

NBER WORKING PAPER SERIES

AVERTING CATASTROPHES THAT KILL

Ian Martin
Robert S. Pindyck

Working Paper 23346
<http://www.nber.org/papers/w23346>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2017

Our thanks to Christian Gollier, Antony Millner and participants at various seminars for helpful comments and suggestions. The authors declare that they have no relevant or material financial interests that related to the research described in this paper. This research was partially supported by the ERC under Starting Grant 639744. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Ian Martin and Robert S. Pindyck. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Averting Catastrophes that Kill
Ian Martin and Robert S. Pindyck
NBER Working Paper No. 23346
April 2017
JEL No. D81,H12,H56,Q50,Q54

ABSTRACT

We face a variety of potential catastrophes; nuclear or bioterrorism, a climate catastrophe, and a "mega-virus" are examples. Martin and Pindyck (AER 2015) showed that decisions to avert such catastrophes are interdependent, so that simple cost-benefit analysis breaks down. They assumed that catastrophic events cause "destruction," i.e., a reduction in the stream of consumption. But some catastrophes cause death instead of, or in addition to, destruction. Here we incorporate death in a model of catastrophe avoidance, and show how it affects the interdependence of catastrophic events and the "willingness to pay" to avoid those events.

Ian Martin
Department of Finance
London School of Economics
Houghton Street
London WC2A 2AE
United Kingdom
i.w.martin@lse.ac.uk

Robert S. Pindyck
MIT Sloan School of Management
100 Main Street, E62-522
Cambridge, MA 02142
and NBER
RPINDYCK@MIT.EDU

1 Introduction

Society faces myriad potential catastrophes; nuclear or bioterrorism, an uncontrolled viral epidemic, or a climate change catastrophe are examples. In Martin and Pindyck (2015) we explored the policy interdependence of catastrophic events, showed that because they can be “non-marginal” considering them in isolation can lead to policies that are far from optimal, and developed a rule for determining which events should be averted and which should not.

In that earlier work, we assumed that a catastrophic event causes “destruction,” i.e., it reduces consumption by destroying part of the capital stock or reducing its productivity. This has been the standard way to analyze catastrophes and their impact.¹ Catastrophic events, however, can and often do kill people. For example, most of the damage from a pandemic would be the deaths of a significant fraction of the population.² Nuclear terrorism would likely result in both death and destruction. Storms, floods, earthquakes, and other natural disasters also result in significant numbers of deaths, as well as destruction. In this paper we show how to incorporate death into a model of catastrophic threats, and we compare the welfare effects of a catastrophe that causes death to one that causes only destruction.

As before, we measure the benefit from averting a catastrophe in terms of “willingness to pay” (WTP), i.e., the maximum fraction of consumption society would be willing to sacrifice, now and throughout the future, to achieve a policy objective. The WTP is thus society’s reservation price for that achieving that objective. Of course whether achieving the objective makes economic sense also depends on its cost. To make comparisons with the WTP measure of benefits, we express cost as a permanent tax τ on consumption, the revenues from which would just suffice to pay for whatever is required to avert the catastrophe.

We want to find the WTP to avert a catastrophe that would kill some fraction ψ of the population (chosen at random), leaving the consumption of those who live unchanged, as opposed to reducing the consumption of everyone by ψ . As one would expect, the WTP to avert the death of a fraction ψ of the population is much greater than the WTP to avert a drop in consumption by the same fraction; most people would pay far more to avoid a 5% chance of dying than they would to avoid a 5% drop in consumption. The difference in WTPs is large because a reduction in consumption causes a marginal loss of utility, whereas

¹Barro and Jin (2011), Martin (2008), and others refer to (generic) catastrophes as “consumption disasters.” The literature on climate change always 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)).

²The Spanish flu of 1918–1919 infected roughly 20 percent of the developed world’s population and killed 3 to 5 percent. Because populations today have greater mobility, a similar virus could spread more easily.

death causes a total loss (albeit for only a fraction of the population).

The actual difference in WTPs depends on the value of a life lost, which is often proxied by the “value of a statistical life” (VSL). The VSL is the marginal rate of substitution between wealth (or consumption, or discounted lifetime consumption) and the probability of survival. Thus it is a local measure, and is highly imperfect in that we would expect an individual or society to be willing to pay much more than the VSL to avoid certain death or a high probability of death. Nonetheless, the VSL is used widely in public policy applications, so we will use it here as an input to our WTP calculation. Many studies have estimated the VSL using data on risk-of-death choices made by individuals, and typically find numbers in the range of 3 to 10 times average lifetime income or consumption.

In our previous paper we assumed that there are N types of potential catastrophes, all of which would cause “destruction,” and that we know for each type the cost (τ_i) and WTP (w_i) of averting it. Now we allow for two types of catastrophes, one of which causes “destruction” and the other causes death. Thus the WTP to avoid the second type results from lives saved, as opposed to consumption saved. We determine the WTP to avert each type of catastrophe and the WTP to avert both types. Although our focus is on specific types of catastrophes (e.g., a pandemic), our approach can also be applied to the generic catastrophes that are part of the “consumption disaster” literature.³

Because they are non-marginal in nature, major potential catastrophes cannot be evaluated independently of each other. Here we will see that catastrophes that result in death can have higher w_i 's, and thus are even more “non-marginal” than catastrophes that only reduce consumption. As a result, a catastrophe that causes death can have a substantial impact on the decision to avert a consumption catastrophe. Averting a major pandemic, for example, might crowd out a policy to avert an extreme climate change outcome.

In the next section we lay out our VSL-based framework for measuring the welfare loss from catastrophes that kill people, but leave the consumption of survivors unchanged. In Section 3 we present a fully dynamic model that allows for two types of catastrophes, one that causes a (random) drop in consumption and the other a (random) number of deaths. These catastrophes arrive independently as Poisson processes with known arrival rates. In Section 4 we specify distributions for the random impacts of each type of catastrophe, and we solve for the WTP to avert each type, and to avert both. We then provide some numerical examples and discuss their policy implications.

³As is clear from Barro and Ursúa (2008), some of the greatest “consumption disasters” during the past century were the result of major wars. Those wars caused vast numbers of deaths (military and civilian), the welfare losses from which were likely much greater than the losses from reduced consumption.

2 Framework

We assume that at time t there are N_t identical consumers who, barring a catastrophe, remain alive indefinitely. As in our earlier work, we use a constant relative risk aversion (CRRA) utility function to measure the welfare accruing from a consumption stream, and denote the index of relative risk aversion by η and rate of time preference by δ . Unless noted otherwise, 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 to 5. We also assume that the occurrence of a catastrophe is a Poisson event with a known mean arrival rate.

