) are widely cited examples; other examples include the many studies that attempt to estimate the social cost of carbon (SCC). For a survey of SCC estimates based on three widely used IAMs, see Greenstone, Kopits, and Wolverton (2013) and Interagency Working Group on Social Cost of Carbon (2010). These studies, however, generally focus on “most likely” climate outcomes, *not* low-probability catastrophic outcomes. See Pindyck (2013a,b) for a critique and discussion. One of the earliest treatments of environmental catastrophes is Cropper (1976).

sacrifice, now and forever, to achieve an objective. We can then address the following two questions: first, how will the WTP for averting Catastrophe A change once we take into account that other potential catastrophes B, C, D, etc., lurk in the background? We show that the WTP to eliminate A will go up.⁴ The reason is that the other potential catastrophes reduce expected future consumption, thereby increasing expected future marginal utility and therefore also the benefit of averting catastrophe A. Likewise, each individual WTP (e.g., to avert just B) will be higher the greater is the “background risk” from the other catastrophes. What about the WTP to avert all of the potential catastrophes? It will be less than the sum of the individual WTPs. The WTPs are not additive; society would probably be unwilling to spend 60 or 80 percent of gross domestic product (GDP) (and *could not* spend 110 percent of GDP) to avert all of these catastrophes.

WTP relates to the demand side of policy: it is society’s reservation price—the *most* it would sacrifice—to achieve some goal. In our case, it measures the *benefit* of averting a catastrophe. It does not tell us whether averting the catastrophe makes economic sense. For that we also need to know the *cost*. There are various ways to characterize such a cost: a fixed dollar amount, a time-varying stream of expenditures, etc. In order to make comparisons with the WTP measure of benefits, we express cost as a permanent tax on consumption at rate τ , the revenues from which would just suffice to pay for whatever is required to avert the catastrophe.

Now suppose we know, for each major type of catastrophe, the corresponding costs and benefits. More precisely, imagine we are given a list (τ_1, w_1) , $(\tau_2, w_2), \dots, (\tau_N, w_N)$ of costs (τ_i) and WTPs (w_i) associated with projects to eliminate N different potential catastrophes. That brings us to our second question: which of the N projects should we implement? If $w_i > \tau_i$ for all i , should we eliminate all N potential catastrophes? Not necessarily. We show how to decide which projects to choose to maximize social welfare.

When the projects are very small relative to the economy, and if there are not too many of them, the conventional cost-benefit intuition prevails: if the projects are not mutually exclusive, we should implement any project whose benefit w_i exceeds its cost τ_i . This intuition might apply, for example, for the construction of a dam to avert flooding in some area. Things are more interesting when projects are large relative to the economy, as might be the case for the global catastrophes mentioned above, or if they are small but large in number (so their aggregate influence is large). Large projects change total consumption and marginal utility, causing the usual intuition to break down: there is an essential interdependence among the projects that must be taken into account when formulating policy.

We are not the first to note the interdependence of large projects; early expositions of this point include Dasgupta, Sen, and Marglin (1972) and Little and Mirrlees (1974). (More recently, Dietz and Hepburn 2013 illustrate this point in the context of climate change policy.) Nor are we the first to note the effects of background risk; see, e.g., Gollier (2001) and Gollier and Pratt (1996). But to our knowledge this paper is the first to address the question of selecting among a set of large projects.

⁴As we will see, this result requires the coefficient of relative risk aversion to exceed 1.

We show how this can be done, and we use several examples to illustrate some of the counterintuitive results that can arise.

For instance, one apparently sensible response to the nonmarginal nature of large catastrophes is to decide which is the most serious catastrophe, avert that, and then decide whether to avert other catastrophes. This approach is intuitive and plausible—and wrong. We illustrate this in an example with three potential catastrophes. The first has a benefit w_1 much greater than the cost τ_1 , and the other two have benefits greater than the costs, but not that much greater. Naïve reasoning suggests we should proceed sequentially: eliminate the first catastrophe and then decide whether to eliminate the other two, but we show that such reasoning is flawed. If only one of the three were to be eliminated, we should indeed choose the first; and we would do even better by eliminating all three. But we would do best of all by eliminating the second and third and *not* the first: the presence of the second and third catastrophes makes it suboptimal to eliminate the first.

In the next section we use two very simple examples to illustrate the general interdependence of large projects, and show why, if faced with two potential catastrophes, it might not be optimal to avert both, even if the benefit of averting each exceeds the cost. In Section II we introduce our framework of analysis by first focusing on the WTP to avert a single type of catastrophe (e.g., nuclear terrorism) considered in isolation. We use a constant relative risk aversion (CRRA) utility function to measure the welfare accruing from a consumption stream, and we assume that the catastrophe arrives as a Poisson event with known mean arrival rate; thus catastrophes occur repeatedly and are homogeneous in time. Each time a catastrophe occurs, consumption is reduced by a random fraction.⁵ These simplifying assumptions make our model tractable, because they imply that the WTP to avoid a given type of catastrophe is constant over time.

This tractability is critical when, in Section III, we allow for multiple types of catastrophes. Each type has its own mean arrival rate and impact distribution. We find the WTP to eliminate a single type of catastrophe and show how it depends on the existence of other types, and we also find the WTP to eliminate several types at once. We show that the presence of multiple catastrophes may make it less desirable to try to mitigate some catastrophes for which action would appear desirable, considered in isolation. Next, given information on the cost of eliminating (or reducing the likelihood of) each type of catastrophe, we show how to find the welfare-maximizing combination of projects that should be undertaken.

Section IV presents some extensions. First, we show that our framework allows for the partial alleviation of catastrophes, i.e., for policies that reduce the likelihood of catastrophes occurring rather than eliminating them completely. The paper's central intuitions apply even if we can choose the amount by which we reduce the arrival rate of each catastrophe optimally. Second, our framework easily handles catastrophes that are directly related to one another: for example, averting nuclear terrorism might also help avert bioterrorism. Third, our results also apply to *bonanzas*, that is, to projects

⁵ Similar assumptions are made in the literature on generic consumption disasters. Examples include Backus, Chernov, and Martin (2011); Barro and Jin (2011); and Pindyck and Wang (2013). Martin (2008) estimates the welfare cost of consumption uncertainty to be about 14 percent, most of which is attributable to higher cumulants (disaster risk) in the consumption process. Barro (2013) examines the WTP to avoid a climate change catastrophe with (unavoidable) generic catastrophes in the background.

such as blue-sky research that increase the probability of events that raise consumption (as opposed to decreasing the probability of events that lower consumption).

The contribution of this paper is largely theoretical: we provide a framework for analyzing different types of catastrophes and deciding which ones should be included as a target of government policy. Determining the actual likelihood of nuclear terrorism or a mega-virus, as well as the cost of reducing the likelihood, is no easy matter. Nonetheless, we want to show how our framework might be applied to real-world government policy formulation. To that end, we survey the (very limited) literature for seven potential catastrophes, discuss how one could come up with the relevant numbers, and then use our framework to determine which of these catastrophes should or should not be averted.

I. Two Simple Examples

Why is it that “large,” i.e., nonmarginal projects are inherently interdependent and cannot be evaluated in isolation? The following simple examples should help convey some of the basic intuition, and also clarify the connection between our work and the prior literature. The first example addresses a (static) decision to undertake a set of projects, and shows how the decision rule changes if the projects are large. The second example asks whether resources should be sacrificed today to avert one or two catastrophes that will otherwise occur in the future. It illustrates the effect of background risk, the interdependence of WTPs, and the connection to cost-benefit analysis.

Static Example: Suppose we are deciding whether to undertake two independent projects.⁶ To make the basic point in the simplest possible case, we assume that these are yes/no projects, so that the resources expended on project i , e_i , equals either 0 or x_i . We can approximate net welfare, W , using a second-order Taylor expansion:

$$(1) \quad W(e_1, e_2) \approx W(0, 0) + \sum_{i=1}^2 e_i \left. \frac{\partial W}{\partial e_i} \right|_{e_1=e_2=0} + \frac{1}{2} \sum_{i=1}^2 \sum_{j=1}^2 e_i e_j \left. \frac{\partial^2 W}{\partial e_i \partial e_j} \right|_{e_1=e_2=0} .$$

If both projects are “marginal,” i.e., the x_i are very small, then we can ignore the second-order term in (1), and the optimal decision is to set $e_i = x_i$ if $\left. \frac{\partial W}{\partial e_i} \right|_{e_1=e_2=0} > 0$ and $e_i = 0$ otherwise. In other words, the standard cost-benefit rule applies: undertake a project if doing so yields an increase in net welfare. But if the projects are *not* marginal, then we cannot ignore the second-order term in (1). Now the standard cost-benefit rule fails. Why? Because of the second derivative terms, the value of project 1 depends on whether project 2 is also being carried out, and vice versa. Thus large projects cannot be evaluated independently of each other.⁷

⁶A version of this example was suggested by an anonymous referee, whom we thank.

⁷This is essentially the idea behind Dasgupta, Sen, and Marglin (1972) and Little and Mirrlees (1974). Also, note that this interdependence does *not* depend on the binary (i.e., $e_i = 0$ or x_i) nature of the projects. As we show in Section IVA, it holds even if the size of each project (i.e., x_i) can be freely chosen.

Two-Period Example: As a second example, suppose there are two potential catastrophes that, if not averted, will surely occur at a future time T . Each catastrophe will reduce consumption at time T by a fraction ϕ . Consumption today is $C_0 = 1$, so consumption at T is $C_T = 1$ if both catastrophes are averted, $C_T = 1 - \phi$ if one is averted, and $C_T = (1 - \phi)^2$ if neither is averted. Each catastrophe can be averted by sacrificing a fraction τ of consumption today and at time T . We assume CRRA utility and ignore discounting, so welfare is

$$V = \frac{1}{1 - \eta} [C_0^{1-\eta} + C_T^{1-\eta}],$$

and for simplicity let $\eta = 2$. If neither catastrophe is averted, welfare is $V_0 = -[1 + (1 - \phi)^{-2}]$.

If we avert one of the two catastrophes by sacrificing a fraction of consumption w_1 , welfare is $V_1 = -(1 - w_1)^{-1}[1 + (1 - \phi)^{-1}]$. The WTP is the fraction w_1 that equates V_0 to V_1 :

$$(2) \quad w_1 = 1 - \left[\frac{1 + (1 - \phi)^{-1}}{1 + (1 - \phi)^{-2}} \right].$$

The WTP to avert both catastrophes, $w_{1,2}$, equates V_0 to $V_{1,2} = -2(1 - w_{1,2})^{-1}$, so

$$(3) \quad w_{1,2} = 1 - \frac{2}{1 + (1 - \phi)^{-2}}.$$

Finally, if there were only one catastrophe, the WTP to avert it would be

$$(4) \quad w'_1 = 1 - \left[\frac{2}{1 + (1 - \phi)^{-1}} \right].$$

We can use equations (2), (3), and (4) to illustrate several points:

- (i) *Background risk increases the WTP to avert a catastrophe.* It is easy to see that $w_1 > w'_1$, i.e., the WTP to avert Catastrophe 1 is increased by the presence of Catastrophe 2. For example, if $\phi = 0.5$, $w_1 = 0.40$, and $w'_1 = 0.33$. Catastrophe 2 reduces C_T , raising marginal utility at time T , and thereby raising the value of averting Catastrophe 1.⁸
- (ii) *WTPs don't add.* Specifically, $w_{1,2} < w_1 + w_2$. For example, if $\phi = 0.5$, $w_{1,2} = 0.60 < w_1 + w_2 = 0.80$. Sacrificing 40 percent of consumption

⁸This result is related to the notion of “risk vulnerability” introduced by Gollier and Pratt (1996). They derive conditions under which adding a zero-mean background risk to wealth will increase an agent’s risk aversion with respect to an additional risky prospect. The conditions are that the utility function exhibits absolute risk aversion that is both declining and convex in wealth, a natural assumption that holds for all hyperbolic absolute risk aversion (HARA) utility functions. Risk vulnerability includes the concept of “standard risk aversion” (Kimball 1993) as a special case. In our model, background risk is not zero-mean: background events reduce consumption in our base-line framework and increase consumption in the extension in Section IVC.

sharply increases the marginal utility loss from any further sacrifice of consumption.

- (iii) *Naïve cost-benefit analysis can be misleading.* More specifically, we might not avert a catastrophe even if the benefit of averting it—considered in isolation—exceeds the cost. For example, suppose $\phi = 0.5$ as before, so that $w_1 = w_2 = 0.4$. If $\tau_1 = \tau_2 = 0.35$, the benefit of averting each catastrophe exceeds the cost. But we should not avert both. For if we avert neither catastrophe, net welfare is $V_0 = -5$; if we avert one, net welfare is $W_1 = -4.62$; and if we avert both, net welfare is $W_{1,2} = -4.73$. Averting both is better than averting neither, but we do best by averting exactly one. To understand this, note that if we avert one catastrophe, what matters is whether the *additional* benefit from averting the second exceeds the cost, i.e., whether $(w_{1,2} - w_1)/(1 - w_1) > \tau_2$. We should not avert #2 because $(w_{1,2} - w_1)/(1 - w_1) = 0.33 < \tau_2 = 0.35$.

