WELFARE COSTS OF CATASTROPHES:
LOST CONSUMPTION AND LOST LIVES

Ian W.R. Martin
Robert S. Pindyck

Welfare Costs of Catastrophes: Lost Consumption and Lost Lives
Ian W.R. Martin and Robert S. Pindyck
NBER Working Paper No. 26068
July 2019
JEL No. D81,H5,Q5,Q58

ABSTRACT

Most of the literature on the economics of catastrophes assumes that such events cause a reduction in the stream of consumption, as opposed to widespread fatalities. Here we show how to incorporate death in a model of catastrophe avoidance, and how a catastrophic loss of life can be expressed as a welfare-equivalent drop in wealth or consumption. We examine how potential fatalities affect the policy interdependence of catastrophic events and "willingness to pay" (WTP) to avoid them. Using estimates of the "value of a statistical life" (VSL), we find the WTP to avoid major pandemics, and show it is large (10% or more of annual consumption) and partly driven by the risk of macroeconomic contractions. Likewise, the risk of pandemics significantly increases the WTP to reduce consumption risk. Our work links the VSL and consumption disaster literatures.

Ian W.R. 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

The welfare costs of unpredictable catastrophic events have been addressed in a variety of studies. In most, the events are "generic," in that the cause is unspecified, and each event reduces GDP and consumption by some (random) amount. One can think of these events as destroying a random part of the capital stock or reducing its productivity, and reducing output accordingly. Welfare costs arise in part from the direct impact of the events on consumption and in part from their unpredictability.[1] In Martin and Pindyck (2015), we considered different types of catastrophic events (e.g., nuclear terrorism, a climate catastrophe, or a financial crisis), explored their policy interdependence, and developed a rule for determining which events it is optimal to avert. But as with generic disasters, we assumed that the impact of each type of catastrophic event is a reduction in consumption.

Some catastrophic events kill people but cause little or no drop in the consumption of those who survive. Perhaps the best example — and the one we focus on in this paper — is a major pandemic that spreads uncontrollably. Most of the damage from such an event would be the deaths of a significant fraction of the population. Indeed, the Spanish Flu pandemic of 1918–1919 infected about 20 percent and 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 survivors. Thus an equally virulent pandemic today might kill more than ten million people in the U.S. alone. Epidemiologists have argued that new and more virulent pandemics are very likely to occur in the next couple of decades, and because of modern travel, could kill an even larger fraction of the population.[2]

We explore the welfare implications of the threat of pandemics, but in the context of a world also threatened by "generic" consumption disasters. As usual, we measure the benefit from averting either threat 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 eliminate the threat. We examine the WTP to avoid future pandemics, and ask how that WTP is affected by ongoing economic uncertainty. In particular, what is the WTP to avoid a pandemic, the WTP to avoid consumption disasters, and the WTP to avoid both?

---

[1]Barro (2006, 2009), Martin (2008), Pindyck and Wang (2013), and others refer to (generic) catastrophes as "consumption disasters." The potential for such catastrophes can help explain the equity premium and low risk-free rate "puzzles." As discussed below, Barro (2006, 2009), Barro and Ursúa (2008) and others use historical data for a panel of countries to estimate the mean arrival rate and impact distribution of these events; Martin (2008) and Pindyck and Wang (2013) use data on aggregate economic and financial variables to get the mean arrival rate and impact distribution as outputs of a general equilibrium model.

[2]See, e.g., Byrne (2008), Kilbourne (2008), Enserink (2004), and Harvard Global Health Institute (2018). Johnson and Mueller (2002) estimated the Spanish Flu U.S. mortality rate to be 6.5 percent.

Why include consumption disasters in a study of the threat of pandemics? First, as we showed in Martin and Pindyck (2015), with multiple potential catastrophes, simple cost-benefit analysis breaks down: because they are "non-marginal," considering them in isolation can lead to sub-optimal policies, even if the catastrophes occur independently of one another. Second, general economic uncertainty is a form of "background risk" that raises the WTP to avert a pandemic, and we want to determine by how much. Third, the welfare effects of consumption disasters provides a scale that can help us evaluate the importance of averting pandemics. And finally, combining the two in one model is natural because the framework we develop works by relating the welfare effects of fatalities to losses of consumption.[3]

Here we show how to incorporate death in a model of catastrophe avoidance, and how a catastrophic loss of life can be expressed as a welfare-equivalent drop in consumption. 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 deaths of some fraction of the population is much greater than the WTP to avert a drop in consumption by the same fraction; a drop in consumption causes a marginal loss of utility, whereas death causes a total loss (albeit for only a fraction of the population).