Without any catastrophe, real per-capita consumption, C_t , will grow at the constant rate g , and we normalize so that at time $t = 0$, $C_0 = 1$. Likewise the population, N_t , grows at the constant rate n , and we normalize so that $N_0 = 1$. Utility for each individual comes only from consumption, and total welfare for society is proportional to the population:

$$V_0 = \mathbb{E} \int_0^\infty \frac{1}{1-\eta} N_t C_t^{1-\eta} e^{-\delta t} dt \quad (1)$$

We consider two general classes or “types” of catastrophes. The first class, which we examined in our earlier paper, consists of those catastrophes which result in “destruction,” i.e., a permanent drop in log consumption by a random fraction ϕ for the entire population (so that ϕ is roughly the fraction by which the level of total 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 $t = t_1$. Let λ_c denote the mean arrival time of this type of event.

The second class consists of catastrophes which result in the death of a fraction ψ of the population, where ψ is a random variable, and where the distribution of deaths is also random, i.e., each individual has a probability ψ of dying.⁴ This type of catastrophe, however, leaves unchanged the consumption path for those who remain alive. Thus the first occurrence of a “death” event causes the population to fall to $N_t = e^{-\psi+nt}$ but leaves C_t unchanged. Let λ_d denote the mean arrival time of this type of event. For now we make no assumptions regarding the distributions for ϕ and ψ (and these distributions need not be the same).

We will want to calculate the WTP to avert each type of catastrophe (and the WTP to avert both).⁵ To do this, we must consider whether the type of catastrophe at issue can

⁴Thus lives lost are “statistical” rather than identified. The WTP for a specified mortality risk reduction to an *identified* person is often greater than for a statistical life. See, e.g., Hammitt and Treich (2007) and Pratt and Zeckhauser (1996).

⁵We focus on completely averting a catastrophe, but as shown in Martin and Pindyck (2015), our framework also allows for partially alleviating the catastrophe, e.g., by reducing its mean arrival rate.

occur once and only once, or can occur repeatedly. For some catastrophes (e.g., a climate catastrophe), it is probably reasonable to assume it will occur only once. For most others, however, it is more reasonable to assume it can occur multiple times. In this paper, as in our earlier one, we assume that each type of catastrophe can occur multiple times.

To handle catastrophes that cause death, we must aggregate the welfare of those who remain alive after the event and the welfare of those who die. How should we evaluate the welfare of those who die? We address this problem below.

2.1 The Welfare of the Living and the Dead

What is the welfare loss for a person that dies? Can we simply assume that the utility of those who die falls to zero? This might make sense for CRRA utility with $\eta < 1$, because then $u(C) = C^{1-\eta}/(1-\eta) = 0$ when $C = 0$. But we have assumed that $\eta > 1$, so utility is negative, and unbounded as $C \rightarrow 0$. Thus simply assuming that upon death the individual loses the utility she enjoyed while alive can underestimate the welfare loss.

To see this, suppose an individual who lives two periods enjoys the same consumption $C_0 = 1$ in each period. Ignore discounting and assume $\eta = 2$, so her welfare is $V = -2C_0^{-1} = -2$. But what if she dies just before the second period? Does she simply lose the utility she otherwise would have gained from consuming 1 unit? Upon death her consumption goes to zero and $U(0) = -\infty$, which implies a loss far greater than -1 . To see why the loss must be greater than the utility she otherwise would have received in Period 2, suppose her consumption in Period 2 fell by 75%, i.e. to 0.25. Then Period 2 utility would be -4 and her welfare change would be $-4 - (-1) = -3$, a much greater loss than the utility otherwise gained from consuming 1 unit. But consuming only 25% of “normal” consumption is still (for most people) far preferable to death.

A common approach, used, e.g., by Hall and Jones (2007) in their study of the value of life-extending health expenditures, is to add a positive constant to the utility function:

$$u(C) = \frac{1}{1-\eta} C_t^{1-\eta} + b .$$

For Hall and Jones (2007) the constant b is essential because it allows them to show that as income and consumption increase, the marginal utility of an extra dollar of consumption falls relative to the marginal benefit (in terms of an increase in the probability of survival) from an extra dollar of health care spending.⁶ In their case, “death” corresponds to a drop in consumption from C_0 to some value ε , such that $u(\varepsilon) = 0$, i.e., $\varepsilon = [(\eta - 1)b]^{1/(1-\eta)}$.

⁶Becker, Philipson and Soares (2005) also use this formulation, normalize utility at death to be zero, and take $\eta \approx 0.8$ so that $b < 0$. With $\eta > 1$, a value of b sufficiently large so that $u(C) > 0$ over any relevant

We are interested in the loss of welfare resulting from death, but not a comparison of marginal benefits from additional health care spending, so we can retain the CRRA utility function without adding a constant. We want to treat death using the same framework that we use to treat destruction, i.e., as a utility loss from a drop in consumption. Thus we assume that “death” results in a percentage drop in consumption from C_t to εC_t , with $\varepsilon \ll 1$, which implies a large (but finite) drop in utility. At issue is what value to use for ε .

2.2 The Value of Life

We are using the drop in consumption from $C_0 = 1$ to ε as a proxy for the welfare loss resulting from death. To determine the value to use for ε , we need to know the value of a life (or more precisely, the value of a life lost). There is a large literature on this topic, which focuses on the “value of a statistical life” (VSL), defined as the marginal rate of substitution between wealth (or income, or consumption, or discounted consumption over a lifetime) and the probability of survival. Thus it is a local (and imperfect) measure that tells us how much an individual (or society) would pay in terms of a small decrease in wealth or consumption in return for a small increase in the probability of survival. It does *not* tell us how much an individual or society would pay to avoid certain death, or a significant probability of death, which we expect would be much more than the VSL. It also need not aggregate consistently; a VSL estimate applied to a country’s entire population can easily exceed the present value of the country’s projected GDP over the next, say, 40 years.⁷

Many studies have sought to estimate the VSL using data on risk-of-death choices made by individuals, such as the decision to take a riskier but higher-paying job rather than a safer job. (See, e.g., Viscusi (1993) and Cropper and Sussman (1990).) Those studies typically find that the VSL is on the order of 3 to 10 times lifetime income or lifetime consumption. Consistent with those findings, the U.S. Environmental Protection Agency uses a value of about \$9 million for the VSL when it conducts cost-benefit analyses of proposed policies.⁸

range of C is also needed to ensure that indifference curves between consumption and life expectancy slope down. See Rosen (1988) and Pratt and Zeckhauser (1996). Arrow and Priebsch (2014) propose a generalized hyperbolic absolute risk aversion utility function which is bounded for any value of C .

⁷For an overview of the VSL and its measurement, see Andersson and Treich (2009) and Hammitt (2007). Other issues arise with the VSL model. For example, the macro and finance literatures put the coefficient of relative risk aversion, η , well above 1. Estimates of the income elasticity of the VSL are around 0.5 to 0.6; see, e.g., Viscusi and Aldy (2003). Kaplow (2005) shows that this low income elasticity is inconsistent with $\eta > 1$. Also, the VSL is a partial equilibrium measure that ignores transfers across society (if I live to 100, the young will have to work harder to produce the goods I will consume); see, e.g., Arthur (1981).