These examples help connect our work to the earlier literature and illustrate why large projects are interdependent. We turn next to a fully dynamic model that includes uncertainty over the arrival and impact of multiple potential catastrophes, and that lets us derive a key result regarding the set of catastrophes that should be averted.

II. The Model with One Type of Catastrophe

We first consider a single type of catastrophe. It might be a climate change catastrophe, a mega-virus, or something else. What matters is that we assume for now that this particular type of catastrophe is the only thing society is concerned about. We want to determine society's WTP to avoid this type of catastrophe, i.e., the maximum fraction of consumption, now and throughout the future, that society would sacrifice. Of course it might be the case that the revenue stream corresponding to this WTP is insufficient to eliminate the risk of the catastrophe occurring, in which case eliminating the risk is economically infeasible. Or, the cost of eliminating the risk might be lower than the corresponding revenue stream, in which case the project would have a positive net social surplus. The WTP applies only to the demand side of government policy. Later, when we examine multiple types of catastrophes, we will also consider the supply (i.e., cost) side.

To calculate a WTP, we must consider whether the type of catastrophe at issue can occur once and only once (if it occurs at all), or can occur repeatedly. For a climate catastrophe, it might be reasonable to assume that it would occur only once—the global mean temperature, for example, might rise much more than expected, causing economic damage far greater than anticipated, and perhaps becoming worse over time as the temperature keeps rising.⁹ But for most potential catastrophes, such as a mega-virus, nuclear terrorism, or nuclear war, it is more reasonable to assume that the catastrophe could occur multiple times. Throughout the paper we will assume

⁹That is why some argue that the best way to avert a climate catastrophe is to invest now in geoengineering technologies that could be used to reverse the temperature increases. See, e.g., Barrett (2008, 2009) and Kousky et al. (2009).

that multiple occurrences are indeed possible. However, in an online Appendix we examine the WTP to eliminate a catastrophe that can occur only once.

We will assume that without any catastrophe, real per-capital consumption will grow at a constant rate g , and we normalize so that at time $t = 0$, $C_0 = 1$. Let c_t denote log consumption. We define a catastrophe as an event that permanently reduces log consumption by a random amount ϕ (so that ϕ is roughly the fraction by which the level of consumption falls). Thus if the catastrophic event first occurs at time t_1 , $C_t = e^{gt}$ for $t < t_1$ and then falls to $C_t = e^{-\phi+gt}$ at $t = t_1$. For now we impose no restrictions on the probability distribution for ϕ . We use a simple CRRA utility function to measure welfare, and denote the index of relative risk aversion by η and rate of time preference by δ . Unless noted otherwise, in the rest of this paper we will assume that $\eta > 1$, so utility is negative. This is consistent with both the finance and macroeconomics literatures, which put η in the range of 2–5 (or even higher). Later we treat the special case of $\eta = 1$, i.e., log utility.

We assume throughout this paper that the catastrophic event of interest occurs as a Poisson arrival with mean arrival rate λ , and that the impact of the n th arrival, ϕ_n , is independent and identically distributed across realizations n . Thus the process for consumption is

$$(5) \quad c_t = \log C_t = gt - \sum_{n=1}^{Q(t)} \phi_n,$$

where $Q(t)$ is a Poisson counting process with known mean arrival rate λ , so when the n th catastrophic event occurs, consumption is multiplied by the random variable $e^{-\phi_n}$. We follow Martin (2013) by introducing the *cumulant-generating function* (CGF),

$$\kappa_t(\theta) \equiv \log E e^{c_t \theta} \equiv \log E C_t^\theta.$$

As we will see, the CGF summarizes the effects of various types of risk in a convenient way. Since the process for consumption given in (5) is a Lévy process, we can simplify $\kappa_t(\theta) = \kappa(\theta)t$, where $\kappa(\theta)$ means $\kappa_1(\theta)$. In other words, the t -period CGF scales the 1-period CGF linearly in t . We show in the Appendix that the CGF is then¹⁰

$$(6) \quad \kappa(\theta) = g\theta + \lambda(E e^{-\theta\phi_1} - 1).$$

Given this consumption process, welfare is

$$(7) \quad E \int_0^\infty \frac{1}{1-\eta} e^{-\delta t} C_t^{1-\eta} dt = \frac{1}{1-\eta} \int_0^\infty e^{-\delta t} e^{\kappa(1-\eta)t} dt \\ = \frac{1}{1-\eta} \frac{1}{\delta - \kappa(1-\eta)},$$

¹⁰We could allow for $c_t = g_t - \sum_{n=1}^{N(t)} \phi_n$, where g_t is any Lévy process, subject to the condition that ensures finiteness of expected utility. (For the special case in (5), $g_t = gt$ for a constant g .) This only requires that the term $g\theta$ in the CGFs is replaced by $g(\theta)$, where $g(\theta)$ is the CGF of g_t , so if there are Brownian shocks with volatility σ , and jumps with arrival rate ω and stochastic impact J , then $g(\theta) = \mu\theta + \frac{1}{2}\sigma^2\theta^2 + \omega(E e^{\theta J} - 1)$. This lets us handle Brownian shocks and unavoidable catastrophes without modifying the framework. Since the generalization has no effect on any of our qualitative results, we stick to the simpler formulation.

where $\kappa(1 - \eta)$ is the CGF of equation (6) with $\theta = 1 - \eta$. Note that equation (7) is quite general and applies to *any* distribution for the impact ϕ . But note also that welfare is finite only if the integrals converge, and for this we need $\delta - \kappa(1 - \eta) > 0$ (Martin 2013).

Eliminating the catastrophe is equivalent to setting $\lambda = 0$ in equation (6). We denote the CGF in this case by $\kappa^{(1)}(\theta)$. (This notation will prove convenient later when we allow for several types of catastrophes.) So if we sacrifice a fraction w of consumption to avoid the catastrophe, welfare is

$$(8) \quad \frac{(1 - w)^{1-\eta}}{1 - \eta} \frac{1}{\delta - \kappa^{(1)}(1 - \eta)}.$$

The WTP to eliminate the event (i.e., set $\lambda = 0$) is the value of w that equates (7) and (8):

$$\frac{1}{1 - \eta} \frac{1}{\delta - \kappa(1 - \eta)} = \frac{(1 - w)^{1-\eta}}{1 - \eta} \frac{1}{\delta - \kappa^{(1)}(1 - \eta)}.$$

Should society avoid this catastrophe? This is easy to answer because with only one type of catastrophe to worry about, we can apply standard cost-benefit analysis. The benefit is w , and the cost is the permanent tax on consumption, τ , needed to generate the revenue to eliminate the risk. We should avoid the catastrophe as long as $w > \tau$. As we will see shortly, when there are multiple potential catastrophes the benefits from eliminating each are interdependent, causing this simple logic to break down.¹¹

III. Optimal Policy with Multiple Catastrophes

We now allow for multiple types of catastrophes, show how to find the WTP to avert each type, and examine the interrelationship among the WTPs. We can then address the issue of choosing which catastrophes to avert. We aim to answer the following question: given a list of costs and benefits of eliminating different types of catastrophes, which ones should we eliminate? The punch line will be given by Result 2 below: there is a fundamental sense in which *benefits add but costs multiply*. This will imply that there may be a substantial penalty associated with implementing several projects. As a result, it may be optimal not to avert catastrophes whose elimination seems justified in naïve cost-benefit terms.

¹¹ A referee suggested that we could have alternatively expressed benefits in terms of the *growth rate* of consumption, rather than as a percentage of its level. Then WTP would be the maximum reduction in the growth rate society would be willing to sacrifice to avert a catastrophe. Expressing benefits this way is certainly reasonable. If the costs of averting catastrophes are likewise modeled as required reductions in the growth rate (which we think is much less reasonable), our dynamic model could be written in a static form. Modeling benefits and costs in terms of levels is the conventional approach, which we have chosen to maintain.

As before, we assume that a catastrophic event causes a drop in consumption. We also assume that these events occur independently of each other. So log consumption is

$$(9) \quad c_t = \log C_t = gt - \sum_{n=1}^{Q_1(t)} \phi_{1,n} - \sum_{n=1}^{Q_2(t)} \phi_{2,n} - \dots - \sum_{n=1}^{Q_N(t)} \phi_{N,n},$$

where $Q_i(t)$ is a Poisson counting process with mean arrival rate λ_i , and the CGF is

$$(10) \quad \begin{aligned} \kappa(\theta) = g\theta + \lambda_1(Ee^{-\theta\phi_1} - 1) \\ + \lambda_2(Ee^{-\theta\phi_2} - 1) + \dots + \lambda_N(Ee^{-\theta\phi_N} - 1). \end{aligned}$$

Here we write ϕ_i for a representative of any of the $\phi_{i,n}$ (since catastrophic impacts are all independent and identically distributed within a catastrophe type). If no catastrophes are eliminated, welfare is again given by equation (7). In the absence of catastrophe type i , welfare is

$$\frac{1}{1 - \eta} \frac{1}{\delta - \kappa^{(i)}(1 - \eta)},$$

where the i superscript indicates that λ_i has been set to zero. Thus willingness to pay to eliminate catastrophe i satisfies

$$\frac{(1 - w_i)^{1-\eta}}{1 - \eta} \frac{1}{\delta - \kappa^{(i)}(1 - \eta)} = \frac{1}{1 - \eta} \frac{1}{\delta - \kappa(1 - \eta)},$$

and hence

$$w_i = 1 - \left(\frac{\delta - \kappa(1 - \eta)}{\delta - \kappa^{(i)}(1 - \eta)} \right)^{\frac{1}{\eta-1}}.$$

Similarly, the WTP to eliminate some arbitrary subset S of the catastrophes, which we will write as w_S , is given by

$$(11) \quad (1 - w_S)^{1-\eta} = \frac{\delta - \kappa^{(S)}(1 - \eta)}{\delta - \kappa(1 - \eta)}.$$

(The superscript S on the CGF indicates that λ_i is set to zero for all $i \in S$.) The next result shows how w_S , the WTP for eliminating the subset of catastrophes, can be connected to the WTPs for each of the individual catastrophes in the subset.

RESULT 1: *The WTP to avert a subset, S , of the catastrophes is linked to the WTPs to avert each individual catastrophe in the subset by the expression*

$$(12) \quad (1 - w_S)^{1-\eta} - 1 = \sum_{i \in S} [(1 - w_i)^{1-\eta} - 1].$$

PROOF:

The result follows from a relationship between $\kappa^{(S)}(\theta)$ and the individual $\kappa^{(i)}(\theta)$. Note that $\kappa^{(i)}(\theta) = \kappa(\theta) - \lambda_i 0(Ee^{-\theta\phi_i} - 1)$ and $\kappa^{(S)}(\theta) = \kappa(\theta) - \sum_{i \in S} \lambda_i (Ee^{-\theta\phi_i} - 1)$. (This is effectively the *definition* of the notation $\kappa^{(i)}$ and $\kappa^{(S)}$.) Thus

$$\sum_{i \in S} \kappa^{(i)}(\theta) = |S|\kappa(\theta) - \sum_{i \in S} \lambda_i (Ee^{-\theta\phi_i} - 1) = (|S| - 1)\kappa(\theta) + \kappa^{(S)}(\theta),$$

where $|S|$ denotes the number of catastrophes in the subset S , and hence

$$\sum_{i \in S} \frac{\delta - \kappa^{(i)}(1 - \eta)}{\delta - \kappa(1 - \eta)} = \frac{(|S| - 1)(\delta - \kappa(1 - \eta)) + (\delta - \kappa^{(S)}(1 - \eta))}{\delta - \kappa(1 - \eta)}.$$

Using (11), we have the result. ■

If, say, there are $N = 2$ types of catastrophes, then Result 1 implies that

$$(13) \quad 1 + (1 - w_{1,2})^{1-\eta} = (1 - w_1)^{1-\eta} + (1 - w_2)^{1-\eta}.$$

Thus we can express the WTP to eliminate both types of catastrophes, $w_{1,2}$, in terms of w_1 and w_2 . But note that these WTPs do not add: since the function $(1 - x)^{1-\eta}$ is convex, equation (13) implies that $w_{1,2} < w_1 + w_2$, by Jensen's inequality.

By the same reasoning, it can be shown that $w_{1,2,\dots,N} < \sum_{i=1}^N w_i$. Likewise, if we divide the N catastrophes into two groups, 1 through M and $M + 1$ through N , then $w_{1,2,\dots,N} < w_{1,2,\dots,M} + w_{M+1,\dots,N}$. The WTP to eliminate all N catastrophes is less than the sum of the WTPs for each of the individual catastrophes, and less than the sum of the WTPs to eliminate any two groups of catastrophes.