We proxy for the value of a life lost 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. 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. (So as society becomes richer, the VSL will increase.) Here we link the VSL literature to the literature on consumption disasters.

In Martin and Pindyck (2015) we assumed that there are N types of potential catastrophes, all of which would cause a drop in consumption. We now allow for just two types of catastrophes, one of which is a generic consumption disaster and the other is an event (e.g., a pandemic) that kills people. Thus the WTP to avoid the second type results from lives

---

[3]Barro (2014) develops a model with two catastrophes, a consumption disaster and also an environmental catastrophe (e.g., from climate change), with different probabilities of occurring. But he assumes the two catastrophes have the same impact distribution and simply adds the probabilities, in effect modeling a single catastrophe, but with a greater likelihood of occurring than the consumption disasters in the earlier papers.

saved, as opposed to consumption saved.[4]

Determining the WTP to avert a catastrophe that kills people introduces a fundamental issue: How should society value the very existence of people, alive now and potentially in the future, relative to the consumption enjoyed by those people? We want to be clear that we do not have anything new to say about the hard philosophical question regarding the value of human life. To our knowledge, there is no consensus regarding the social value of more or fewer people, although economists tend to value changes in population asymmetrically. *Increases* in population are seen by some as "good" (on purely ethical grounds, but also by contributing to technological change), and by some as "bad" (because of crowding, congestion, and environmental stress). Based on social norms, *decreases* in population are seen as purely "bad," in that we go to great lengths to save lives and prevent life-threatening disasters.[5] Consistent with these social norms, the model we develop treats deaths caused by catastrophes as a pure "bad" that reduces social welfare. Consistent with the views of environmentalists, our model also treats natural increases in population as "bad," although our focus is on the welfare effects of catastrophes.

Using middle-of-the-road estimates of the VSL and conservative estimates of the likelihood and impact distribution, we find the WTP to avert future pandemics to be large — more than 10% of consumption. Also, a good portion of that WTP comes from "background risk," here the threat of a consumption catastrophe. Likewise, the threat of pandemics substantially affects the WTP to avert consumption catastrophes. This is another example of the policy interdependence of major catastrophic threats.

In the next section we lay out our framework for measuring the welfare loss from catastrophes that kill people, but leave the consumption of survivors unchanged. In Sections 3 and 4 we present a fully dynamic model with two types of catastrophes that arrive independently as Poisson processes, one that causes a (random) drop in consumption and the other a (random) number of deaths. We specify distributions for the random impacts of each type, and we find the WTPs to avert each and to avert both. In Section 5 we calibrate the model using parameter values consistent with the recent literatures on pandemics and macroeconomic contractions, calculate the WTPs, and discuss the policy implications. Section 6 concludes.

---

[4]As 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 greater than the losses from reduced consumption.

[5]But see Young (2005) and Voigtländer and Voth (2013), who show how pandemics and plagues can raise welfare over time. For estimates of the negative population externality associated with climate change, see Bohn and Stuart (2015).

2   Framework

We assume that at time t there are N_t identical consumers who, barring a catastrophe, remain alive indefinitely. 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 $\eta > 1$. This is consistent with 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, that can occur repeatedly.

Without any catastrophe, real per-capita consumption, C_t, grows 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; $u_t = C_t^{1-\eta}/(1 - \eta)$. We assume that total welfare for society is proportional to the population:

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

Note that with $\eta > 1$, $V_0 < 0$, so a higher natural rate of population growth n reduces total social welfare. This is not the result of an explicit population externality, but simply an implication of scaling social welfare by population.[6] It might appear, then, that eqn. (1) implies that a catastrophe that reduces N_t will increase welfare. But as we will see, our treatment of deaths from a catastrophe ensures the opposite.

We consider two types of catastrophes. The first consists of catastrophes which cause a permanent drop in log consumption by a random fraction ϕ for the entire population (so that ϕ is roughly the fraction by which 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 rate of this type of event.