⁸U.S. Environmental Protection Agency (2014), page 7-8: “EPA currently recommends a default central VSL of \$7.9 million (in 2008 dollars) to value reduced mortality for all programs and policies.” This translates

We translate the VSL into a value for ε as follows. Let w represent wealth or lifetime consumption, which we will take to be a multiple of current consumption, and let p be the *ex ante* probability of death (so $1 - p$ is the probability of survival). We assume that an individual can reduce p at the cost of some reduction in w . Let $u(w)$ be the individual's utility if alive, and $v(w)$ be her utility if dead, with $u(w) > v(w)$ and $u'(w) > v'(w)$. The VSL measures the trade-off between wealth (or lifetime consumption) and the probability of survival, i.e., it is the marginal rate of substitution between w and $1 - p$:

$$\text{VSL} = -\frac{dw}{d(1-p)} = \frac{dw}{dp} = \frac{u(w) - v(w)}{(1-p)u'(w) + pv'(w)}. \quad (2)$$

Note that $u(w)$ and $v(w)$ are measured in utils, and $u'(w)$ and $v'(w)$ are measured in utils per dollar, so that the VSL is measured in dollars. The VSL is a cardinal measure, invariant to linear transformations of u or v . Doubling the VSL implies a doubling of the compensation required to incur an incremental increase in the risk of death.

The VSL is increasing in the *ex ante* probability of death p ; if p is high there is less incentive to limit spending to reduce p because it is unlikely the individual will survive and have the opportunity to enjoy whatever wealth remains.⁹ In our case, the *ex ante* probability of death, whether from natural causes or a catastrophe, is low. Likewise, most empirical studies of the VSL are based on the behavior of populations for which p varies, but over a small range and a very low base value. Thus it is reasonable to evaluate the VSL at $p = 0$, which we do below.

As is often done in the VSL literature, we take lifetime consumption to be a multiple m of current time- t consumption C_t . Annual consumption when dead is assumed to be a fraction $\varepsilon \ll 1$ of what it was when alive, so the equivalent “lifetime” consumption when dead is $m\varepsilon C_t$. Then $u(w) = u(mC)$ and $v(w) = u(m\varepsilon C)$, so eqn. (2) becomes:

$$\text{VSL} = \frac{mC}{1-\varepsilon} [1 - \varepsilon^{1-\eta}]. \quad (3)$$

Empirical studies usually express the VSL as a multiple s of lifetime consumption, with s in the range of 3 to 10. Because lifetime consumption is multiple m of annual consumption C ,

to \$9.3 million in 2015 dollars. In 2015, U.S. per capita consumption was about \$38,000, which corresponds to lifetime consumption of roughly \$1,500,000 (i.e., a factor of 40). The average estimate of the VSL is around 7 times lifetime consumption, which yields \$10.5 million, and is close to the EPA number. A recent study by Rohlfs, Sullivan and Kniesner (2015), using airbag regulations during the 1990s and used car prices, finds a median VSL between \$9 and \$11 million.

⁹This is what Pratt and Zeckhauser (1996) call the “dead anyway” effect. The VSL is also increasing in wealth w , because a wealthier individual has more utility to lose should she die.

smC is the VSL that goes on LHS of eqn. (3). Cancelling mC from both sides, ε is

$$\varepsilon = [s(\eta - 1) + 1]^{\frac{1}{1-\eta}}. \quad (4)$$

Although estimates of s vary, the most common ones put it at around 7, and we will use that value in what follows. Thus if $\eta = 2$, $\varepsilon = 1/(s + 1) = .125$, i.e, death is equivalent in welfare terms to an 88% drop in consumption. If $\eta = 3$, $\varepsilon = .27$, and if $\eta = 4$, $\varepsilon = .42$. Note that ε is increasing in η because a larger value of η implies a larger utility loss from any given reduction in C . The utility loss in (2) is fixed by the VSL, so the welfare-equivalent reduction in consumption upon death must be smaller if η is larger.

2.3 Welfare Loss

Initially, static welfare is $U_0 = N_0 C_0^{1-\eta}/(1-\eta)$. With $\eta > 1$ so $U_0 < 0$, it might appear that a drop in N_0 raises welfare, but that is not the case. Setting $C_0 = N_0 = 1$, if a fraction ψ of the population dies, welfare becomes

$$U_N = \frac{1}{1-\eta} [(1-\psi) + \psi \varepsilon^{1-\eta}] < U_0$$

Note that the first term, $(1-\psi)/(1-\eta)$, is the welfare of the people still living, and $\psi \varepsilon^{1-\eta}/(1-\eta)$ is the welfare of those that have died. This last term ensures that $U_N < U_0$. From (4), $\varepsilon^{1-\eta} = s(\eta - 1) + 1$, so $U_N = \frac{1}{1-\eta} [1 + (\eta - 1)\psi s] < U_0$ as long as $s > 0$. And if instead consumption falls by a fraction ϕ , welfare becomes $U_C = (1-\phi)^{1-\eta}/(1-\eta) < U_0$.

We can also compare the welfare loss for an event that kills a fraction ϕ of the population with the loss for an event that reduces the consumption of everyone by the same fraction ϕ . Let LR denote the ratio of the first welfare loss to the second. Using (4), the (annual) welfare loss for each person who dies is $L_d = u(\varepsilon) - u(C_0) = [\varepsilon^{1-\eta} - 1]/(1-\eta) = -s$. Since a fraction ϕ die, the total loss is $-\phi s$. This is independent of η , because the drop in consumption (from 1 to ε) is constrained by (3) to yield a welfare loss equal to the VSL. If instead the consumption of everyone falls by the same fraction ϕ , the welfare loss is $L_c = [(1-\phi)^{1-\eta} - 1]/(1-\eta)$. (This loss depends on η because the drop in consumption is fixed.)

The ratio of the welfare loss from death to the loss from reduced consumption is then:

$$\text{LR} = L_d/L_c = \frac{(\eta - 1)\phi s}{(1-\phi)^{1-\eta} - 1} \quad (5)$$

If $s = 7$ and $\eta = 2$, for “low” values of ϕ , e.g., $\phi = .1$, the welfare loss from death is more than six times the welfare loss from destruction. (As $\phi \rightarrow 0$, $\text{LR} \rightarrow s$, i.e., LR approaches the VSL.) The ratio is large because compared to a drop in consumption, death causes a

much larger loss of utility, albeit for only a fraction ϕ of the population. This result depends on the VSL parameter, but $s = 7$ is near the center of the range of VSL estimates.

How large would the drop in consumption for all members of society have to be to yield the same welfare loss as an event that kills a fraction ϕ ? Denoting the equivalent drop in consumption by ϕ_c , using (4), and setting $u[(1 - \phi_c)C_0] - u(C_0) = \phi[u(\varepsilon C_0) - u(C_0)]$,

$$\phi_c = 1 - [s\phi(\eta - 1) + 1]^{\frac{1}{1-\eta}} \quad (6)$$

If $s = 7$ and $\phi = .05$, when $\eta = 2$, $\phi_c = .26$, and when $\eta = 4$, $\phi_c = .21$. (Setting $\phi = .05$ corresponds roughly to the Spanish Flu of 1918-19, which killed some 4 to 5 percent of the population in the U.S. and Europe.) If $\phi = .1$, $\phi_c = .41$ when $\eta = 2$ and $.31$ when $\eta = 4$. So ϕ_c is 3 to 5 times as large as ϕ when $\phi \leq .10$. However, the multiple is smaller when ϕ is large, and can be less than 1 if ϕ is sufficiently large. (If $\phi = .8$ and $\eta = 4$, $\phi_c = .62$.) Once again, the VSL constrains L_d , but L_c is unconstrained as ϕ_c increases.