A. Which Catastrophes to Avert?

The WTP, w_i , measures the benefit of averting Catastrophe i as the *maximum* fraction of consumption society would sacrifice to achieve this result. We measure the corresponding cost as the *actual* fraction of consumption that would have to be sacrificed, via a permanent consumption tax τ_i , to generate the revenue needed to avert the catastrophe. Thus we could avert all the catastrophes in some set S at the cost of multiplying consumption by $\prod_{i \in S} (1 - \tau_i)$ forever.^{12, 13}

¹²This *multiplicative cost* assumption implies that it is cheaper in absolute terms to avert a given catastrophe if the economy is small than if it is large. We think this is the natural formulation, because we also model the impact of catastrophes as multiplicative (and thus additive in logs, as in (9)), but we could alternatively have assumed that consumption was multiplied by $1 - \sum_{i \in S} \tau_i$. When costs are small relative to the aggregate economy we have $\prod_{i \in S} (1 - \tau_i) \approx 1 - \sum_{i \in S} \tau_i$, so the two assumptions are essentially identical. When costs are not small, our multiplicative cost assumption is conservative, because it implies a smaller cost of averting groups of catastrophes than the alternative additive assumption would. But even with our multiplicative formulation, it will often not be optimal to avert catastrophes that, considered in isolation, appear to pass a cost-benefit hurdle.

¹³As a referee pointed out, if the model is reformulated so that both benefits and costs affect the growth rate rather than level of consumption, then the problem becomes separable and the standard cost-benefit rule applies. With respect to costs, we would then have consumption multiplied by a factor $\exp(-\sum_{i \in S} \tau_i t)$. We have chosen

Thus, if we eliminate some subset S of the catastrophes, welfare (net of taxes) is

$$(14) \quad \frac{\prod_{i \in S} (1 - \tau_i)^{1-\eta}}{(1 - \eta)(\delta - \kappa^{(S)}(1 - \eta))} = \frac{\prod_{i \in S} (1 - \tau_i)^{1-\eta}}{(1 - \eta)(\delta - \kappa(1 - \eta))(1 - w_S)^{1-\eta}},$$

where the equality follows from (11). Our goal is to pick the set of catastrophes to be eliminated to maximize this expression. To do so, it will be convenient to define

$$(15) \quad K_i = (1 - \tau_i)^{1-\eta} - 1 \quad \text{and} \quad B_i = (1 - w_i)^{1-\eta} - 1.$$

Here K_i is the percentage loss of utility that results when consumption is reduced by τ_i percent, and likewise for B_i .

These utility-based definitions of costs and benefits are positive and increasing in τ_i and w_i , respectively, and $K_i > B_i$ if and only if $\tau_i > w_i$. For small τ_i , we have the linearization $K_i \approx (\eta - 1)\tau_i$; and for small w_i , we have $B_i \approx (\eta - 1)w_i$. The utility-based measures have the nice property that the B_i s across catastrophes are additive (by Result 1) and the K_i s are multiplicative. That is, the benefit from eliminating, say, three catastrophes is $B_{1,2,3} = B_1 + B_2 + B_3$, and the cost is $K_{1,2,3} = (1 + K_1)(1 + K_2)(1 + K_3) - 1$. This allows us to state our main result in a simple form.

RESULT 2 (Benefits Add, Costs Multiply): *It is optimal to choose the subset, S , of catastrophes to be eliminated to solve the problem*

$$(16) \quad \max_{S \subseteq \{1, \dots, N\}} V = \frac{1 + \sum_{i \in S} B_i}{\prod_{i \in S} (1 + K_i)},$$

where if no catastrophes are eliminated (i.e., if S is the empty set) then the objective function in (16) is taken to equal 1.

PROOF:

If we choose some subset S then, using Result 1 to rewrite the denominator of expression (14) in terms of the individual WTPs, w_i , expected utility equals

$$\frac{\prod_{i \in S} (1 - \tau_i)^{1-\eta}}{(1 - \eta)(\delta - \kappa(1 - \eta)) \left(1 + \sum_{i \in S} [(1 - w_i)^{1-\eta} - 1]\right)}$$

to model costs as taking the form $\prod_{i \in S} (1 - \tau_i)$ because this is the conventional assumption made in the literature, and to us seems more reasonable.

or, rewriting in terms of B_i and K_i ,

$$\frac{\prod_{i \in S} (1 + K_i)}{(1 - \eta)(\delta - \kappa(1 - \eta))(1 + \sum_{i \in S} B_i)}.$$

Since $(1 - \eta)(\delta - \kappa(1 - \eta)) < 0$, the optimal set S that maximizes the above expression is the same as the set S that solves the problem (16). ■

It is problem (16) that generates the strange economics of the title. To understand how the problem differs from what one might naïvely expect, notice that the set S solves

$$\max_S \log\left(1 + \sum_{i \in S} B_i\right) - \sum_{i \in S} \log(1 + K_i).$$

One might think that if costs and benefits K_i and B_i are all small, then—since $\log(1 + x) \approx x$ for small x —this problem could be closely approximated by the simpler problem

$$(17) \quad \max_S \sum_{i \in S} (B_i - K_i).$$

This linearized problem is separable, which vastly simplifies its solution: a catastrophe should be averted if and only if the benefit of doing so, B_i , exceeds the cost, K_i . But the linearized problem is only a tolerable approximation to the true problem if the *total* number of catastrophes is limited, and in particular, if $\sum_{i \in S} B_i$ is small. It is not enough for the B_i s to be individually small. The reason is that averting a large number of small catastrophes has the same aggregate impact on consumption (and marginal utility) as does averting a few large catastrophes. We illustrate this with the following example.

Example 1 (Many Small Catastrophes): Suppose we have a large number of identical (but independent) small potential catastrophes, each with $B_i = B$ and $K_i = K$. The naïve intuition is to eliminate all if $B > K$, and none if $B \leq K$. As Result 3 below shows, the naïve intuition is correct in the latter case; but if $B > K$ we should not eliminate all of the catastrophes. Instead, the number to eliminate, m , must solve the problem

$$(18) \quad \max_m \frac{1 + mB}{(1 + K)^m}.$$

In reality, m must be an integer, but we will ignore this constraint for simplicity. The optimal choice, m^* , is then determined by the first-order condition associated with (18),

$$\frac{B}{(1 + K)^{m^*}} - \frac{(1 + m^*B) \log(1 + K)}{(1 + K)^{m^*}} = 0.$$

Solving this equation for m^* , we find that $m^* = 1 / \log(1 + K) - 1/B$.

If $w = 0.020$, $\tau = 0.015$, and $\eta = 2$, $B \approx 0.020$, $K \approx 0.015$, and $m^* = 17$. But if $\eta = 3$, $m^* = 9$. And if $\eta = 4$, $B \approx 0.062$, $K \approx 0.031$, and $m^* = 6$. A larger value of η implies a smaller number m^* , because the percentage drop in consumption, $1 - (1 - \tau)^m$, results in a larger increase in marginal utility, and thus a greater loss of utility from averting one more catastrophe.

Does it matter how large is the “large number” of catastrophes in this example (assuming it is larger than the number we will avert)? No, because we fixed the values of w and τ (and hence B and K) for each catastrophe. But if we go back a step and consider what determines w , it could indeed matter. The catastrophes we do not avert represent “background risk,” and more background risk makes w larger. Thus w (and hence B) will be larger if we face 200 small catastrophes than if we face only 50.

B. Scylla and Charybdis

Suppose there are $N = 2$ types of catastrophes, and B_1 is sufficiently greater than K_1 that we will definitely avert Catastrophe 1. Should we also avert Catastrophe 2? Result 2 provides the answer: only if the benefit-cost ratio B_2/K_2 exceeds the following hurdle rate:

$$(19) \quad \frac{B_2}{K_2} > 1 + B_1.$$

Thus the fact that society is going to avert Catastrophe 1 *increases the hurdle rate* for Catastrophe 2. Furthermore, the greater is the benefit B_1 , the greater is the increase in the hurdle rate for Catastrophe 2. Notice that this logic also applies if $B_1 = B_2$ and $K_1 = K_2$; it might be the case that only one of two identical catastrophes should be averted.

As we saw in the two-period example of Section I, what matters is the *additional* benefit from averting Catastrophe 2, i.e., $(w_{1,2} - w_1)/(1 - w_1)$. Substituting in the definitions of K_i and B_i , we can see that equation (19) is equivalent to $(w_{1,2} - w_1)/(1 - w_1) > \tau_2$. It can easily be the case that $w_2 > \tau_2$ but $(w_{1,2} - w_1)/(1 - w_1) < \tau_2$. The reason is that these are not marginal projects, so $w_{1,2} < w_1 + w_2$. This is what raises the hurdle rate in equation (19). To avert Catastrophe 1, society is willing to sacrifice up to a fraction w_1 of consumption, so the remaining consumption is lower and marginal utility is higher, increasing the utility loss from the second tax τ_2 .

Example 2 (Two Catastrophes): To illustrate this result, suppose $\tau_1 = 20\%$ and $\tau_2 = 10\%$. Figure 1 shows which catastrophes should be averted for different values of w_1 and w_2 . When $w_i < \tau_i$ for both catastrophes (the bottom-left rectangle), neither should be averted. We should avert both only for combinations (w_1, w_2) in the middle lozenge-shaped region. That region shrinks considerably when we increase η . In the context of equation (19), the larger is η the larger is B_1 , and thus the larger is the hurdle rate for averting the second catastrophe.

Consider the point $(w_1, w_2) = (60\%, 20\%)$ in panel B of Figure 1. As shown, we should avert only the first catastrophe even though $w_2 > \tau_2$. Here $B_1 = 5.25$,

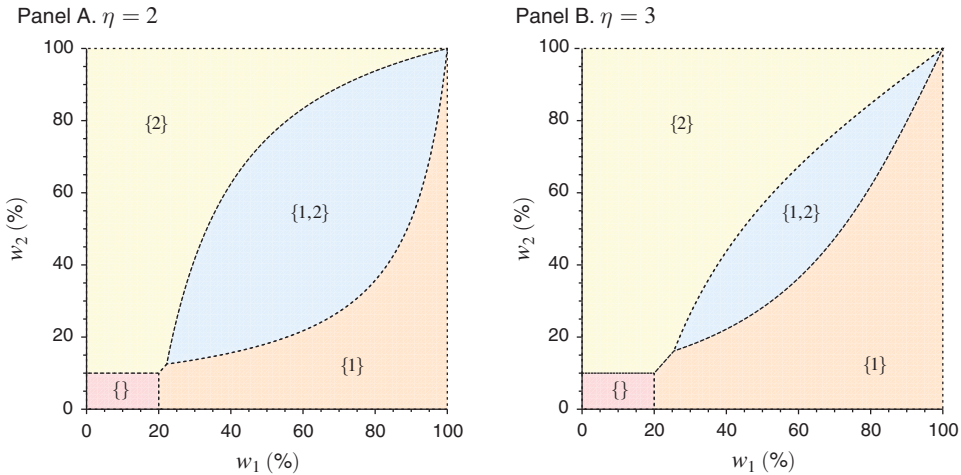


FIGURE 1

Notes: There are two potential catastrophes, with $\tau_1 = 20\%$ and $\tau_2 = 10\%$. The figures show, for all possible values of w_1 and w_2 , which catastrophes should be averted (in curly brackets). We should avert both catastrophes only for combinations (w_1, w_2) in the middle shaded region. That region shrinks considerably when risk aversion, η , increases.

$B_2 = 0.56$, and $K_2 = 0.23$, so $B_2/K_2 = 2.39 < 1 + B_1 = 6.25$. Equivalently, $w_{1,2} = 61.7\%$, so $(w_{1,2} - w_1)/(1 - w_1) = 4.3\% < \tau_2 = 10\%$. The additional benefit from averting Catastrophe 2 is less than the cost.

How is the WTP to avert Catastrophe 1 affected by the existence of Catastrophe 2? Catastrophe 2 is a kind of “background risk” that (i) reduces expected future consumption; and (ii) thereby raises future expected marginal utility. Because each catastrophic event reduces consumption by some percentage ϕ , the first effect *reduces* the WTP; there is less (future) consumption available, so the event causes a smaller absolute drop in consumption. The second effect *raises* the WTP because the loss of utility is greater when total consumption has been reduced. If $\eta > 1$ so that expected marginal utility rises sufficiently when consumption falls, the second effect dominates, and the existence of Catastrophe 2 will on net increase the benefit of averting Catastrophe 1, and raise its WTP.

C. Multiple Catastrophes of Arbitrary Size

With multiple catastrophes of arbitrary size, the solution of problem (16) is much more complicated. How does one find the set S in practice? In general, one can search over every possible subset of the catastrophes to find the subset that maximizes the objective function in (16). With N catastrophes there are 2^N possible subsets to evaluate. There is a stark contrast here with conventional cost-benefit analysis, in which an individual project can be evaluated in isolation.

The next result shows that we can eliminate certain projects from consideration, before checking all subsets of the remaining projects.

RESULT 3 (Do No Harm): *A project with $w_i \leq \tau_i$ should never be implemented.*

Example 3 (Sequential Choice Is Not Optimal): Suppose that there are three catastrophes with $(K_1, B_1) = (0.5, 1)$ and $(K_2, B_2) = (K_3, B_3) = (0.25, 0.6)$; these numbers apply if, say, $\eta = 2$ and $(\tau_1, w_1) = (\frac{1}{3}, \frac{1}{2})$ and $(\tau_2, w_2) = (\tau_3, w_3) = (\frac{1}{5}, \frac{3}{8})$. If only one were to be eliminated, we should choose the first (so that in equation (16), $V = 1.33$); and we would do even better by eliminating all three (so that $V = 1.37$). But we would do best of all by eliminating the second and third catastrophes and *not* the first (so that $V = 1.41$).