The second type consists of catastrophes which result in the death of a random fraction ψ of the population, where the distribution of deaths is also random, i.e., each individual has a probability ψ of dying. 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

---

[6]We could introduce a population externality by writing individual utility as $u_t = C_t^{1-\eta} N_t^\alpha/(1-\eta)$, so the externality is positive (negative) if $\alpha > (<) 0$. Then N_t in eqn. (1) is replaced by $N_t^{1+\alpha}$, so $\partial V_0/\partial N > (<) 0$ if $\alpha < (>) -1$. Population externalities are not the focus of this paper, so we will not introduce one.

ϕ and ψ (and these distributions need not be the same).

Later, when we calibrate and apply the model, we will be more specific about the nature of these two types of catastrophes (e.g., a "death" event will be a pandemic). For now we simply assume that one causes a drop in C_t and the other a drop in N_t. We will want to calculate the WTP to avert each type, and the WTP to avert both.[7]

2.1 The Quick and the Dead

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.

What is the welfare loss for a person that dies? One approach to this question is to simply assume that the consumption 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 $C \to 0$.

An alternative, "critical-level utilitarianism," is to find a consumption level C_0 just high enough to sustain (tolerable) life, and then measure utility relative to $u(C - C_0)$. (See Broome (2004) and Blackorby, Bossert and Donaldson (2005).) But as a matter of calibration, we see no obvious way to determine C_0. More importantly, in our model C could fall below C_0, and we need a way to compare this to an outcome in which C falls just to C_0.

Another and perhaps simplest approach is to assume that upon death, utility (as opposed to consumption) falls to zero, i.e., the individual loses the utility she enjoyed while alive. But with CRRA utility, this can considerably underestimate the welfare loss from death.

To see this, consider an individual who lives two periods and 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 not to zero, but only 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.

Nonetheless, the usual approach in the literature is to assume that upon death utility

---

[7]We 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.

falls to zero. For example, Hall and Jones (2007), in their study of the value of life-extending health expenditures, do this by adding a positive constant b to the utility function:

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

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.[8] In their case, death corresponds to a drop in consumption from C_0 to a value ω such that $u(\omega) = 0$, i.e., $\omega = [(\eta - 1)b]^{1/(1-\eta)}$.

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 can then treat death using the same framework used to treat the utility loss from a drop in consumption. To do this we adopt the VSL measure that is widely used in public policy.

2.2 The Value of Life

To determine the welfare effects of fatalities, 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 future lifetime consumption) and the probability of survival. Thus the VSL 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.[9]

---

[8]Becker, 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 range of C is also needed to ensure that indifference curves between consumption and life expectancy slope down. See Pratt and Zeckhauser (1996). In a related paper, Jones (2016) develops a representative agent model in which research effort can lead to consumption growth, but also the chance of a catastrophe that will kill the agent, whose utility then falls to zero. Arrow and Priebsch (2014) propose a generalized hyperbolic absolute risk aversion utility function which is bounded for any value of C.

[9]For an overview of the VSL and its measurement and use, see Andersson and Treich (2009), Ashenfelter (2006), and Viscusi (2018). The VSL has other issues. For example, the macro/finance literatures put the coefficient of relative risk aversion, η, well above 1. Estimates of the income elasticity of the VSL are 0.5 to

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 one. (See, e.g., Viscusi (1993) and Cropper and Sussman (1990).) The range of estimates is wide: 3 to 10 times lifetime income or lifetime consumption, with an average value of around 7. Consistent with this average value, the U.S. Environmental Protection Agency set the VSL at about $9 million when it conducted cost-benefit analyses of proposed policies.[10]

Let w represent wealth or lifetime consumption, and let p be the *ex ante* probability of death (so $1 - p$ is the probability of survival). Suppose an individual can reduce p at the cost of some reduction in w. Let $u(w)$ be the individual's utility if alive, and $v(w)$ 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 $1 - p$:[11]

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

Note that $u(w)$ and $v(w)$ are measured in utils, and $u'(w)$ and $v'(w)$ are measured in utils per dollar, so the VSL is measured in dollars. The VSL is a cardinal measure, invariant to affine 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. Consistent with the literature, we take the VSL to be an observable multiple of wealth: $\text{VSL} = sw$.