3 Dynamic Model

We turn now to a fully dynamic model, and begin by specifying the processes for per-capita consumption C_t and population N_t . We assume that catastrophic events that reduce consumption occur as Poisson arrivals with mean arrival rate λ_c , and the impact of the k th arrival, ϕ_k , is i.i.d. across realizations k . Thus the process for consumption is:

$$c_t = \log C_t = gt - \sum_{k=1}^{Q(t)} \phi_k \quad (7)$$

where g is the normal growth rate of per capita consumption, and $Q(t)$ is a Poisson counting process with known mean arrival rate λ_c . When the k th catastrophic event occurs, per-capita consumption is multiplied by the random variable $e^{-\phi_k}$. (We are scaling consumption so that $C_0 = 1$ and hence $c_0 = 0$.) We can then define the cumulant-generating function (CGF),

$$\kappa_{C,t}(\theta) \equiv \log \mathbb{E} e^{c_t\theta} \equiv \log \mathbb{E} C_t^\theta.$$

We showed in Martin and Pindyck (2015) that since log consumption follows a Lévy process, the CGF is linear in t —that is, $\kappa_{C,t}(\theta) = \kappa_{C,1}(\theta)t$ —and we can write

$$\kappa_C(\theta) \equiv \kappa_{C,1}(\theta) = g\theta + \lambda_c (\mathbb{E} e^{-\theta\phi} - 1) , \quad (8)$$

where ϕ is a representative of the (i.i.d.) ϕ_k .¹⁰

¹⁰We could easily make the normal growth of consumption stochastic, i.e., replace gt in (7) by any Lévy process g_t (e.g., a Brownian motion). Then $g\theta$ in (8) becomes $g(\theta)$, the CGF of g_1 . The normal rate of population growth can likewise be stochastic. For simplicity, we limit uncertainty to catastrophic events.

We assume that N_t , the population alive at time t , follows an analogous process:

$$\log N_t = nt - \sum_{k=1}^{X(t)} \psi_k, \quad (9)$$

where n is the natural rate of population growth, $X(t)$ is a Poisson counting process with arrival rate λ_d that determines how many “death events” have happened by time t , and ψ_k measures the size of the k th event. Thus when the k th event occurs, N_t is multiplied by the random variable $e^{-\psi_k}$. We scale N_t so that $N_0 = 1$, and we define the corresponding CGF,

$$\kappa_{N,t}(\theta) = \log \mathbb{E} e^{\theta \log N_t}.$$

Once again, the CGF scales linearly with t , so we write

$$\kappa_N(\theta) \equiv \kappa_{N,1}(\theta) = n\theta + \lambda_d (\mathbb{E} e^{-\theta\psi} - 1) \quad (10)$$

where ψ is a representative of the (i.i.d.) ψ_k . Thus

$$\mathbb{E} N_t = e^{\kappa_N(1)t} = e^{nt + \lambda_d(\mathbb{E} e^{-\psi} - 1)t} = e^{nt - \tilde{\lambda}_d t}.$$

The term $\tilde{\lambda}_d = \lambda_d(1 - \mathbb{E} e^{-\psi})$ can be interpreted as the adjusted death arrival rate, taking into account that only a fraction of the population dies when there is a death event. Depending on the relative sizes of n and $\tilde{\lambda}_d$, the population may grow or shrink in expectation: if disasters are sufficiently frequent (large λ_d) or cataclysmic (small $\mathbb{E} e^{-\psi}$), and population growth n is sufficiently low, then the expected population is declining in t , in which case $\kappa_N(1) < 0$.

We will use an asterisk to denote the absence of catastrophes. So if there are no death catastrophes ($\lambda_d = 0$), the population evolves as $N_t^* = e^{nt}$, and the CGF for N_t is $\kappa_N^*(\theta) = n\theta$.

Recall that death implies a drop in consumption to a fraction ε of what it was before, i.e., from C_t to to εC_t . So if no catastrophes are averted, total welfare is

$$V = \mathbb{E} \left\{ \int_0^\infty e^{-\delta t} \left[\frac{N_t C_t^{1-\eta}}{1-\eta} + \frac{(N_t^* - N_t) \varepsilon^{1-\eta} C_t^{1-\eta}}{1-\eta} \right] dt \right\}, \quad (11)$$

where $(N_t^* - N_t)$ is the number of people that have died.

Because C_t is an exponential Lévy process, it evolves independently of N_t . This makes it easy to calculate the expectations inside the integral above. In particular,

$$\mathbb{E}(N_t C_t^{1-\eta}) = \mathbb{E} N_t \mathbb{E} C_t^{1-\eta} = e^{\kappa_N(1)t} \cdot e^{\kappa_C(1-\eta)t}$$

and

$$\mathbb{E} [(N_t^* - N_t) \varepsilon^{1-\eta} C_t^{1-\eta}] = (e^{\kappa_N^*(1)t} - e^{\kappa_N(1)t}) \varepsilon^{1-\eta} e^{\kappa_C(1-\eta)t}.$$

Substituting these expressions into the integral above,

$$\begin{aligned}
V &= \frac{1}{1-\eta} \left\{ \frac{1}{\delta - \kappa_N(1) - \kappa_C(1-\eta)} + \frac{\varepsilon^{1-\eta}}{\delta - \kappa_N^*(1) - \kappa_C(1-\eta)} - \frac{\varepsilon^{1-\eta}}{\delta - \kappa_N(1) - \kappa_C(1-\eta)} \right\} \\
&= \frac{1}{1-\eta} \left\{ \frac{1 - \varepsilon^{1-\eta}}{\delta - \kappa_N(1) - \kappa_C(1-\eta)} + \frac{\varepsilon^{1-\eta}}{\delta - \kappa_N^*(1) - \kappa_C(1-\eta)} \right\}. \tag{12}
\end{aligned}$$

To interpret this equation, note that the second term captures the welfare associated with the guaranteed consumption stream εC_t (which is received even after death). We can think of $\delta - \kappa_N^*(1) = \delta - n$ as the social rate of time preference (i.e., discount rate on future utility) associated with this consumption stream: it adjusts, via $\kappa_N^*(1) = n$, for the fact that the population is increasing (so that the larger is n , the greater is future total utility, and thus the lower is the discount rate). The first term captures the welfare associated with the ongoing consumption stream $(1 - \varepsilon)C_t$ received by those who are still alive (i.e., those alive receive $(1 - \varepsilon + \varepsilon)C = C$). Since there is a risk of death, this latter consumption stream is discounted at the higher rate $\delta - \kappa_N(1) > \delta - \kappa_N^*(1)$.¹¹