The next example shows, again with three types of catastrophes, how the choice of which to avert can vary considerably with the costs and benefits and with risk aversion.

Example 4 (Choosing among Three Catastrophes): We now extend Example 2 by adding a third catastrophe. Specifically, suppose that there are three potential catastrophes with $\tau_1 = 20\%$, $\tau_2 = 10\%$, and $\tau_3 = 5\%$. Figure 2 shows, for various different values of w_3 and η , which potential catastrophes should be averted as w_1 and w_2 vary between 0 and 1. (Figure 2 is analogous to Figure 1, except that now there is a third potential catastrophe.)

When η is close to 1, as in panel A of Figure 2, the usual intuition applies: Catastrophe 3 is always averted (since $w_3 > \tau_3$), and Catastrophes 1 and 2 should be averted if $w_i > \tau_i$. Panel B of Figure 2 shows that this intuition fails when $\eta = 2$; now it is *never* optimal to avert all three catastrophes. In panel C of Figure 2, we increase w_3 to 20 percent, and the decision becomes complicated. Consider what happens as we move horizontally across the figure, keeping w_2 fixed at 50 percent. For $w_1 < 30\%$, we avert Catastrophes 2 and 3 but not Catastrophe 1, even when $w_1 > \tau_1 = 20\%$. The reason is that the additional benefit from including Catastrophe 1, $(w_{1,2,3} - w_{2,3}) / (1 - w_{2,3})$, is less than the cost, τ_1 . If $w_1 > 30\%$, the additional benefit exceeds the cost, so we should avert Catastrophe 1. But when w_1 is greater than 70 percent (but less than 90 percent), we should avert Catastrophes 1 and 2, but not 3; the additional benefit of also averting Catastrophe 3, i.e., $(w_{1,2,3} - w_{1,2}) / (1 - w_{1,2})$, is less than the cost, τ_3 . Finally, when we increase η to 3, in panel D of Figure 2, the range of values of w_1 and w_2 for which all three catastrophes should be averted is much smaller.

We now show that the presence of many small potential catastrophes raises the hurdle rate required to prevent a large one.

Example 5 (Multiple Small Catastrophes Can Crowd Out a Large Catastrophe): Suppose that there are many small, independent, catastrophes, each with cost k and benefit b , and one large catastrophe with cost K and benefit B . Then we must compare

$$\max_m \frac{1 + mb}{(1 + k)^m} \quad \text{with} \quad \max_m \frac{1 + B + mb}{(1 + K)(1 + k)^m}.$$

Ignoring the integer constraint, and assuming that it is optimal to eliminate at least one small catastrophe, the optimized values of these problems are

$$\frac{b(1 + k)^{1/b}}{e \log(1 + k)} \quad \text{and} \quad \frac{b(1 + k)^{(1+B)/b}}{e(1 + K) \log(1 + k)},$$




FIGURE 2

Notes: There are three catastrophe types with $\tau_1 = 20\%$, $\tau_2 = 10\%$, and $\tau_3 = 5\%$. Different panels make different assumptions about w_3 and η . Numbers in brackets indicate which catastrophes should be averted for different values of w_1 and w_2 .

respectively. It follows that we should eliminate the large catastrophe if and only if

$$(20) \quad \frac{B}{\log(1 + K)} > \frac{b}{\log(1 + k)}.$$

Thus the hurdle rate for elimination of the large catastrophe is increased by the presence of the small catastrophes.

Figure 3 shows this graphically. Here $\eta = 4$ and the small catastrophes, indicated on each figure by a small solid circle, have $w_i = 1\%$ and $\tau_i = 0.5\%$ (on the left) or $w_i = 1\%$ and $\tau_i = 0.25\%$ (on the right). If the large catastrophe lies in the shaded region determined by (20), it should not be averted. In contrast, absent the small catastrophes, the major one would be averted if it lies anywhere above the dashed 45-degree line.

Example 6 (Choosing among Eight Catastrophes): Figure 4 shows some numerical experiments. Each panel plots randomly chosen (from a uniform distribution on $[0, 50\%]$) WTPs and costs, w_i and τ_i , for eight catastrophes. Fixing these w_i s and τ_i s,




FIGURE 3

Notes: Illustration of Example 5. The presence of many small catastrophes (each with cost τ_i and WTP w_i , indicated by a solid circle) expands the region of inaction for a larger catastrophe, which should not be averted if its cost τ and WTP w lie in the shaded region.

we calculate B_i and K_i for a range of values of η . We then find the set of catastrophes that should be eliminated to maximize (16). These are indicated by dots in each panel; catastrophes that should *not* be eliminated are indicated by crosses. The 45° line is shown in each panel; points below it have $w_i < \tau_i$ and hence should never be averted. Points above the line have $w_i > \tau_i$, so the benefit of averting exceeds the cost. Even so, it is often not optimal to avert.

Panel A of Figure 4 shows that when η is close to 1, every catastrophe above the 45° line should be averted, consistent with Result 5. As η increases above 1.2, the optimal project selection depends in a complicated way on the level of risk aversion. When $\eta = 5$, it is optimal to avert just one “doomsday” catastrophe. When $\eta = 4$, it is optimal to avert two *different* catastrophes. When $\eta = 3$, three should be averted—but still not the doomsday catastrophe. As η declines further, it again becomes optimal to avert the doomsday catastrophe.

IV. Extensions

Thus far, we have made various assumptions to keep things simple. We have taken an “all-or-nothing” approach in which a catastrophe is averted entirely or not at all. We have assumed that a policy to avert catastrophe A has no effect on the likelihood of catastrophe B. And we have assumed that catastrophes are, well, *catastrophes*: that is, bad news. This section shows that all three assumptions are inessential. We can allow for partial, as opposed to total, alleviation of catastrophes; we can allow for the possibility that a policy to avert (say) nuclear terrorism decreases the likelihood of bioterrorism; and we can use the framework to consider optimal policies with respect to potential *bonanzas*—projects such as blue-sky research or infrastructure investment that increase the probability of something *good* happening (as opposed to decreasing the probability of something bad happening).




FIGURE 4

Notes: Eight catastrophes. Optimal project choice at different levels of risk aversion, η .

A. Partial Alleviation of Catastrophes

As a practical matter, the complete elimination of some catastrophes may be impossible or prohibitively expensive. A more feasible alternative may be to reduce the likelihood that the catastrophe will occur, i.e., to reduce the Poisson arrival rate λ . For example, Allison (2004) suggests that the annual probability of a nuclear terrorist attack is $\lambda \approx 0.07$. While reducing the probability to zero may not be possible, we might be able to reduce λ substantially at a cost that is less than the benefit. Should we do that, and how would the answer change if we are also considering reducing the arrival rates for other potential catastrophes?

Our analysis of multiple catastrophes makes this problem easy to deal with. We consider the possibility of reducing the arrival rate of some catastrophe from λ to $\lambda(1 - p)$, which we call “alleviating the catastrophe by probability p .” We write $w_1(p)$ for the WTP to do just that for the first type of catastrophe. Thus w_1 , in our earlier notation, is equal to $w_1(1)$.

We consider two forms of partial alleviation. First, suppose there are specific policies that alleviate a given catastrophe type by some probability; an example is the rigorous inspection of shipping containers. This implies a discrete set of policies to consider, and the previous analysis goes through essentially unmodified. Second, we allow the probability by which the catastrophe is alleviated to be chosen optimally. Perhaps surprisingly, the discrete flavor of our earlier results still holds, and those results are almost unchanged.

Discrete Partial Alleviation.—To find the WTP to alleviate the first type of catastrophe by probability p , that is, $w_1(p)$, we make use of a property of Poisson processes. We can split the “type-1” catastrophe into two subsidiary types: 1a (arriving at rate $\lambda_{1a} \equiv \lambda_1 p$) and 1b (arriving at rate $\lambda_{1b} \equiv \lambda_1(1 - p)$).¹⁴ Thus we can rewrite the CGF (10) in the equivalent form:

$$\begin{aligned} \kappa(\theta) = g\theta + & \underbrace{\lambda_{1a}(Ee^{-\theta\phi_1} - 1)}_{\text{type 1a, arriving at rate } \lambda_{1a}} + \underbrace{\lambda_{1b}Ee^{-\theta\phi_1} - 1}_{\text{type 1b, arriving at rate } \lambda_{1b}} \\ & + \underbrace{\sum_{i=2}^N \lambda_i(Ee^{-\theta\phi_i} - 1)}_{\text{all other types}}, \end{aligned}$$

so that alleviating catastrophe 1 by probability p corresponds to setting λ_{1a} to zero, and alleviating catastrophe 1 by probability $1 - p$ corresponds to setting λ_{1b} to zero. This fits the partial alleviation problem into our framework. For example, Result 1 implies that $1 + (1 - w_1(1))^{1-\eta} = (1 - w_1(p))^{1-\eta} + (1 - w_1(1 - p))^{1-\eta}$, and the argument below equation (13) implies that $w_1(p) + w_1(1 - p) > w_1(1)$ for all $p \in (0, 1)$. For example, $w_1(\frac{1}{2}) > \frac{1}{2} w_1(1)$; the WTP to reduce the likelihood by one-half is more than one-half of the WTP to eliminate it entirely.

More generally, we can split each type of catastrophe into two or more subtypes. Suppose Catastrophe 2 can be alleviated by 20 percent at some cost, and by 30 percent at a higher cost, but it cannot be fully averted. We can split this into type 2a arriving at rate $0.2 \times \lambda_2$, which can be averted at cost $\tau_{2a} < 1$; type 2b arriving at rate $0.3 \times \lambda_2$, which can be averted at cost $\tau_{2b} < 1$; and type 2c, arriving at rate $0.5 \times \lambda_2$, which cannot be averted.

To summarize, our framework can accommodate without modification policies that alleviate catastrophes by some probability, if catastrophe types are defined appropriately.

Optimal Partial Alleviation.—Now we allow the probability by which a given catastrophe is alleviated to be chosen freely. We assume that for each catastrophe i we are given the cost function $\tau_i(p)$ associated with alleviating by probability p . For

¹⁴The mathematical fact in the background is that if we start with a single Poisson process with arrival rate λ , and independently color each arrival red with probability p and blue otherwise, the red and blue processes are each Poisson processes, with arrival rates λp and $\lambda(1 - p)$, respectively.

now we do not specify the particular form of $\tau_i(p)$, but below we will consider a natural special case in which $\tau_i(p)$ is determined as a function of $\tau_i \equiv \tau_i(1)$ and p . The next result, whose proof is in the Appendix, links the WTP to alleviate a catastrophe by some probability to the WTP to avert fully: it shows that $w_i(p)$ is determined by $w_i \equiv w_i(1)$ and p .

RESULT 6: *The WTP to avert catastrophe i by probability $p \in [0, 1]$ is given in terms of $w_i = w_i(1)$ by the formula*

$$w_i(p) = 1 - \left\{ 1 + p \left[(1 - w_i)^{1-\eta} - 1 \right] \right\}^{\frac{1}{1-\eta}}.$$

In terms of $B_i(p)$, defined, analogous to (15), by $B_i(p) = [1 - w_i(p)]^{1-\eta} - 1$, we have

$$B_i(p) = pB_i.$$

Defining $K_i(p) = (1 - \tau_i(p))^{1-\eta} - 1$, the optimization problem is to

$$\max_{p_j \in [0, 1]} \frac{1 + \sum_{j=1}^N B_j(p_j)}{\prod_{j=1}^N (1 + K_j(p_j))},$$

or equivalently, using Result 6 to write $B_j(p_j) = p_j B_j$ and defining $k_i(p) = \log(1 + K_i(p))$,

$$\max_{p_j \in [0, 1]} \log \left(1 + \sum_{j=1}^N p_j B_j \right) - \sum_{j=1}^N k_j(p_j).$$

If the functions $k_j(\cdot)$ are convex, which we now assume is the case, then this is a convex problem, so that the Kuhn-Tucker conditions are necessary and sufficient. Attaching multipliers γ_j to the constraints $p_j - 1 \leq 0$ and μ_j to the constraints $-p_j \leq 0$, we have the following necessary and sufficient conditions: for all j , we have $\gamma_j \geq 0$ and $\mu_j \geq 0$, and

$$\frac{B_j}{1 + \sum_i p_i B_i} - k'_j(p_j) = \gamma_j - \mu_j \quad \text{where} \quad \gamma_j(p_j - 1) = 0 \quad \text{and} \quad \mu_j p_j = 0.$$

To go further, we consider two alternative assumptions about the cost functions $k_i(p)$.