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. (Pratt and Zeckhauser (1996) call this the "dead anyway" effect.) 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

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 support me); see, e.g., Arthur (1981).

[10]U.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 to $9.7 million in 2017 dollars. In 2017, U.S. per capita consumption was about $41,000, which corresponds to lifetime consumption of roughly $1,600,000 (i.e., a factor of 40). The average VSL estimate of 7 times lifetime consumption yields $11.2 million. Rohlfs, Sullivan and Kniesner (2015), using airbag regulations during the 1990s and used car prices, find a median VSL between $9 and $11 million, and Aldy (2019), using employment decisions of married couples, estimates the VSL to be between $9 and $13 million. But Greenstone, Ryan and Yankovich (2012), using variation in occupation-specific mortality risks and occupation-specific re-enlistment bonuses, found that soldiers' re-enlistment decisions implies a VSL of only $3 to $4 million.

[11]To get eqn. (2), note that expected utility is $V = (1 - p)u(w) + pv(w)$. Setting the total derivative $dV = (\partial V/\partial p)dp + (\partial V/\partial w)dw = 0$ yields dw/dp.

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.

We will think of $u(w)$ as the utility associated with a wealth level of w in the hands of a living person, and $v(w)$ as the (social) utility associated with the same wealth if the person dies and their assets are redistributed, either via intentional bequests to children or others, or through some other reallocation mechanism. As s is observable and we can approximate $p \approx 0$, the relationship between $v(w)$ and $u(w)$ must satisfy

$$v(w) = u(w) - swu'(w). \tag{3}$$

Each person's utility while alive is ultimately derived from a flow of consumption, C_t:

$$\text{Expected utility} = \mathbb{E} \int_0^\infty \frac{1}{1-\eta} C_t^{1-\eta} e^{-\delta t} \, dt \,.$$

As discussed below, we assume that consumption growth is i.i.d., which implies that expected utility is proportional to $C_0^{1-\eta}/(1-\eta)$ and that consumption and wealth are proportional to one another (Martin (2008), Barro (2009)). Thus we can simply work with utility $u(w) = w^{1-\eta}/(1-\eta)$. As long as we are consistent in whether we work with wealth or consumption, the constant of proportionality itself is unimportant, given the invariance of the VSL to affine transformations.

Equation (3) implies that we must have

$$
\begin{aligned}
v(w) &= \frac{w^{1-\eta}}{1-\eta} - sw^{1-\eta} \\
&= \frac{w^{1-\eta}}{1-\eta} \left[1 + s\,(\eta - 1) \right] \\
&= \frac{\left\{ w \left[1 + s\,(\eta-1) \right]^{\frac{1}{1-\eta}} \right\}^{1-\eta}}{1-\eta}\,.
\end{aligned}
$$

It follows immediately that the social welfare loss associated with a death is equivalent to the loss that would occur following a drop in wealth for a living person from w to εw, where ε is given by:

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

Or, if thinking in consumption terms, it is equivalent to the loss following a drop in consumption from C to εC.

As explained above, estimates of s vary, but the average is around 7, and we will use that value in much of what follows. Thus if $\eta = 2$, $\varepsilon = 1/(s+1) = .125$, i.e, death is equivalent

in welfare terms to an 88% drop in wealth or 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 w or 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.

Note that in our model the social welfare loss from a death is equivalent to a *percentage* drop in wealth, which implies that this loss depends on the person's wealth level. In fact, the VSL itself is increasing in wealth w, because a wealthier individual has more utility to lose should she die.[12] This would be a problem if our objective were to make interpersonal (or cross-country) comparisons, or to estimate à la Hall and Jones (2007) how the relative marginal benefit of health care spending varies with income and consumption. We address a simpler problem by assuming in eqn. (1) that prior to a catastrophe consumption is the same for everyone, so total welfare is proportional to the population.