We want to find the WTPs to avert each of the two types of catastrophes individually, and the WTP to avert both. If consumption C_t is exposed to a “destruction” catastrophe, then averting that catastrophe corresponds to setting $\lambda_c = 0$, and thus replacing $\kappa_C(1 - \eta)$ by $\kappa_C^*(1 - \eta) \equiv g(1 - \eta)$. What is the WTP to avert this type of catastrophe? If it is averted at the cost of a permanent loss of a fraction w_c of consumption, then welfare is

$$V_c = \frac{(1 - w_c)^{1-\eta}}{1 - \eta} \left\{ \frac{1 - \varepsilon^{1-\eta}}{\delta - \kappa_N(1) - \kappa_C^*(1 - \eta)} + \frac{\varepsilon^{1-\eta}}{\delta - \kappa_N^*(1) - \kappa_C^*(1 - \eta)} \right\}. \tag{13}$$

The WTP to avoid this catastrophe is the value of w_c that equates (12) and (13):

$$(1 - w_c)^{1-\eta} = A \times B \times C \tag{14}$$

where

$$\begin{aligned}
A &= \frac{\delta - \kappa_N^*(1) - \kappa_C^*(1 - \eta)}{\delta - \kappa_N^*(1) - \kappa_C(1 - \eta)} \\
B &= \frac{\delta - \kappa_N(1) - \kappa_C^*(1 - \eta)}{\delta - \kappa_N(1)\varepsilon^{1-\eta} - (1 - \varepsilon^{1-\eta})\kappa_N^*(1) - \kappa_C^*(1 - \eta)} \\
C &= \frac{\delta - \kappa_N(1)\varepsilon^{1-\eta} - (1 - \varepsilon^{1-\eta})\kappa_N^*(1) - \kappa_C(1 - \eta)}{\delta - \kappa_N(1) - \kappa_C(1 - \eta)}.
\end{aligned}$$

¹¹The expressions $1/[\delta - \kappa_N(1) - \kappa_C(1 - \eta)]$ and $1/[\delta - \kappa_N^*(1) - \kappa_C(1 - \eta)]$ can be interpreted as the valuation ratios of the two consumption streams, as shown by Martin (2013).

What is the WTP to avoid only the death catastrophe? If a fraction w_d of consumption is sacrificed to avert this catastrophe, welfare is

$$V_d = \mathbb{E} \left\{ \int_0^\infty e^{-\delta t} \frac{(1-w_d)^{1-\eta} N_t^* C_t^{1-\eta}}{1-\eta} dt \right\} = \frac{(1-w_d)^{1-\eta}}{1-\eta} \left\{ \frac{1}{\delta - \kappa_N^*(1) - \kappa_C(1-\eta)} \right\}. \quad (15)$$

Equating (12) and (15), w_d satisfies

$$(1-w_d)^{1-\eta} = \frac{\delta - \kappa_N(1)\varepsilon^{1-\eta} - (1-\varepsilon^{1-\eta})\kappa_N^*(1) - \kappa_C(1-\eta)}{\delta - \kappa_N(1) - \kappa_C(1-\eta)}. \quad (16)$$

Lastly, if a fraction $w_{c,d}$ of consumption is sacrificed to avert both types of catastrophes, welfare is

$$V_{c,d} = \frac{(1-w_{c,d})^{1-\eta}}{1-\eta} \left\{ \frac{1}{\delta - \kappa_N^*(1) - \kappa_C^*(1-\eta)} \right\}. \quad (17)$$

Equating (12) and (17), $w_{c,d}$ satisfies

$$(1-w_{c,d})^{1-\eta} = \frac{\delta - \kappa_N^*(1) - \kappa_C^*(1-\eta)}{\delta - \kappa_N^*(1) - \kappa_C(1-\eta)} \cdot \frac{\delta - \kappa_N(1)\varepsilon^{1-\eta} - (1-\varepsilon^{1-\eta})\kappa_N^*(1) - \kappa_C(1-\eta)}{\delta - \kappa_N(1) - \kappa_C(1-\eta)}. \quad (18)$$

Now we can compare expressions (14), (16), and (18) for the WTPs w_c , w_d , and $w_{c,d}$, respectively. Using the three factors A , B , and C that are multiplied together in (14), we have $(1-w_c)^{1-\eta} = ABC$, $(1-w_d)^{1-\eta} = C$, and $(1-w_{c,d})^{1-\eta} = AC$.

We assumed that $\eta > 1$, so (i) $A > 1$, (ii) $B < 1 < C$, and (iii) $BC > 1$, using the facts that $\kappa_C^*(1-\eta) < \kappa_C(1-\eta)$, $\kappa_N^*(1) > \kappa_N(1)$, and $\varepsilon^{1-\eta} > 1$ (for $\eta > 1$).¹² These inequalities imply that $w_{c,d} > \max\{w_c, w_d\}$ and, more interestingly (using the fact that $BC > 1$),

$$(1-w_{c,d})^{1-\eta} < (1-w_c)^{1-\eta}(1-w_d)^{1-\eta}, \quad \text{or (equivalently)} \quad w_{c,d} < w_c + w_d - w_c w_d.$$

Note that the WTPs do not add, and the two catastrophes are interdependent.¹³

4 Applying the Model

Our results so far are quite general. Here we show how the model can be parameterized and used to evaluate the WTPs to avert one or both types of catastrophes.

¹²This implies, for example, that $\kappa_N(1)\varepsilon^{1-\eta} + \kappa_N^*(1)(1-\varepsilon^{1-\eta}) < \kappa_N(1)$.

¹³The inequality in Martin and Pindyck (2015) (page 2957) that applied to two types of *consumption* catastrophes (1 and 2) said only that $w_{1,2} < w_1 + w_2$; in fact the stronger result that $w_{1,2} < w_1 + w_2 - w_1 w_2$ also applies in that context.