Alternative 1: Inada-type conditions on $k_i(p)$.—Suppose that $k'_j(0) = 0$ and $k'_j(1) = \infty$. Then we can rule out corner solutions, so all Lagrange multipliers are zero and

$$(21) \quad \frac{B_j}{k'_j(p_j)} = 1 + \sum_i p_i B_i \quad \text{for each } j.$$

If it is optimal to avert at least one catastrophe, then $1 + \sum_i p_i B_i > \prod_{j=1}^N (1 + K_j(p_j))$ and hence $1 + \sum_i p_i B_i > 1 + K_j(p_j)$ for all j . But then, using the fact that $k'_j(p_j) = K'_j(p_j)/[1 + K_j(p_j)]$, condition (21) implies that $B_j > K'_j(p_j)$ at any interior optimum.¹⁵ Compare this with the corresponding condition in the naïve problem $\max_{p_j} \sum_j B_j(p_j) - \sum_j K_j(p_j)$, which is that $B_j = K'_j(p_j)$. Once again, the presence of multiple catastrophes raises the hurdle rate, but now for an increase in p_j , i.e., greater alleviation.

Alternative 2: A benchmark functional form for $k_i(p)$.—Suppose that

$$(1 - \tau_i(p))(1 - \tau_i(q)) = 1 - \tau_i(p + q) \quad \text{for all } p, q, \text{ and } i,$$

so that “alleviating by p ” and then “alleviating by q ” is as costly as “alleviating by $p + q$ in one go.” (This might hold if, e.g., a deadly virus comes from goats or chimps, and funds can be devoted to goat research, chimp research, or both. It would not hold if, e.g., there is a finite cost of alleviating by 0.5 but an infinite cost of fully averting.) This assumption pins down the form of the cost function:¹⁶ writing $\tau_i(1) = \tau_i$, we must have $\tau_i(p) = 1 - (1 - \tau_i)^p$ or, equivalently, $1 + K_i(p) = (1 + K_i)^p$. Then the functions $k_i(\cdot)$ defined above are linear:

$$(22) \quad k_i(p_i) = \log(1 + K_i(p_i)) = p_i k_i,$$

where $k_i \equiv \log(1 + K_i)$. Thus $k'_i(p_i) = k_i$, an exogenous constant independent of p_i .

By analyzing the Kuhn-Tucker conditions, the set of catastrophes can be divided into three groups. First, there are catastrophes j that should not be averted even partially (so that $p_j = 0$). For these catastrophes the cost-benefit trade-off is unattractive, in that

$$\frac{B_j}{k_j} < 1 + \sum_i p_i B_i.$$

Then there are the catastrophes that should be fully averted. These are catastrophes j that pass a certain hurdle rate,

$$\frac{B_j}{k_j} > 1 + \sum_i p_i B_i.$$

Finally, there may be catastrophes that are partially averted. These must satisfy

$$\frac{B_j}{k_j} = 1 + \sum_i p_i B_i.$$

¹⁵Remember that $B_j = B_j(1)$ is a number, not a function; since $B_j(p_j) = p_j B_j$, from Result 6, we can also interpret B_j as the marginal benefit of an increase in p_j , that is, $B'_j(p_j)$.

¹⁶To see this, note that the defining assumption can be rewritten as $h(p) + h(q) = h(p + q)$, where $h(x) = \log(1 - \tau_i(x))$. This is Cauchy’s functional equation, whose solution (given that $\tau_i(x)$ and hence $h(x)$ is monotonic) is that $h(x) = cx$ for some constant c . The result follows on imposing $\tau_i(1) = \tau_i$.

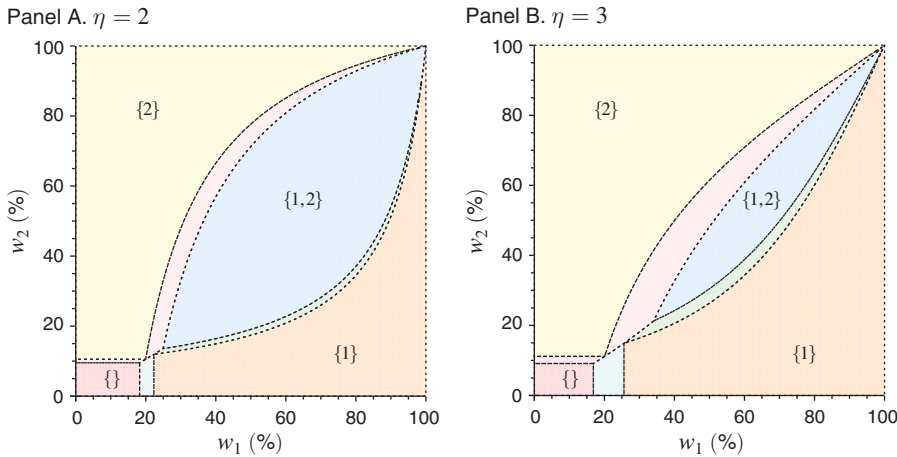


FIGURE 5

Notes: Modifying the example illustrated in Figure 1 to allow for partial alleviation with cost functions as in (22). There are two potential catastrophes with $\tau_1 = 20\%$, and $\tau_2 = 10\%$. Unnumbered zones are areas where one of the catastrophes should be partially alleviated (and it should be obvious from the location which one).

Catastrophes are therefore characterized by their benefit-cost ratios B_j/k_j . These can be thought of as parameters of the policy choice problem. If, by coincidence, two or more different types of catastrophes have the same ratio B_j/k_j , then we may have two or more types of catastrophe that are partially alleviated. But generically, all catastrophes will have different values of B_j/k_j and so *at most one catastrophe should be partially alleviated*; the remainder are all-or-nothing, and should be fully averted if their benefit-cost ratio exceeds the critical hurdle rate $X \equiv 1 + \sum_i p_i B_i$, and not averted at all if their benefit-cost ratio is less than X .¹⁷ The interdependence manifests itself through the fact that the hurdle rate X is dependent on the characteristics of, and optimal policies regarding, *all* of the catastrophes.

This is illustrated in Figure 5, which makes the same assumptions about w_i and τ_i as in Figure 1; the only difference is that we now allow for optimal partial alleviation, with cost functions $k_i(p)$ as in (22). The basic intuition is not altered by partial alleviation.

B. Related Catastrophes

Thus far, we have thought of policy responses to one catastrophe as having no effect on the likelihood of another catastrophe. We might expect, however, that a policy to avert nuclear terrorism may also help to avert bioterrorism. As in Section IVA, our framework allows for this possibility, once catastrophe types are defined appropriately.

¹⁷This characterization fails in the all-or-nothing case, as can be seen by considering an example with two catastrophes and $B_1 = 8, K_1 = 0.4, B_2 = 36, K_2 = 4$. The optimal policy with partial alleviation is to avert Catastrophe 1 fully, and Catastrophe 2 with probability 0.371. Correspondingly, Catastrophe 2 has a lower B_j/k_j . But in the all-or-nothing case, it is best to avert Catastrophe 2 and not Catastrophe 1.

For example, we may want to bundle nuclear and bioterrorism together into a single catastrophe type that can be averted at some cost. When a terrorist attack occurs, it may be either a biological attack or a nuclear attack. The distribution of damages associated with biological attacks may differ from the distribution of damages associated with nuclear attacks; the resulting distribution for ϕ , the loss associated with the “bundled” catastrophe, is then simply a mixture of the two distributions.

To illustrate how our framework can accommodate this possibility, suppose nuclear and bioterrorism are the only two types of catastrophe, with arrival rates λ_1 and λ_2 and stochastic impacts ϕ_1 and ϕ_2 , respectively. If the two are entirely independent, and policies to avert them are independent (as we have been implicitly assuming thus far) then the CGF is

$$(23) \quad \kappa(\theta) = g\theta + \lambda_1(Ee^{-\theta\phi_1} - 1) + \lambda_2(Ee^{-\theta\phi_2} - 1).$$

Alternatively, if we believe that the same policy action will avert both nuclear and bioterrorism, we can think of there being a single catastrophe¹⁸ that arrives at rate $\lambda \equiv \lambda_1 + \lambda_2$, and such that a fraction $p \equiv \lambda_1/(\lambda_1 + \lambda_2)$ of arrivals correspond to nuclear attacks with stochastic impact ϕ_1 , and a fraction $1 - p = \lambda_2/(\lambda_1 + \lambda_2)$ correspond to bio-attacks with stochastic impact ϕ_2 . This ensures that the arrival rate of nuclear attacks is λ_1 , as before, and similarly for bio-attacks. Then we can think of the CGF as

$$(24) \quad \kappa(\theta) = g\theta + \lambda(Ee^{-\theta\phi} - 1).$$

Equations (23) and (24) describe the same CGF, since $Ee^{-\theta\phi} = pEe^{-\theta\phi_1} + (1 - p)Ee^{-\theta\phi_2}$. If policies to avert nuclear and bioterrorism are best thought of separately, then it is natural to work with (23); averting nuclear terrorism corresponds to setting $\lambda_1 = 0$. If, on the other hand, a policy to avert nuclear terrorism will also avert bioterrorism, then it is more natural to work with (24); averting both corresponds to setting $\lambda = 0$.

Lastly, we can combine the results of this section and Section IVA to allow a single policy to avert multiple catastrophes partially (and potentially by different probabilities). Consider a policy that alleviates catastrophe type 1 with probability p_1 and type 2 with probability p_2 . Then split types 1 and 2 into four separate types: types 1a and 1b have arrival rates $\lambda_1 p_1$ and $\lambda_1(1 - p_1)$, respectively, and types 2a and 2b have arrival rates $\lambda_2 p_2$ and $\lambda_2(1 - p_2)$. Now view types 1a and 2a as an amalgamated Poisson process with arrival rate $\tilde{\lambda} \equiv \lambda_1 p_1 + \lambda_2 p_2$ (with impact distribution equal to a mixture of distributions ϕ_1 and ϕ_2 with weights $\lambda_1 p_1/(\lambda_1 p_1 + \lambda_2 p_2)$ and $\lambda_2 p_2/(\lambda_1 p_1 + \lambda_2 p_2)$). The policy option then is to set $\tilde{\lambda}$ to zero, and the previous results go through unchanged.

¹⁸As in footnote 13, if we have a “red” Poisson process with arrival rate λ_1 and a “blue” Poisson process with arrival rate λ_2 , we can define a “color-blind” stochastic process that does not distinguish between blue and red arrivals. This stochastic process is also a Poisson process, with arrival rate $\lambda_1 + \lambda_2$.

C. Bonanzas

Our framework also applies to projects that may lead to good outcomes. For simplicity, suppose that log consumption is $c_t = gt$ in the absence of any action. There are also projects $j = 1, \dots, m$ that can be implemented. If project j is implemented, log consumption is augmented by the process $\sum_{i=1}^{Q_j(t)} \phi_{j,i}$; if they are all implemented, log consumption follows

$$c_t = gt + \sum_{i=1}^{Q_1(t)} \phi_{1,i} + \dots + \sum_{i=1}^{Q_m(t)} \phi_{m,i},$$

where the processes $Q_1(t), \dots, Q_m(t)$ are Poisson processes as before. For consistency with previous sections, we define $\kappa(\theta) = g\theta$ to be the CGF of log consumption growth if no policies are implemented, $\kappa^{(j)}(\theta) = g\theta + \lambda_j(Ee^{\theta\phi_{j,1}} - 1)$ to be the CGF of log consumption growth if project j is implemented, and $\kappa^{(S)}(\theta) = g\theta + \sum_{j \in S} \lambda_j(Ee^{\theta\phi_{j,1}} - 1)$ to be the CGF of log consumption growth if projects $j \in S$ are implemented.

If no projects are implemented, expected utility is $1/[(1 - \eta)(\delta - \kappa(1 - \eta))]$. If projects $j \in S$ are implemented, expected utility is $1/[(1 - \eta)(\delta - \kappa^{(S)}(1 - \eta))]$. The WTP for the set S of projects, w_S , therefore satisfies

$$(1 - w_S)^{1-\eta} = \frac{\delta - \kappa^{(S)}(1 - \eta)}{\delta - \kappa(1 - \eta)}.$$

This is the analog of equation (11). Similarly, because we have the key formula $\sum_{i \in S} \kappa^{(i)}(\theta) = (|S| - 1)\kappa(\theta) + \kappa^{(S)}(\theta)$, we immediately have the formula (12). Results 1 and 2 therefore also hold for bonanzas, by the same logic as before.

V. Some Rough Numbers

This paper is largely theoretical in nature; our objective has been to show that policies or projects to avert or reduce the likelihood of major catastrophes cannot be analyzed in isolation, and the problem of deciding which catastrophes should be averted and which should not is nontrivial. However, it is useful to examine some rough numbers to see how our framework could be applied in practice. To that end, we examine some of the potential catastrophes that we believe are important, along with some very rough estimates of the likelihood, potential impact, and cost of averting each of them.