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 catastrophe-induced 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

$$U_N = U_0 - \psi \frac{1 - \varepsilon^{1-\eta}}{1 - \eta} < U_0$$

The adjustment term $\psi(1-\varepsilon^{1-\eta})/(1-\eta)$ captures the welfare loss associated with the fraction of the population that has died, and ensures that $U_N < U_0$. From (4), $\varepsilon^{1-\eta} = s(\eta - 1) + 1$, so $U_N = [1 + (\eta - 1)\psi s]/(1 - \eta) < U_0$ 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. 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 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.)

---

[12]Think of $D(w) = v(w) - u(w)$ as the "cost of death," so expected utility is $V = u(w) - pD(w)$ and we can rewrite eqn. (2) as $dw/dp = D(w)/[u'(w) + pD'(w)]$. $D(w)$ is increasing in w, as is the VSL. (Thus Bill Gates has more to lose by dying than most other people.)

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

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

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 $\phi \to 0$, LR $\to s$, 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.

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 $u[(1-\phi_c)C_0] - u(C_0) = \phi[u(\varepsilon C_0) - u(C_0)]$,

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

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.) 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   A Dynamic Model

To determine the welfare loss from a drop in per-capital consumption C_t or a drop in population N_t, we need to describe how these events can unfold over time. We thus turn to a fully dynamic model, and specify the processes for C_t and N_t. We assume that catastrophic events that reduce C_t occur as Poisson arrivals with mean arrival rate λ_c, and the impact of the kth arrival, ϕ_k, is i.i.d. across realizations k. Then the process for consumption is:

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

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 kth 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),

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

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

$$\boldsymbol{\kappa}_C(\theta) \equiv \boldsymbol{\kappa}_{C,1}(\theta) = g\theta + \lambda_c \left(\mathbb{E}\, e^{-\theta\phi} - 1 \right), \tag{8}$$

where ϕ is a representative of the (i.i.d.) ϕ_k.[13]

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

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

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 kth event. Thus when the kth 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,

$$\boldsymbol{\kappa}_{N,t}(\theta) \equiv \log \mathbb{E}\, e^{n_t\theta} \equiv \log \mathbb{E}\, N_t^\theta.$$

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

$$\boldsymbol{\kappa}_N(\theta) \equiv \boldsymbol{\kappa}_{N,1}(\theta) = n\theta + \lambda_d \left(\mathbb{E}\, e^{-\theta\psi} - 1 \right) \tag{10}$$

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

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

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 a death event occurs. 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 $\boldsymbol{\kappa}_N(1) < 0$.

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 $\boldsymbol{\kappa}_N^*(\theta) = n\theta$. Likewise, if there are no consumption catastrophes ($\lambda_c = 0$), consumption evolves as $C_t^* = e^{gt}$, and the CGF for C_t is $\boldsymbol{\kappa}_C^*(\theta) = g\theta$.

---

[13]We 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.

Recall that the social welfare loss associated with a death is equivalent to the loss that would occur from a drop in consumption to a fraction ε of what it was before, i.e., from C_t to to εC_t, where ε is given by eqn. (4). So if no catastrophes are averted, total welfare is

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

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,

$$\mathbb{E}(N_t C_t^{1-\eta}) = \mathbb{E}\,N_t\,\mathbb{E}\,C_t^{1-\eta} = e^{\boldsymbol{\kappa}_N(1)t} \cdot e^{\boldsymbol{\kappa}_C(1-\eta)t}$$

and

$$\mathbb{E}\left[(N_t^* - N_t)\varepsilon^{1-\eta} C_t^{1-\eta}\right] = \left(e^{\boldsymbol{\kappa}_N^*(1)t} - e^{\boldsymbol{\kappa}_N(1)t}\right)\varepsilon^{1-\eta} e^{\boldsymbol{\kappa}_C(1-\eta)t}.$$

Substituting these expressions into the integral above,

$$\begin{aligned}
V &= \frac{1}{1-\eta}\left\{\frac{1}{\delta - \boldsymbol{\kappa}_N(1) - \boldsymbol{\kappa}_C(1-\eta)} + \frac{\varepsilon^{1-\eta}}{\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C(1-\eta)} - \frac{\varepsilon^{1-\eta}}{\delta - \boldsymbol{\kappa}_N(1) - \boldsymbol{\kappa}_C(1-\eta)}\right\} \\
&= \frac{1}{1-\eta}\left\{\frac{1 - \varepsilon^{1-\eta}}{\delta - \boldsymbol{\kappa}_N(1) - \boldsymbol{\kappa}_C(1-\eta)} + \frac{\varepsilon^{1-\eta}}{\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C(1-\eta)}\right\}. \qquad (12)
\end{aligned}$$