The CGFs of eqns. (8) and (10) apply to *any* probability distributions for the impacts ϕ and ψ . In order to explore some numerical examples, we will assume that ϕ and ψ are exponentially distributed:¹⁴

$$f_\phi(x) = \beta_c e^{-\beta_c x} \quad \text{for } x \geq 0 \quad \text{and} \quad f_\psi(x) = \beta_d e^{-\beta_d x} \quad \text{for } x \geq 0.$$

Note that $\mathbb{E}(\phi) = 1/\beta_c$ and $\mathbb{E}(z_c) = \mathbb{E}e^{-\phi} = \beta_c/(\beta_c + 1)$, and similarly for ψ and z_d . Thus large values of β_c and β_d imply small expected impacts, i.e., small values of $\mathbb{E}(\phi)$ and $\mathbb{E}(\psi)$ and large values of $\mathbb{E}(z_c)$ and $\mathbb{E}(z_d)$. Given these distributions for ϕ and ψ , the CGFs are

$$\begin{aligned} \kappa_C(1 - \eta) &= g(1 - \eta) - \frac{\lambda_c(1 - \eta)}{\beta_c + (1 - \eta)} \\ \kappa_C^*(1 - \eta) &= g(1 - \eta) \\ \kappa_N(1) &= n - \frac{\lambda_d}{\beta_d + 1} \\ \kappa_N^*(1) &= n \end{aligned}$$

We can substitute these CGFs into the expressions for the three factors A , B , and C that are multiplied together in (14) to determine each of the WTPs. Define $\rho \equiv \delta - n + g(\eta - 1)$. One can think of ρ as the discount rate on future total consumption (as opposed to $\delta - n$, the discount rate on future utility). It accounts for the growth of total consumption via n and also the decline in marginal utility of per capita consumption via g . Also define

$$\lambda'_c \equiv \lambda_c(\eta - 1)/(\beta_c + 1 - \eta) \tag{19}$$

$$\lambda'_d \equiv \lambda_d/(\beta_d + 1) \tag{20}$$

One can think of λ'_c and λ'_d as risk- and impact-adjusted arrival rates for the two types of catastrophes. For example, increasing β_d reduces the expected impact of a death catastrophe, which is welfare-equivalent to reducing its expected arrival rate. And λ'_c further adjusts for risk aversion; increasing η increases the utility loss from a catastrophe that reduces consumption, which is welfare-equivalent to increasing the expected arrival rate.

Substituting in the CGFs, ρ , λ'_c and λ'_d , the factors A , B , and C in eqn. (14) become:

$$\begin{aligned} A &= \frac{\rho}{\rho - \lambda'_c} \\ B &= \frac{\rho + \lambda'_d}{\rho + \lambda'_d \varepsilon^{1-\eta}} \\ C &= \frac{\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c}{\rho + \lambda'_d - \lambda'_c} \end{aligned}$$

¹⁴This is equivalent to assuming that $z_c = e^{-\phi}$ and $z_d = e^{-\psi}$ follow the power distributions $b(z_c) = \beta_c z_c^{\beta_c - 1}$ for $0 \leq z_c \leq 1$ and $b(z_d) = \beta_d z_d^{\beta_d - 1}$ for $0 \leq z_d \leq 1$.

4.1 WTPs

Recall that $(1 - w_c)^{1-\eta} = ABC$, $(1 - w_d)^{1-\eta} = C$, and $(1 - w_{c,d})^{1-\eta} = AC$. Using the expressions above for A , B , and C , the WTPs are:

$$w_c = 1 - \left[\frac{(\rho - \lambda'_c)(\rho + \lambda'_d \varepsilon^{1-\eta})(\rho + \lambda'_d - \lambda'_c)}{\rho(\rho + \lambda'_d)(\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c)} \right]^{\frac{1}{\eta-1}} \quad (21)$$

$$w_d = 1 - \left[\frac{(\rho + \lambda'_d - \lambda'_c)}{(\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c)} \right]^{\frac{1}{\eta-1}} \quad (22)$$

$$w_{c,d} = 1 - \left[\frac{(\rho - \lambda'_c)(\rho + \lambda'_d - \lambda'_c)}{\rho(\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c)} \right]^{\frac{1}{\eta-1}} \quad (23)$$

From eqn. (4), $\varepsilon^{1-\eta} = s(\eta - 1) + 1$. Note that $w_c < 1$ (and $w_{c,d} < 1$) only if $\rho > \lambda'_c$. Recall that ρ is a discount rate on future consumption, but ignoring consumption catastrophes. One can think of λ'_c as a depreciation rate that accounts for the risk- and impact-adjusted arrival rate of consumption catastrophes, and thereby reduces the expected future welfare from consumption. If the net discount rate $\rho - \lambda'_c$ is zero, the welfare loss from the catastrophes is unbounded, pushing the WTP to avoid the catastrophes to one.

To illustrate the characteristics of the WTPs, we can compare two catastrophes that are identical except that one causes a drop in consumption and the other causes death. We will assume that both catastrophes have the same mean arrival rates and same impact distributions. We set $\lambda_c = \lambda_d = .05$, so that each type of catastrophe has a 50 percent chance of occurring in the next 14 years (and there is a 75 percent chance of at least one type occurring in the next 14 years). We set $\beta_c = \beta_d = 10$, so should an event occur, the mean impact would be roughly a 10-percent reduction in consumption or a 10-percent fatality rate, and there is a 35-percent probability of an impact greater than 10 percent.

Figure 1 shows the WTPs as functions of η for two values of the population growth rate, $n = .02$ and $n = 0$. The other parameter values are $\delta = .02$, $g = .02$, and $s = 7$. Note that as η is increased, the effective discount rate becomes larger and the WTPs decline. Setting $n = 0$ makes the WTPs lower for every value of η (because the effective discount rate is higher). Also, $w_d > w_c$ for all values of η , but the percentage difference declines with η (for all η if $n = 0$ and for $\eta > 2.5$ for $n = .02$) because w_d is constrained by the VSL.

Figure 2 shows the WTPs as functions of the VSL parameter s , for $\eta = 2$ and 4. (The other parameters values are $\delta = .03$ and $n = g = .02$.) When $s = 0$ there is no value for lost lives, so $w_d = 0$ and $w_{c,d} = w_c$. Once again, increasing η reduces all of the WTPs, because it increases the effective discount rate. As s is increased, w_d and $w_{c,d}$ increase, but note

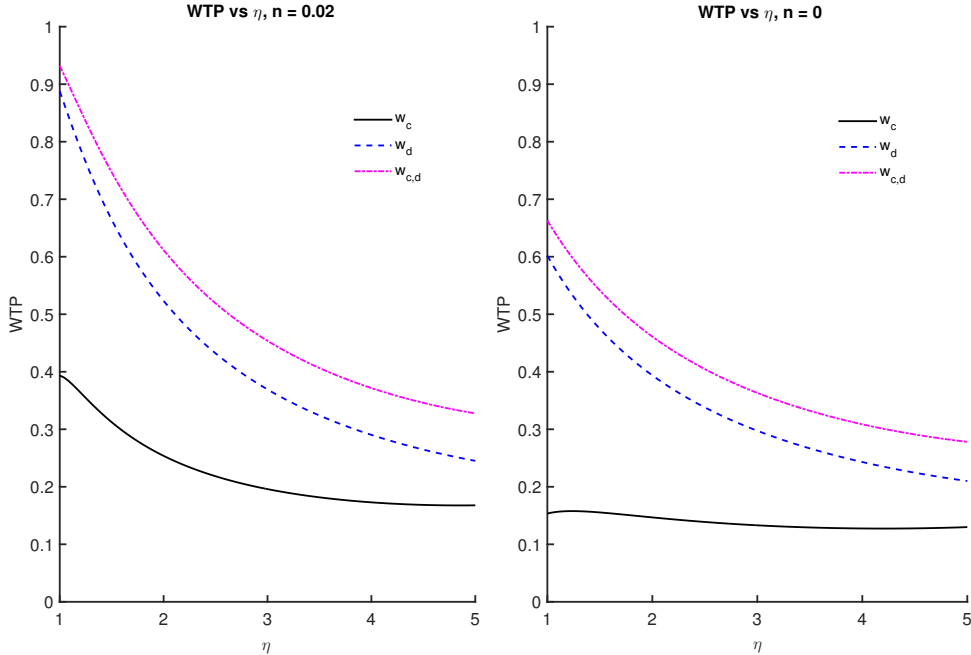


Figure 1: Dependence of WTP to Avert One or Both Types of Catastrophes on η

that w_c also increases (though not by as much). This is because life and consumption are complementary; If nothing is done to avoid the death catastrophe, increasing s has the same effect as increasing the *ex ante* probability of death (which we took to be zero). This in turn increases the WTP to avoid a drop in consumption because expected lifetimes are shorter.