The CGF of (10) applies to *any* distribution for the impacts ϕ_i . In our case, we will assume that $z_i = e^{-\phi_i}$ is distributed according to a power distribution with parameter $\beta_i > 0$:

$$(25) \quad b(z_i) = \beta_i z_i^{\beta_i - 1}, \quad 0 \leq z_i \leq 1.$$

TABLE 1—CHARACTERISTICS OF SEVEN POTENTIAL CATASTROPHES

Potential Catastrophe	$\eta = 2$			$\eta = 4$		
	w_i	B_i	K_i	w_i	B_i	K_i
Mega-virus $\lambda_i = 0.02, \beta_i = 5, \tau_i = 0.02$	0.159	0.189	0.020	0.309	2.030	0.062
Climate $\lambda_i = 0.004, \beta_i = 4, \tau_i = 0.04$	0.048	0.050	0.042	0.180	0.812	0.130
Nuclear terrorism $\lambda_i = 0.04, \beta_i = 17, \tau_i = 0.03$	0.086	0.095	0.031	0.141	0.580	0.096
Bioterrorism $\lambda_i = 0.04, \beta_i = 32, \tau_i = 0.03$	0.047	0.049	0.031	0.079	0.280	0.096
Floods $\lambda_i = 0.17, \beta_i = 100, \tau_i = 0.02$	0.061	0.065	0.020	0.096	0.356	0.062
Storms $\lambda_i = 0.14, \beta_i = 100, \tau_i = 0.02$	0.051	0.053	0.020	0.082	0.293	0.062
Earthquakes $\lambda_i = 0.03, \beta_i = 100, \tau_i = 0.01$	0.011	0.011	0.010	0.020	0.063	0.031
Avert all seven:	0.339	0.513	0.188	0.442	4.415	0.677

Notes: For each catastrophe, table shows estimate of mean arrival rate λ , impact distribution parameter β , and prevention cost τ , as discussed in Appendix C. The impact of each event is assumed to follow equation (25); $\beta = Ez/(1 - Ez)$, where $z = e^{-\phi}$ is the fraction of consumption remaining following the event (so a large β implies a small expected impact). For each value of η , the table shows w_i , the WTP to avert catastrophe i as a fraction of consumption, and the benefit and cost in utility terms, B_i and K_i , assuming $g = \delta = 0.02$. The bottom row shows the WTP to avert all seven catastrophes, and the corresponding benefit $B_{1,\dots,7}$ and cost $K_{1,\dots,7}$ in utility terms.

Thus $E(e^{-\phi_i\theta}) = \beta_i/(\beta_i + \theta)$. A large value of β_i implies a large Ez_i and thus a small expected impact of the event.¹⁹ Given this distribution, the CGF for N types of catastrophes is

$$(26) \kappa(1 - \eta) = g(1 - \eta) - \frac{\lambda_1(1 - \eta)}{\beta_1 - \eta + 1} - \frac{\lambda_2(1 - \eta)}{\beta_2 - \eta + 1} - \dots - \frac{\lambda_N(1 - \eta)}{\beta_N - \eta + 1}.$$

This CGF tends to infinity as $\eta \rightarrow 1 + \min_i \beta_i$ from below. In order that $\delta - \kappa(1 - \eta) > 0$, we must therefore assume that $\beta_i > \eta - 1$ for all i : catastrophes cannot be too fat-tailed.

To apply this power distribution (25) for $z_i = e^{-\phi_i}$, we determine β_i for each type of catastrophe from an estimate of Ez_i . Using estimates of λ_i and τ_i , we set $g = \delta = 0.02$ and calculate w_i, B_i , and K_i for $\eta = 2$ and $\eta = 4$. The w_i for each catastrophe is calculated taking into account the presence of the other catastrophes. The estimates of λ_i, β_i , and τ_i are summarized in Table 1 along with the calculated values of w_i, B_i , and K_i . The last row of the table shows the WTP to avert all seven catastrophes ($w_{1,\dots,7} < \sum_i w_i$) and the corresponding benefit and cost in utility terms, $B_{1,\dots,7} = \sum_i B_i$ and $1 + K_{1,\dots,7} = \prod_i (1 + K_i)$. Note that for both

¹⁹ A power distribution of this sort has often been used in modeling (albeit smaller) catastrophic events such as floods and hurricanes: see, e.g., Woo (1999). Barro and Jin (2011) show that the distribution provides a good fit to panel data on the sizes of major consumption contractions. Note $E(z_i) = \beta_i/(\beta_i + 1)$, and the variance of z_i around its mean is $\text{var}(z_i) = \beta_i/[(\beta_i + 2)(\beta_i + 1)^2]$.

values of η , $B_{1,\dots,7} > K_{1,\dots,7}$, but as we will see, it is *not* optimal to avert all seven catastrophes.

The estimates of λ_i , β_i , and τ_i are explained in Appendix C. For some of the catastrophes (floods, storms, and earthquakes), the estimates are based on a relatively large amount of data. For others (e.g., nuclear terrorism), they are based on the subjective estimates of several authors, and readers may disagree with some of the numbers. As a result, they should be viewed as speculative and largely illustrative.

Some of the catastrophes we consider involve death as opposed to a drop in consumption. In Martin and Pindyck (2014), we show that the WTP to avert the death of a fraction ϕ of the population is much greater than the WTP to avert a drop in consumption by the same fraction.²⁰ This should not be surprising; most people would pay far more to avoid a 5 percent chance of dying than they would to avoid a 5 percent drop in consumption. The difference in WTPs depends on the value of a life lost, which is often proxied by the value of a statistical life (VSL). Estimates of the VSL are in the range of 3–10 times lifetime consumption. We find that a VSL of six times lifetime consumption implies that the WTP to avoid a probability of death of ϕ is equal to the WTP to avoid a drop in consumption of at least 5ϕ . We use this multiple to translate a ϕ for death into a welfare-equivalent ϕ for lost consumption.

The estimates of w_i , B_i , and K_i in Table 1 depend on δ and η . What are the “correct” values of these two parameters? We have chosen values consistent with the macroeconomics and finance literatures, but we view δ and η as policy parameters, i.e., reflecting the choices of policymakers. Thus there are no single values that we can deem “correct.”²¹

Which of the seven potential catastrophes summarized in Table 1 should be averted? We can answer this using Result 2. Although $B_{1,\dots,7} > K_{1,\dots,7}$, it is not optimal to avert all seven. As Figure 6 shows, the correct answer depends partly on the coefficient of risk aversion, η . If $\eta = 2$, five of the seven catastrophes should be averted; earthquakes and a climate catastrophe should *not* be averted. But if $\eta = 4$, a climate catastrophe *should* be averted, but not bioterrorism, storms, or earthquakes, even though the benefit of averting each of these three catastrophes exceeds the cost.

Note from Table 1 and Figure 6 that if $\eta = 2$, the WTP to avert a climate catastrophe is just slightly greater than the WTP to avert bioterrorism (0.048 versus 0.047). However, if $\eta = 4$ the WTP for climate becomes much greater than the WTP for bioterrorism (0.180 versus 0.079). Why the sharp increase in the WTP for climate? The reason is that a climate catastrophe has a relatively low arrival rate ($\lambda = 0.004$, which implies it is very unlikely to occur in the next 20 or so years) but a large expected impact ($\beta = 4$), whereas the opposite is true for bioterrorism

²⁰To our knowledge, the literature on climate change, and in particular the use of IAMs to assess climate change policies, utilizes consumption-based damages, i.e., climate change reduces GDP and consumption directly (as in Nordhaus 2008 and Stern 2007), or reduces the growth rate of consumption (as in Pindyck 2012). Millner (2013) discusses welfare frameworks which incorporate death.

²¹The rate of time preference δ matters because catastrophic events are expected to occur infrequently, so long time horizons are involved. The macroeconomics and finance literatures suggest $\delta \approx 2\text{--}5$ percent. Some economists (e.g., Stern 2008) argue that on ethical grounds, δ should be zero. Likewise, η reflects aversion to consumption inequality across generations. In the end, δ and η are (implicitly) chosen by policymakers, who might or might not believe (or care) that their decisions reflect the values of voters. For interesting discussions of social discounting, see Caplin and Leahy (2004) and Gollier (2013). For a wide-ranging and insightful discussion of economic policy-making under uncertainty, see Manski (2013).

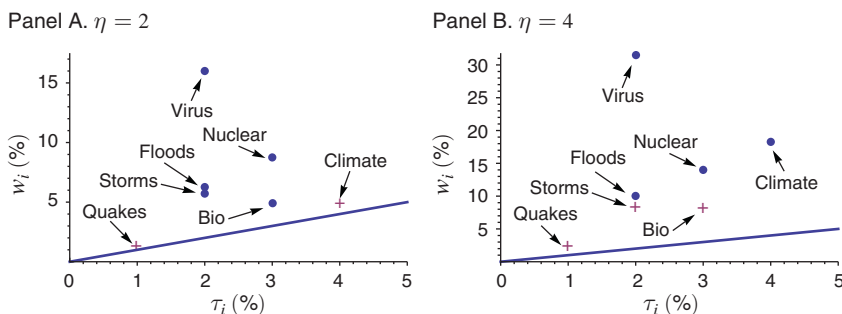


FIGURE 6

Notes: The figures show which of the seven catastrophes summarized in Table 1 should be averted. Catastrophes that should be averted are indicated by dots in each panel; catastrophes that should not be averted are indicated by crosses.

($\lambda = 0.04$ and $\beta = 32$). With high risk aversion, the impact on marginal utility of a very severe (even if unlikely) climate outcome is scaled up considerably.

We want to stress the word “Rough” in the title of this section. Some readers will disagree with the numbers in Table 1, or how we modeled some catastrophes. For example, we set the expected drop in consumption from a climate catastrophe at 20 percent. We think this is consistent with recent assessments, but those assessments are widely dispersed. Also, we assumed the arrival rate for a climate catastrophe is constant, but it is more likely to increase over time. Nonetheless, the results in Table 1 and Figure 6 illustrate our key points: policies to avert major catastrophes should not be evaluated in isolation, not all catastrophes should necessarily be averted, and the choice of which ones to avert is complex. Figure 7 makes this last point in a different way, by showing how the set of catastrophes that should be averted depends on risk aversion, η , and the time preference rate, δ . Only if η and δ are low should all seven be averted, and the optimal choice varies widely for larger values of η and δ .

VI. Conclusion

How should economists evaluate projects or policies to avert major catastrophes? We have shown that if society faces more than just one catastrophe (which it surely does), conventional cost-benefit analysis breaks down; if applied to each catastrophe in isolation, it can lead to policies that are far from optimal. The reason is that the costs and benefits of averting a catastrophe are not marginal, in that they have significant impacts on total consumption. This creates an interdependence among the projects that must be taken into account when formulating policy. In fact, as we demonstrated in Example 1, cost-benefit analysis can even fail when applied to small catastrophes if they have a nonmarginal aggregate impact.

Using WTP to measure benefits and a permanent tax on consumption as the measure of cost (both a percentage of consumption), we derived a decision rule (Result 2) to determine the optimal set of catastrophes that should be averted. And we have shown that this decision rule can yield “strange” results. For instance, as we demonstrated in Example 3, although naïve reasoning would suggest using a

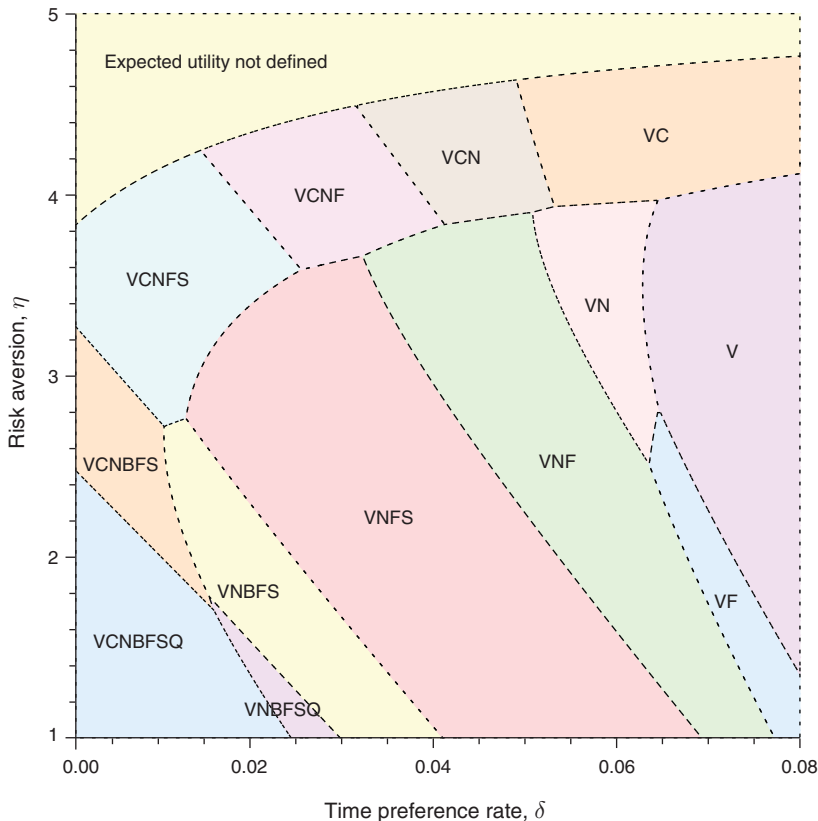


FIGURE 7

Notes: The set of catastrophes to be averted depends in a complex way on η and δ . V: virus; C: climate; N: nuclear terrorism; B: bioterrorism; F: floods; S: storms; Q: quakes.

sequential decision rule (e.g., avert the catastrophe with the largest benefit/cost ratio, then decide on the one with the next-largest ratio, etc.), such a rule is not optimal. In general, in fact, there is no simple decision rule. Instead, determining the optimal policy requires evaluating the objective function (16) of Result 2 for every possible combination of catastrophes. In a strong sense, then, the policy implications of different catastrophe types are *inextricably intertwined*.