To interpret this equation, note that the second term captures the welfare associated with the guaranteed consumption stream εC_t (which is "received" via bequests, etc., even after death). We can think of $\delta - \boldsymbol{\kappa}_N^*(1) = \delta - n$ as the social rate of time preference (i.e., discount rate on future utility) associated with this consumption stream: it adjusts, via $\boldsymbol{\kappa}_N^*(1) = n$, for the fact that the population is increasing (so that the larger is n, the lower is the discount rate, and because $V < 0$, the lower is total welfare). The first term captures the welfare associated with the additional ongoing consumption stream $(1-\varepsilon)C_t$ received by those who are still alive (i.e., those alive receive $(1-\varepsilon)C + \varepsilon C = C$). Since there is a risk of death, this latter consumption stream is discounted at the higher rate $\delta - \boldsymbol{\kappa}_N(1) > \delta - \boldsymbol{\kappa}_N^*(1)$.[14]

We want to find the WTPs to avert each type of catastrophes individually, and the WTP to avert both. Averting a consumption catastrophe corresponds to setting $\lambda_c = 0$, and thus replacing $\boldsymbol{\kappa}_C(1-\eta)$ by $\boldsymbol{\kappa}_C^*(1-\eta) \equiv g(1-\eta)$. What is the WTP to avert this catastrophe? If it is averted at the cost of a permanent loss of a fraction w_c of consumption, welfare is

$$V_c = \frac{(1-w_c)^{1-\eta}}{1-\eta}\left\{\frac{1 - \varepsilon^{1-\eta}}{\delta - \boldsymbol{\kappa}_N(1) - \boldsymbol{\kappa}_C^*(1-\eta)} + \frac{\varepsilon^{1-\eta}}{\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C^*(1-\eta)}\right\}. \qquad (13)$$

---

[14]The expressions $1/[\delta - \boldsymbol{\kappa}_N(1) - \boldsymbol{\kappa}_C(1-\eta)]$ and $1/[\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C(1-\eta)]$ can be interpreted as the valuation ratios of the two consumption streams, as shown by Martin (2013). Because $V < 0$, a higher value of δ implies a higher value of V. However, we treat δ as fixed and ignore how it might be determined.

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

$$A = \frac{\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C^*(1-\eta)}{\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C(1-\eta)} \tag{15}$$

$$B = \frac{\delta - \boldsymbol{\kappa}_N(1) - \boldsymbol{\kappa}_C^*(1-\eta)}{\delta - \boldsymbol{\kappa}_N(1)\varepsilon^{1-\eta} - (1-\varepsilon^{1-\eta})\boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C^*(1-\eta)} \tag{16}$$

$$C = \frac{\delta - \boldsymbol{\kappa}_N(1)\varepsilon^{1-\eta} - (1-\varepsilon^{1-\eta})\boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C(1-\eta)}{\delta - \boldsymbol{\kappa}_N(1) - \boldsymbol{\kappa}_C(1-\eta)}. \tag{17}$$

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

$$V_d = \mathbb{E}\left\{\int_0^\infty e^{-\delta t} \frac{(1-w_d)^{1-\eta} N_t^* C_t^{1-\eta}}{1-\eta} \, dt\right\} = \frac{(1-w_d)^{1-\eta}}{1-\eta}\left\{\frac{1}{\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C(1-\eta)}\right\}. \tag{18}$$

Equating (12) and (18), w_d satisfies

$$(1-w_d)^{1-\eta} = C. \tag{19}$$

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

$$V_{c,d} = \frac{(1-w_{c,d})^{1-\eta}}{1-\eta}\left\{\frac{1}{\delta - \boldsymbol{\kappa}_N^*(1) - \boldsymbol{\kappa}_C^*(1-\eta)}\right\}. \tag{20}$$