4.2 Example

As an example, we re-consider two of the catastrophes examined in Martin and Pindyck (2015): a “mega-virus” and nuclear terrorism. A mega-virus is clearly a type of catastrophe that causes death as opposed to destruction. For example, the Spanish Flu pandemic of 1918–1919 killed about 4 to 5 percent of the populations of Europe and the U.S., but had a minimal impact on GDP and the consumption of those that survived.¹⁵ As for nuclear terrorism, a Hiroshima-grade bomb detonated in a major city could kill several hundred thousand people, but its biggest impact would be economic in nature. It would cause a major shock to the capital stock and GDP from a worldwide reduction in trade and economic activity, and vast resources would have to be devoted to averting further attacks.¹⁶ Thus in this example we

¹⁵Some studies suggest that because of modern travel, another pandemic is likely in the next couple of decades, and could kill a larger fraction of the population. See, e.g., Byrne (2008) and Kilbourne (2008).

¹⁶See, e.g., Allison (2004) and Ackerman and Potter (2008).




Figure 2: Dependence of WTP to Avert One or Both Types of Catastrophes on VSL Parameter s

will treat nuclear terrorism as a pure “consumption” catastrophe.

In Martin and Pindyck (2015), we presented rough estimates of the mean arrival rates and impact parameters for these two catastrophes. For a mega-virus, we argued that a reasonable value for the mean arrival rate is $\lambda_d = .02$, i.e., there is a roughly 10 percent chance of a pandemic occurring in the next 10 years, and a 50 percent chance of one occurring in the next 35 years. We will assume that the mean impact is the death of 5 percent of the population (corresponding to the 1918-1919 Spanish Flu), which implies that $\beta_d = 20$. For nuclear terrorism, we argued that while opinions differ, a reasonable value for the mean arrival rate is $\lambda_c = .04$, which implies a 50-percent chance of an occurrence in the next 17 years, and that the mean impact would be a 5.5 percent drop in consumption, so that $\beta_c = 17$.

We will calculate the WTP to avert each of these two catastrophes and the WTP to avert both. We set $\delta = n = g = .02$, and calculate the WTPs for $\eta = 2$ and $\eta = 4$. Then if $\eta = 2$, $\rho = .02$, and from eqns. (19) and (20), $\lambda'_c = .0025$ and $\lambda'_d = .0010$. If $\eta = 4$, $\rho = .06$, $\lambda'_c = .0086$ and $\lambda'_d = .0010$. Using eqns. (21), (22) and (23), and setting $s = 7$ so that $\varepsilon^{1-\eta} = s(\eta - 1) + 1 = 7\eta - 6$, we can calculate the WTPs for each value of η .

Table 1 shows these three WTPs and how they vary as we change some of the parameter values. (Note that the constraint $w_{c,d} < w_c + w_d - w_c w_d$ always holds.) In every case, increasing η from 2 to 4 causes every WTP to fall. Increasing η has two off-setting effects. First, it increases the risk- (and impact-) adjusted arrival rate for the consumption catastrophe, λ'_c

Table 1: WTPS AND NET WELFARE FOR AVERTING TWO TYPES OF CATASTROPHES

Parameters		w_c	w_d	$w_{c,d}$	W_0	W_c	W_d	$W_{c,d}$
Base Case	$\eta = 2$.1527	.2654	.3572	-77.8	-69.4	-60.2	-55.4*
	$\eta = 4$.0626	.1022	.1472	-8.96	-8.61	-7.56	-7.55*
$n = 0$	$\eta = 2$.0710	.1478	.2010	-31.3	-30.6	-28.1	-27.7*
	$\eta = 4$.0445	.0781	.1123	-5.96	-6.06	-5.44*	-5.67
$s = 3$	$\eta = 2$.1390	.1341	.2423	-66.0	-59.8	-60.2	-55.4*
	$\eta = 4$.0564	.0493	.0969	-7.54	-7.39*	-7.56	-7.56
$s = 10$	$\eta = 2$.1605	.3404	.4229	-86.6	-76.6	-60.2	-55.4*
	$\eta = 4$.0661	.1351	.1784	-10.02	-9.52	-7.56	-7.55*
$\lambda_d = 0$	$\eta = 2$.1250	0	.1250	-57.1	-52.6*	-60.2	-55.4
	$\eta = 4$.0501	0	.0501	-6.48	-6.47*	-7.56	-7.55

Note: Base case parameter values: $\delta = g = n = .02$ and $s = 7$. Also, $\lambda_c = .04$, $\beta_c = 17$, $\lambda_d = .02$, and $\beta_d = 20$. Costs of averting: $\tau_c = \tau_d = .05$. Asterisk indicates maximum net welfare.

(but has no effect on the impact-adjusted arrival rate for the death catastrophe, λ'_d), which increases all three WTPs. Second, it increases ρ , the effective discount rate on future total consumption $N_t C_t$, and this reduces the WTPs. Thus in general, $\partial w / \partial \eta$ is indeterminate.

Which of these catastrophes (if any) should be averted? To answer that question we need to know the cost of averting each one, expressed as a permanent tax on consumption at a rate just sufficient to pay for whatever is required to avert the catastrophe. We denote these costs by τ_c and τ_d . In our earlier work, we argued that a reasonable value for the τ_d needed to avert a mega-virus is about 0.02, and for the τ_c needed to avert nuclear terrorism is about .03. Some have claimed, however, that modern transportation, mobility, and access to enriched uranium have made both costs much higher, so we will set $\tau_d = \tau_c = .05$.