Given that the complete elimination of some catastrophes may be impossible or prohibitively expensive, a more realistic alternative may be to reduce the likelihood that the catastrophe will occur, i.e., reduce the Poisson arrival rate λ . We have shown how that alternative can easily be handled in our framework. In the previous section we examined the costs and benefits of completely averting seven catastrophes, but we could have just as easily considered projects to reduce the likelihood of each, and given the amounts of reduction and corresponding costs, determined the optimal set of projects to be undertaken.

The theory we have presented is quite clear. (We hope most readers will agree.) But there remain important challenges when applying it as a tool for government policy, as should be evident from Section V. First, one must identify all of the relevant potential catastrophes; we identified seven, but there might be others. Second,

for each potential catastrophe, one must estimate the mean arrival rate λ_i , and the probability distribution for the impact ϕ_i . Finally, one must estimate the cost of averting or alleviating the catastrophe, which we expressed as a permanent tax on consumption at the rate τ_i . As we explained, for some catastrophes (floods, storms, and earthquakes), a relatively large amount of data are available. But for others (nuclear and bioterrorism, or a mega-virus), estimates of λ_i , β_i , and τ_i are likely to be subjective and perhaps speculative. On the other hand, one can use our framework to determine optimal policies for ranges of parameter values, and thereby determine which parameters are particularly critical, and should be the focus of research.

APPENDIX

A. The CGF in (6)

We need to calculate EC_1^θ . To do so, use the law of iterated expectations:

$$EC_1^\theta = E[E(C_1^\theta | Q(1))] = \sum_{m=0}^\infty P(Q(1) = m) E[e^{\theta(g - \sum_{n=1}^{Q(1)} \phi_n)} | Q(1) = m].$$

Since $Q(1)$ is distributed according to a Poisson distribution with parameter λ ,

$$P(Q(1) = m) = e^{-\lambda} \frac{\lambda^m}{m!}.$$

Meanwhile,

$$E[e^{\theta(g - \sum_{n=1}^{Q(1)} \phi_n)} | Q(1) = m] = E[e^{\theta(g - \sum_{n=1}^m \phi_n)}] = e^{\theta g} (E e^{-\theta \phi_1})^m$$

because the catastrophe sizes ϕ_1, \dots, ϕ_m are independent and identically distributed. Thus

$$\begin{aligned} EC_1^\theta &= \sum_{m=0}^\infty e^{-\lambda} \frac{\lambda^m}{m!} e^{\theta g} (E e^{-\theta \phi_1})^m \\ &= e^{-\lambda + \theta g} \sum_{m=0}^\infty \frac{(\lambda E e^{-\theta \phi_1})^m}{m!} = e^{-\lambda + \theta g} \exp\{\lambda E e^{-\theta \phi_1}\}. \end{aligned}$$

Taking logs, $\kappa(\theta) = g\theta + \lambda(E e^{-\theta \phi_1} - 1)$, as required.

B. Proof of Results 4, 5, and 6

PROOF OF RESULT 4:

Property (i) follows immediately from (16) and the fact that $(1 + B_i)/(1 + K_i) > 1$, so that we do better by eliminating i than by doing nothing at all. (This does not imply, however, that it is optimal to avert i itself—there may be even better options.)

Property (ii) follows by contradiction: if j were not included in the set of catastrophes to be eliminated, then we could increase the objective function in (16) by replacing i with j .

Property (iii) follows because if $B_i > K_i$, some catastrophe should be averted, by (i). And now by (ii), we see that catastrophe i should be averted.

To prove that (iv) holds, note first that by Result 3, it is never optimal to avert a catastrophe with $\tau_i \geq w_i$. So restrict attention to catastrophes with $w_i > \tau_i$. Then, by Result 2, we seek to

$$\max_S \frac{1 + \sum_{i \in S} \alpha_i^{1-\eta} - 1}{\prod_{i \in S} \beta_i^{1-\eta}},$$

where $\alpha_i = 1 - w_i$ and $\beta_i = 1 - \tau_i$ are fixed and satisfy $0 < \alpha_i < \beta_i < 1$ for all i . Since $\beta_i < 1$ and $\eta > 1$, the denominator explodes as $\eta \rightarrow \infty$. Thus, the problem is equivalent to

$$\max_S \prod_{i \in S} \beta_i^{\eta-1} \sum_{j \in S} \alpha_j^{1-\eta},$$

or

$$\max_S \prod_{i \in S} \beta_i \left(\sum_{j \in S} \alpha_j^{1-\eta} \right)^{\frac{1}{\eta-1}}.$$

Now we use the fact that for arbitrary positive x_1, \dots, x_N , we have $\lim_{\theta \rightarrow \infty} (x_1^\theta + \dots + x_N^\theta)^{\frac{1}{\theta}} = \max_i x_i$. This means that for sufficiently large η , the problem is equivalent to

$$\max_S \left(\max_{k \in S} \frac{1}{\alpha_k} \right) \prod_{i \in S} \beta_i.$$

Notice that for a fixed set S ,

$$\left(\max_{k \in S} \frac{1}{\alpha_k} \right) \prod_{i \in S} \beta_i \leq \max_{k \in S} \frac{\beta_k}{\alpha_k},$$

because $\beta_i < 1$ for all i . So given a candidate set S , we can increase the objective function by averting only the catastrophe $k \in S$ that maximizes β_k/α_k . This holds for arbitrary S , so the unconstrained optimum is achieved by averting only a single catastrophe that maximizes β_k/α_k . This is equivalent to the conditions provided in the statement of the result. ■

PROOF OF RESULT 5:

With log utility, we can use the property of the CGF that $\kappa'_i(0) = E \log C_i$ to write expected utility as

$$E \int_0^\infty e^{-\delta t} \log C_t dt = \int_0^\infty e^{-\delta t} \kappa'_i(0) dt = \kappa'_i(0) \int_0^\infty t e^{-\delta t} dt = \frac{\kappa'_i(0)}{\delta^2}.$$

If we eliminate catastrophes 1 through N costlessly, expected utility is $\kappa^{(1,\dots,N)'}(0)/\delta^2$.

So WTPs satisfy

(27)

$$\log(1 - w_{1,\dots,N}) = \frac{\kappa'(0) - \kappa^{(1,\dots,N)'}(0)}{\delta} \quad \text{and} \quad \log(1 - w_i) = \frac{\kappa'(0) - \kappa^{(i)'}(0)}{\delta}.$$

Exploiting the same relationship between CGFs as before, we find that

$$\sum_{i=1}^N \frac{\kappa'(0) - \kappa^{(i)'}(0)}{\delta} = \frac{\kappa'(0) - \kappa^{(1,\dots,N)'}(0)}{\delta},$$

and so

$$(28) \quad \sum_{i=1}^N \log(1 - w_i) = \log(1 - w_{1,\dots,N}).$$

Now, suppose we eliminate catastrophes 1 through N at cost τ_i (i.e., as before, consumption is multiplied by $(1 - \tau_i)$ to eliminate catastrophe i), then expected utility is

$$\begin{aligned} E \int_0^\infty e^{-\delta t} \log [C_t^{(1,\dots,N)}(1 - \tau_1) \cdots (1 - \tau_N)] dt \\ = \frac{1}{\delta^2} \kappa^{(1,\dots,N)'}(0) + \frac{1}{\delta} [\log(1 - \tau_1) + \cdots + \log(1 - \tau_N)], \end{aligned}$$

where $C_t^{(1,\dots,N)}$ is notation for the consumption process after eliminating catastrophes 1 through N . Using equation (27), this becomes

$$\frac{1}{\delta^2} [\kappa'(0) - \delta \log(1 - w_{1,\dots,N})] + \frac{1}{\delta} [\log(1 - \tau_1) + \cdots + \log(1 - \tau_N)],$$

and using (28), this becomes

$$\frac{1}{\delta^2} \kappa'(0) + \frac{1}{\delta} \left[\log \left(\frac{1 - \tau_1}{1 - w_1} \right) + \cdots + \log \left(\frac{1 - \tau_N}{1 - w_N} \right) \right].$$

This means that the problem is separable: we should eliminate projects i with $w_i > \tau_i$. ■

PROOF OF RESULT 6:

As in Section IVA, we can split the original catastrophe (arriving at rate λ) into N different catastrophes, each arriving at rate λ/N , and link the cost of eliminating each of these individually to the cost of eliminating the overall catastrophe. From equation (12), $N[(1 - w_i(1/N))^{1-\eta} - 1] = (1 - w_i)^{1-\eta} - 1$, and hence

$$(29) \quad w_i(1/N) = 1 - \left\{ 1 + \frac{1}{N}[(1 - w_i)^{1-\eta} - 1] \right\}^{\frac{1}{1-\eta}}.$$

This establishes the result when $p = 1/N$, for integer N . Next we extend to rationals, M/N . But this follows immediately because, using equation (12) again, $M[(1 - w_i(1/N))^{1-\eta} - 1] = (1 - w_i(M/N))^{1-\eta} - 1$, and so, using (29),

$$w_i(M/N) = 1 - \left\{ 1 + \frac{M}{N}[(1 - w_i)^{1-\eta} - 1] \right\}^{\frac{1}{1-\eta}}.$$

This establishes the result for arbitrary rational p . Finally, since WTP is a continuous function of p , and since the rationals are dense in the reals, the result holds for all p ; and it is immediate that the formula for $w_i(p)$ is equivalent to the formula for $B_i(p)$. ■

C. Catastrophe Characteristics

Here we explain the numbers in Table 1. For catastrophes such as floods, storms, and earthquakes, a relatively large amount of data are available. For others (e.g., nuclear terrorism), there is little or no data, so the numbers are based on subjective estimates of several authors.

Mega-Viruses: Numerous authors view major pandemics as both likely and having a catastrophic impact, but do not estimate probabilities of occurrence. Instead, an occurrence is simply viewed by several authors as “likely.” For a range of possibilities, see Byrne (2008) and Kilbourne (2008). For detailed discussions of how such mega-viruses could start (and maybe end), see, e.g., Beardsley (2006) and Enserink (2004). A mega-virus would directly reduce GDP and consumption by reducing trade, travel, and economic activity worldwide, but its greatest impact would be the deaths of many people. In related work (Martin and Pindyck 2014), we show that the WTP to avert an event that kills a random ϕ percent of the population is much larger than the WTP to avert an event that reduces everyone’s consumption by the same fraction ϕ .

The last major pandemic to affect developed countries was the Spanish flu of 1918–1919, which infected roughly 20 percent of the world’s population and killed 3–5 percent. Because populations today have greater mobility, a similar virus could spread more easily. We take the average mortality rate of the next pandemic to be 3.5 percent, which we estimate is equivalent in welfare terms to a roughly 17.5 percent drop in consumption.²² This corresponds to a value of $0.825/0.175 = 4.7$ for β , which we round to 5. We assume $\lambda = 0.02$, i.e., there is roughly a 20 percent chance of a pandemic occurring in the next 10 years.

What would be required to avert such an event? There is nothing that can be done to prevent new viruses from evolving and infecting humans (most likely from

²²We use a multiple of five, which, as discussed in Martin and Pindyck (2014), would apply if we use the value of a statistical life (VSL) to represent the value of a life lost, and take the VSL to be six times lifetime consumption. A great many studies have sought to estimate the VSL using data on risk-of-death choices made by individuals, and typically find that the VSL is on the order of 3–10 lifetime income or lifetime consumption. See, for example, Viscusi (1993) and Cropper and Sussman (1990).

an animal host). If a new virus is extremely virulent and contagious, containment involves (i) the implementation of systems to identify and isolate infected individuals (e.g., before they board planes or trains); and (ii) the rapid production of a vaccine (which would require yet-to-be developed technologies, and government investment in a large-scale production facility). Both of these elements involve substantial ongoing costs, but we are not aware of any estimates of how large those costs would be. As a best guess, we will assume that those costs could amount to 2 percent or more of GDP, and set $\tau = 0.02$.

Climate: A consensus estimate of the increase in global mean temperature that would be catastrophic is about 5–7°C. A summary of 22 climate science studies surveyed by Intergovernmental Panel on Climate Change (IPCC 2007) puts the probability of this occurring by the end of the century at around 5–10 percent. Preliminary drafts of the 2014 IPCC report suggest a somewhat higher probability. Weitzman (2009, 2011) argued that the probability distribution is fat-tailed, making the actual probability 10 percent or more. We will use the “pessimistic” end of the range and assume that there is a 20 percent chance that a catastrophic climate outcome could occur in the next 50 or 60 years, which implies that $\lambda = 0.004$. What would be the impact of a catastrophic increase in temperature? Estimates of the effective reduction in (world) GDP from catastrophic warming range from 10–30 percent; we will take the middle of this range, which puts ϕ at 0.20 (so that $\beta = 4$).²³

What would be the cost of averting a climate catastrophe? Some have argued that this would require limiting the atmospheric greenhouse gas (GHG) concentration to 450 ppm, and estimates of the cost of achieving this target vary widely. A starting point would be the GHG emission reductions mandated by the Kyoto Protocol (which the United States never signed). Estimates of the cost of compliance with the Protocol range from 1–3 percent of GDP.²⁴ Using the middle of that range (2 percent) and assuming that stabilizing the GHG concentration at 450 ppm would be twice as costly as Kyoto gives a cost of $\tau = 0.04$.