Equating (12) and (20), $w_{c,d}$ satisfies

$$(1-w_{c,d})^{1-\eta} = A \times C. \tag{21}$$

Now we can compare expressions (14), (19), and (21) for the WTPs w_c, w_d, and $w_{c,d}$. We assumed that $\eta > 1$, so (i) $A > 1$, (ii) $B < 1 < C$, and (iii) $BC > 1$, using the facts that $\boldsymbol{\kappa}_C^*(1-\eta) < \boldsymbol{\kappa}_C(1-\eta)$, $\boldsymbol{\kappa}_N^*(1) > \boldsymbol{\kappa}_N(1)$, and $\varepsilon^{1-\eta} > 1$ (for $\eta > 1$).[15] 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.[16]

---

[15] This implies, for example, that $\boldsymbol{\kappa}_N(1)\varepsilon^{1-\eta} + \boldsymbol{\kappa}_N^*(1)(1-\varepsilon^{1-\eta}) < \boldsymbol{\kappa}_N(1)$.

[16] The 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.

4 Impact Distributions and WTPs

Our results so far are quite general, in that the CGFs of eqns. (8) and (10) apply to *any* probability distributions for the impacts ϕ and ψ. In order to apply the model, we will assume that ϕ and ψ are exponentially distributed. This distribution has often been used to model catastrophic events, and considerably simplifies the expressions for the WTPs:

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

This is equivalent to assuming that $z_c = e^{-\phi}$ and $z_d = e^{-\psi}$, the surviving (post-impact) fractions of consumption and people, 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 $0 \leq z_d \leq 1$. 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:

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

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

$$\lambda_c' \equiv \lambda_c(\eta - 1)/(\beta_c + 1 - \eta) \tag{22}$$

$$\lambda_d' \equiv \lambda_d/(\beta_d + 1) \tag{23}$$

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 its expected arrival rate.

14

Substituting the CGFs, λ'_c, λ'_d, and $\rho \equiv \delta - n + g(\eta - 1)$ into (15), (16), and (17), the factors A, B, and C become:

$$A = \frac{\rho}{\rho - \lambda'_c}$$

$$B = \frac{\rho + \lambda'_d}{\rho + \lambda'_d \varepsilon^{1-\eta}}$$

$$C = \frac{\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c}{\rho + \lambda'_d - \lambda'_c}$$

We can use these expressions for A, B, and C to get the WTPs. Recall that $(1 - w_c)^{1-\eta} = ABC$, $(1 - w_d)^{1-\eta} = C$, and $(1 - w_{c,d})^{1-\eta} = AC$. Thus:

$$w_c = 1 - \left[\frac{(\rho - \lambda'_c)(\rho + \lambda'_d \varepsilon^{1-\eta})(\rho + \lambda'_d - \lambda'_c)}{\rho(\rho + \lambda'_d)(\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c)} \right]^{\frac{1}{\eta - 1}} \tag{24}$$

$$w_d = 1 - \left[\frac{(\rho + \lambda'_d - \lambda'_c)}{(\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c)} \right]^{\frac{1}{\eta - 1}} \tag{25}$$

$$w_{c,d} = 1 - \left[\frac{(\rho - \lambda'_c)(\rho + \lambda'_d - \lambda'_c)}{\rho(\rho + \lambda'_d \varepsilon^{1-\eta} - \lambda'_c)} \right]^{\frac{1}{\eta - 1}} \tag{26}$$

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. Thus 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.

Also note that $\partial w_d / \partial \lambda'_c > 0$ and $\partial w_c / \partial \lambda'_d > 0$, i.e., an increase in the likelihood of catastrophe i increases the WTP to avert catastrophe j. Each potential catastrophe creates "background risk" that raises the WTP to avert the other catastrophe. As we will see in the next section when we examine pandemics and consumption contractions, these effects of "background risk" can be substantial.

To determine which catastrophes (if any) should be averted, we would also need to know the cost of averting each one, expressed as a permanent tax on consumption at rates denoted by τ_c and τ_d. We would then calculate the net (of taxes) welfare of doing nothing (W_0), averting only consumption catastrophes (W_c), averting only death catastrophes (W_d), and averting both ($W_{c,d}$), and choose the one that is highest.[17] Using eqns. (12), (13), (18),

---

[17]Because 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.