To determine the optimal policy, we calculate the net (of taxes) welfare of doing nothing (W_0), averting only nuclear terrorism (W_c), averting only the mega-virus (W_d), and averting both ($W_{c,d}$).¹⁷ Using eqns. (12), (13), (15), (17) and the expressions above for the CGFs, the net welfare for each policy is:

$$W_0 = \frac{1}{1-\eta} \left[\frac{-s(\eta-1)}{\rho + \lambda'_d - \lambda'_c} + \frac{s(\eta-1)+1}{\rho - \lambda'_c} \right] \quad (24)$$

$$W_c = \frac{(1-\tau_c)^{1-\eta}}{1-\eta} \left[\frac{\rho + \lambda'_d(s(\eta-1)+1)}{\rho(\rho + \lambda'_d)} \right] \quad (25)$$

¹⁷Because a death catastrophe divides the population into two groups with different levels of consumption, Result 2 of Martin and Pindyck (2015) does not hold. Thus we must calculate the net welfare for each possible policy to find the optimal one.

$$W_d = \frac{(1 - \tau_d)^{1-\eta}}{(1 - \eta)(\rho - \lambda'_c)} \quad (26)$$

$$W_{c,d} = \frac{(1 - \tau_c)^{1-\eta}(1 - \tau_d)^{1-\eta}}{(1 - \eta)\rho} \quad (27)$$

Note from (25) that an increase in λ_d (and thus λ'_d) *raises* W_c , net welfare when the consumption catastrophe is eliminated. This is an example of the “dead anyway” effect; expected lifetimes are shorter, so it is less “costly” to devote resources to protecting consumption. Conversely, from (26), an increase in λ_c *reduces* W_d ; the net benefit of avoiding death is smaller when future consumption is more likely to fall.

Table 1 shows W_0 , W_c , W_d , and $W_{c,d}$ for our example of nuclear terrorism and a virus, with net welfare for the optimal policy noted by an asterisk. Observe that it is usually optimal to avert both catastrophes. One exception is when we set the VSL parameter $s = 3$, which reduces the value of averting the virus, so if $\eta = 4$ it becomes optimal to only avert nuclear terrorism. (The same is true for the trivial case of $\lambda_d = 0$.)

5 Conclusions

In our earlier work we showed that decisions to avert major catastrophes are interdependent, and we showed how to determine which ones to avert. We assumed, however, that catastrophes only cause a reduction in the flow of consumption. Here we show how catastrophes that cause death can be incorporated into our framework.

Our approach is to treat death as a welfare-equivalent reduction in consumption, not to zero but to a small value ε (so we can retain the use of CRRA utility). We rely on an estimate of the VSL to determine the utility loss from death, and thus the value of ε . We can then find expressions for the WTPs to avert a “death” catastrophe (w_d), a consumption catastrophe (w_c), and both ($w_{c,d}$). Our results hold for *any* probability distributions for the random reductions ϕ and ψ of $\log C$ and $\log N$. However, by assuming that ϕ and ψ are exponentially distributed, these WTPs become very easy to calculate.

Using middle-of-the-road estimates of the VSL, we find the utility loss from a “death” catastrophe is large, as is the WTP to avert such a catastrophe. Although our model can be applied to any number of catastrophes (of both the consumption and death variety), we illustrated our results with just two, nuclear terrorism (which we treated as a consumption catastrophe) and a mega-virus (a death catastrophe), and showed how the optimal choice of which ones to avert depends on the VSL and other parameters.

References

- Ackerman, Gary, and William C. Potter.** 2008. “Catastrophic Nuclear Terrorism: A Preventable Peril.” In *Global Catastrophic Risks*, ed. Nick Bostrom and Milan Ćirković. Oxford University Press.
- Allison, Graham.** 2004. *Nuclear Terrorism: The Ultimate Preventable Catastrophe*. Henry Holt & Company.
- Andersson, Henrik, and Nicolas Treich.** 2009. “The Value of a Statistical Life.” Toulouse School of Economics Working Paper 09-015.
- Arrow, Kenneth J., and Marcel Priebisch.** 2014. “Bliss, Catastrophe, and Rational Policy.” *Environmental and Resource Economics*, 58: 491–509.
- Arthur, W.B.** 1981. “The Economics of Risks to Life.” *American Economic Review*, 71(1): 54–64.
- Barro, Robert J., and José Ursúa.** 2008. “Macroeconomic Contractions since 1870.” *Brookings Papers on Economic Activity*, 255–335.
- Barro, Robert J., and Tao Jin.** 2011. “On the Size Distribution of Macroeconomic Disasters.” *Econometrica*, 79(5): 1567–1589.
- Becker, Gary S., Tomas J. Philipson, and Rodrigo R. Soares.** 2005. “The Quantity and Quality of Life and the Evolution of World Inequality.” *American Economic Review*, 95(1): 277–291.
- Byrne, Joseph,** ed. 2008. *Encyclopedia of Pestilence, Pandemics, and Plagues*. Greenwood.
- Cropper, Maureen L., and Frances G. Sussman.** 1990. “Valuing Future Risks to Life.” *Journal of Environmental Economics and Management*, 19: 160–174.
- Hall, Robert E., and Charles I. Jones.** 2007. “The Value of Life and the Rise in Health Spending.” *Quarterly Journal of Economics*, 122: 39–72.
- Hammit, James K.** 2007. “Statistical vs. Identified Lives in Benefit-Cost Analysis.” *Journal of Risk and Uncertainty*, 35: 45–66.
- Hammit, James K., and Nicolas Treich.** 2007. “Valuing Changes in Mortality Risk: Lives Saved Versus Life Years Saved.” *Review of Environmental Economics and Policy*, 1(2): 228–240.
- Kaplow, Louis.** 2005. “The Value of a Statistical Life and the Coefficient of Relative Risk Aversion.” *Journal of Risk and Uncertainty*, 31(1): 23–34.

- Kilbourne, Edwin Dennis.** 2008. “Plagues and Pandemics: Past, Present, and Future.” In *Global Catastrophic Risks*, ed. Nick Bostrom and Milan Ćirković. Oxford 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: 745–773.
- Martin, Ian W.R., and Robert S. Pindyck.** 2015. “Averting Catastrophes: The Strange Economics of Scylla and Charybdis.” *American Economic Review*, 105(10): 2947–2985.
- Nordhaus, William D.** 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- Pindyck, Robert S.** 2012. “Uncertain Outcomes and Climate Change Policy.” *Journal of Environmental Economics and Management*, 63: 289–303.
- Pratt, John W., and Richard J. Zeckhauser.** 1996. “Willingness to Pay and the Distribution of Risk and Wealth.” *Journal of Political Economy*, 104(4): 747–763.
- Rohlf, Chris, Ryan Sullivan, and Thomas Kniesner.** 2015. “New Estimates of the Value of a Statistical Life Using Air Bag Regulations as a Quasi-Experiment.” *American Economic Journal: Economic Policy*, 7(1): 331–359.
- Rosen, Sherwin.** 1988. “The Value of Changes in Life Expectancy.” *Journal of Risk and Uncertainty*, 1: 285–304.
- Stern, Nicholas.** 2007. *The Economics of Climate Change: The Stern Review*. Cambridge University Press.
- U.S. Environmental Protection Agency.** 2014. “Guidelines for Preparing Economic Analyses.” U.S. Government Report.
- Viscusi, W. Kip.** 1993. “The Value of Risks to Life and Health.” *Journal of Economic Literature*, 31(4): 1912–1946.
- Viscusi, W. Kip, and Joseph E. Aldy.** 2003. “The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World.” *Journal of Risk and Uncertainty*, 27: 5–76.