Nuclear Terrorism: Various studies have assessed the likelihood and impact of the detonation of one or several nuclear weapons (with the yield of the Hiroshima bomb) in major cities. At the high end, Allison (2004) put the probability of this occurring in the United States in the subsequent *ten* years at about 50 percent, which would imply a mean arrival rate $\lambda = 0.07$. Others put the probability for the subsequent ten years at around 5 percent, which implies $\lambda = 0.005$. For a survey, see Ackerman and Potter (2008). We take an average of these two arrival rates and round it to $\lambda = 0.04$.

What would be the impact? Possibly a million or more deaths in the United States, which is 0.3 percent of the population. In welfare terms, this would be equivalent to a roughly 1.5 percent drop in consumption. But the main impact would be a shock to the capital stock and GDP from a reduction in trade and economic activity worldwide, as vast resources would have to be devoted to averting further events.

²³ See, e.g., Pindyck (2012) and Stern (2007, 2008, 2013).

²⁴ See the survey of cost studies by Energy Information Administration (1998) and the more recent country cost studies surveyed in Intergovernmental Panel on Climate Change (2007).

This could easily result in a 4 percent drop in GDP and consumption, for a total (effective) drop of 5.5 percent. This corresponds to $\beta = 17$.

The cost of completely averting a nuclear terrorist attack would be considerable. One element is increased surveillance and intelligence (which presumably is already taking place). But in addition, we would need to thoroughly inspect all of the containers shipped into the United States daily; currently almost none are inspected. The combined cost of these two activities could be 3 percent of GDP, so we set $\tau = 0.03$.

Bioterrorism: Rough (and largely subjective) estimates of the likelihood of a bioterrorist attack and the costs of reducing or eliminating the risks can be found in Nouri and Chyba (2008), Lederberg (1999), and references therein. We will assume that the likelihood of a bioterrorist attack is about the same as a nuclear attack, and set $\lambda = 0.04$. Bioterrorism is unlikely to result in large numbers of deaths; instead the major impact would be a shock to GDP from panic and a reduction in trade and economic activity worldwide. As with a nuclear attack, vast resources would have to be devoted to averting further attacks. We assume that a bioterrorist attack would be less disruptive than a nuclear attack, and estimate that it could result in a 3 percent drop in GDP and consumption. This implies that $\beta = 0.97/0.03 = 32$.

The cost of averting bioterrorism includes increased surveillance and intelligence (which, as with nuclear terrorism, is presumably already taking place), but also the development of and capacity to rapidly produce vaccines and antiviral agents to counter whatever virus or other organisms were released. We will assume that the cost of completely averting bioterrorism is the same as for nuclear terrorism, so $\tau = 0.03$.

Floods, Storms, and Earthquakes: We make use of the recent study by Cavallo et al. (2013) of natural disasters and their economic impact. They utilized a dataset covering 196 countries over the period 1970–2008, which combined World Bank data on real GDP per capita with data on natural disasters and their impacts from the EM-DAT database.²⁵ Cavallo et al. (2013) estimated the effects of disasters occurring in 196 countries during 1970–1999 on the countries' GDP in the following years. There were 2,597 disasters during 1970–2008, out of a total of $39 \times 196 = 7,664$ country-year observations, which implies an average annual rate of $2,597/7,664 = 0.34$. Of these disasters, about one-half were floods, about 40 percent storms (including hurricanes), and about 10 percent earthquakes. Thus we set $\lambda = 0.17, 0.14,$ and 0.03 for floods, storms, and earthquakes, respectively.

These disasters resulted in deaths, but the number was almost always very small relative to the country's population. (For example, Hurricane Katrina caused 1,833 deaths, which was less than 0.001 percent of the US population.) We therefore ignore the death tolls from these events and focus on the impact on GDP. Cavallo et al. (2013) found that only the largest disasters (the ninety-ninth percentile in terms of deaths per million people) had a statistically significant impact on GDP ten years

²⁵The EM-DAT database was created by the Centre for Research on the Epidemiology of Disasters at the Catholic University of Louvain, and has data on the occurrence and effects of natural disasters from 1900 to the present. The data can be accessed at <http://www.cred.be/>.

following the event (reducing GDP by 28 percent relative to what it would have been otherwise). But although not statistically significant, smaller disasters (at the seventy-fifth percentile) reduced GDP by 5–10 percent. Assuming that events below the seventy-fifth percentile had no impact, we take the average impact for all three types of disasters to be a 1 percent drop in consumption, which implies $\beta = 100$. Thus floods, storms and earthquakes are relatively common catastrophes, but have relatively small impacts on average.

Storms cannot be prevented, but their impact can be reduced or completely averted. This involves relocating coastline homes and other buildings, retrofitting homes, putting power lines underground, etc. Similar steps would have to be taken to avert the impact of floods. We assume the cost of *completely* averting each of these disasters is about 2 percent of consumption. The cost of averting the impact of earthquakes should be lower—we assume 1 percent of consumption—because many buildings in vulnerable areas have already been built to withstand earthquakes. Thus we set $\tau = 0.02$ for storms and floods, and 0.01 for earthquakes.

Other Catastrophic Risks: Much less likely, but certainly catastrophic, events include nuclear war, gamma ray bursts, an asteroid hitting the Earth, and unforeseen consequences of nanotechnology. For an overview, see Bostrom and Čirković (2008), and see Posner (2004) for a further discussion, including policy implications. We ignore these other risks.

REFERENCES

- Ackerman, Gary, and William C. Potter. 2008. "Catastrophic Nuclear Terrorism: A Preventable Peril." In *Global Catastrophic Risks*, edited by Nick Bostrom and Milan M. Čirković, 402–49. Oxford: Oxford University Press.
- Allison, Graham. 2004. *Nuclear Terrorism: The Ultimate Preventable Catastrophe*. New York: Henry Holt & Company.
- Backus, David, Mikhail Chernov, and Ian Martin. 2011. "Disasters Implied by Equity Index Options." *Journal of Finance* 66 (6): 1969–2012.
- Barrett, Scott. 2008. "The Incredible Economics of Geoengineering." *Environmental and Resource Economics* 39 (1): 45–54.
- Barrett, Scott. 2009. "The Coming Global Climate-Technology Revolution." *Journal of Economic Perspectives* 23 (2): 53–75.
- Barro, Robert J. 2013. "Environmental Protection, Rare Disasters, and Discount Rates." National Bureau of Economic Research Working Paper 19258.
- Barro, Robert J., and Tao Jin. 2011. "On the Size Distribution of Macroeconomic Disasters." *Econometrica* 79 (5): 1567–89.
- Beardsley, Timothy M. 2006. "Catastrophe in the Wings?" *BioScience* 56 (3): 179.
- Bostrom, Nick, and Milan M. Čirković, eds. 2008. *Global Catastrophic Risks*. Oxford: Oxford University Press.
- Byrne, Joseph P., ed. 2008. *Encyclopedia of Pestilence, Pandemics, and Plagues*. Westport, CT: Greenwood Press.
- Caplin, Andrew, and John Leahy. 2004. "The Social Discount Rate." *Journal of Political Economy* 112 (6): 1257–68.
- Cavallo, Eduardo, Sebastian Galiani, Ilan Noy, and Juan Pantano. 2013. "Catastrophic Natural Disasters and Economic Growth." *Review of Economics and Statistics* 95 (5): 1549–61.
- Cropper, Maureen L. 1976. "Regulating Activities with Catastrophic Environmental Effects." *Journal of Environmental Economics and Management* 3 (1): 1–15.
- Cropper, Maureen L., and Frances G. Sussman. 1990. "Valuing Future Risks to Life." *Journal of Environmental Economics and Management* 19 (2): 160–74.
- Dasgupta, Partha, Amartya Kumar Sen, and Stephen A. Marglin. 1972. *Guidelines for Project Evaluation*. New York: United Nations.

- Dietz, Simon, and Cameron Hepburn.** 2013. "Benefit-Cost Analysis of Non-marginal Climate and Energy Projects." *Energy Economics* 40: 61–71.
- Energy Information Administration.** 1998. "Comparing Cost Estimates for the Kyoto Protocol." US Government Report 09-45.
- Enserink, Martin.** 2004. "Avian Influenza: Catastrophe Waiting in the Wings?" *Science* 306 (5704): 2016.
- Gollier, Christian.** 2001. *The Economics of Risk and Time*. Cambridge, MA: MIT Press.
- Gollier, Christian.** 2013. *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World*. Princeton, NJ: Princeton University Press.
- Gollier, Christian, and John W. Pratt.** 1996. "Risk Vulnerability and the Tempering Effect of Background Risk." *Econometrica* 64 (5): 1109–23.
- Greenstone, Michael, Elizabeth Koppits, and Ann Wolverton.** 2013. "Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation." *Review of Environmental Economics and Policy* 7 (1): 23–46.
- Interagency Working Group on Social Cost of Carbon.** 2010. "Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis: Under Executive Order 12866." Washington, DC: US Government.
- Intergovernmental Panel on Climate Change.** 2007. *Climate Change 2007*. Cambridge, UK: Cambridge University Press.
- Kilbourne, Edwin Dennis.** 2008. "Plagues and Pandemics: Past, Present, and Future." In *Global Catastrophic Risks*, edited by Nick Bostrom and Milan M. Ćirković, 287–307. Oxford: Oxford University Press.
- Kimball, Miles S.** 1993. "Standard Risk Aversion." *Econometrica* 61 (3): 589–611.
- Kousky, Carolyn, Olga Rostapshova, Michael A. Toman, and Richard Zeckhauser.** 2009. "Responding to Threats of Climate Change Mega-Catastrophes." Resources for the Future Working Paper 09-45.
- Lederberg, Joshua, ed.** 1999. *Biological Weapons: Limiting the Threat*. Cambridge, MA: MIT Press.
- Little, Ian M. D., and James A. Mirrlees.** 1974. *Project Appraisal and Planning for Developing Countries*. New York: Heinemann.
- Manski, Charles F.** 2013. *Public Policy in an Uncertain World: Analysis and Decisions*. Cambridge, MA: Harvard University Press.
- Martin, Ian W. R.** 2008. "Disasters and the Welfare Cost of Uncertainty." *American Economic Review* 98 (2): 74–78.
- Martin, Ian W. R.** 2013. "Consumption-Based Asset Pricing with Higher Cumulants." *Review of Economic Studies* 80 (2): 745–73.
- Martin, Ian W. R., and Robert S. Pindyck.** 2014. "Death and Destruction in the Economics of Catastrophes." Unpublished.
- Martin, Ian W. R., and Robert S. Pindyck.** 2015. "Averting Catastrophes: The Strange Economics of Scylla and Charybdis: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.20140806>.
- Millner, Antony.** 2013. "On Welfare Frameworks and Catastrophic Climate Risks." *Journal of Environmental Economics and Management* 65 (2): 310–25.
- Nordhaus, William.** 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale University Press.
- Nouri, Ali, and Christopher F. Chyba.** 2008. "Biotechnology and Biosecurity." In *Global Catastrophic Risks*, edited by Nick Bostrom and Milan M. Ćirković, 450–80. Oxford: Oxford University Press.
- Pindyck, Robert S.** 2012. "Uncertain Outcomes and Climate Change Policy." *Journal of Environmental Economics and Management* 63 (3): 289–303.
- Pindyck, Robert S.** 2013a. "Climate Change Policy: What Do the Models Tell Us?" *Journal of Economic Literature* 51 (3): 860–72.
- Pindyck, Robert S.** 2013b. "The Climate Policy Dilemma." *Review of Environmental Economics and Policy* 7 (2): 219–37.
- Pindyck, Robert S., and Neng Wang.** 2013. "The Economic and Policy Consequences of Catastrophes." *American Economic Journal: Economic Policy* 5 (4): 306–39.
- Posner, Richard A.** 2004. *Catastrophe: Risk and Response*. New York: Oxford University Press.
- Stern, Nicholas.** 2007. *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge University Press.
- Stern, Nicholas.** 2008. "The Economics of Climate Change." *American Economic Review* 98 (2): 1–37.
- Stern, Nicholas.** 2013. "The Structure of Economic Modeling of the Potential Impacts of Climate Change: Grafting Gross Underestimation of Risk onto Already Narrow Science Models." *Journal of Economic Literature* 51 (1): 838–59.

- Viscusi, W. Kip.** 1993. "The Value of Risks to Life and Health." *Journal of Economic Literature* 31 (4): 1912–46.
- Weitzman, Martin L.** 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics* 91 (1): 1–19.
- Weitzman, Martin L.** 2011. "Fat-Tailed Uncertainty and the Economics of Catastrophic Climate Change." *Review of Environmental Economics and Policy* 5 (2): 275–92.
- Woo, Gordon.** 1999. *The Mathematics of Natural Catastrophes*. London: Imperial College Press.