(20) and the expressions above for the CGFs, the net welfare for each policy is:

$$W_0 = \frac{1}{1-\eta}\left[\frac{-s(\eta-1)}{\rho+\lambda_d'-\lambda_c'} + \frac{s(\eta-1)+1}{\rho-\lambda_c'}\right] \tag{27}$$

$$W_c = \frac{(1-\tau_c)^{1-\eta}}{1-\eta}\left[\frac{\rho+\lambda_d'(s(\eta-1)+1)}{\rho(\rho+\lambda_d')}\right] \tag{28}$$

$$W_d = \frac{(1-\tau_d)^{1-\eta}}{(1-\eta)(\rho-\lambda_c')} \tag{29}$$

$$W_{c,d} = \frac{(1-\tau_c)^{1-\eta}(1-\tau_d)^{1-\eta}}{(1-\eta)\rho} \tag{30}$$

Note from (28) and (29) that an increase in λ_d (and thus λ_d') reduces W_c, net welfare when the consumption catastrophe is eliminated, and an increase in λ_c (and thus λ_c') reduces W_d, net welfare when the death catastrophe is eliminated. This is the loss of welfare created by "background risk," here the risk created by the second catastrophe. (And of course an increase in λ_d or λ_c also reduces W_0, net welfare when neither catastrophe is eliminated.)

5   Lost Consumption vs. Lost Lives

We turn now to the threat of major pandemics and its interdependence with the threat of macroeconomic contractions. What is the WTP to avoid each, and the connection between the two? To address these questions we need values for the mean arrival times λ_d and λ_c, the impact parameters β_d and β_c, and the parameters s, g, n, δ and η.

5.1   Pandemics

According to an assessment by the Centers for Disease Control and Prevention (CDC), the threat of a "mega-virus" that could cause widespread fatalities is considerable: "While we can't predict exactly when or where the next epidemic or pandemic will begin, we know one is coming." The main reasons: "Increased risk of infectious pathogens spilling over from animals to humans; development of antimicrobial resistance; spread of infectious diseases through global travel and trade; acts of bioterrorism; and weak public health infrastructures."[18]

---

[18]https://www.cdc.gov/globalhealth/healthprotection/fieldupdates/winter-2017. The CDC is especially concerned with the possibility of a pandemic resulting from mutations of the Asian H7N9 avian influenza virus: "... the pandemic potential of this virus is concerning. Influenza viruses constantly change and it is possible that this virus could gain the ability to spread easily and sustainably among people, triggering a global outbreak of disease." See https://www.cdc.gov/flu/avianflu/h7n9-virus.htm

Our experience with major pandemics is (fortunately) limited, making the calibration of λ_d and β_d somewhat speculative. During the past century the Spanish Flu was by far the worst pandemic, but there were several others, though not as severe. For example, the 1957-1958 H2N2 pandemic killed 1.1 million worldwide and 116,000 in the U.S. (for a U.S. fatality rate of .0064); the 1968 H3N2 pandemic killed about 1 million worldwide and 100,000 in the U.S.(for a U.S. fatality rate of .0006); and the 2009 H1N1 pandemic killed 12,469 in the U.S. And then there is HIV/AIDS, which killed more than 700,000 in the U.S. since the epidemic began in 1985, and worldwide close to a million people in 2017 alone.

Based on these numbers and assuming that the likelihood and potential virulence has remained about the same, we would set $\lambda_d = .02$ (i.e., an event on average every 50 years and a 33-percent chance of at least one more in the next 20 years) and $\beta_d = 24$ (consistent with the roughly 4-percent U.S. mortality rate of the Spanish Flu). These numbers would imply that the probability of a pandemic over the next 20 years with at least a 4-percent (8-percent) mortality rate is .14 (.05).[19] However, epidemiologists have argued that major pandemics have become much more likely and could be much more virulent. For example, Jones et al. (2008) analyze 335 "emerging infectious diseases" from 1940 through 2004, and show sharp increases in the numbers of events over time, after controlling for reporting bias.[20] Although these studies do not provide specific estimates of the current likelihood of a major pandemic or possible mortality rates, a doubling of λ_d to .04, while holding β_d fixed at 24, is consistent with the literature. This would imply a 55-percent chance of another pandemic (of some size) in the next 20 years, and implies that the probability of a pandemic occurring in the next 20 years with at least a 4-percent (8-percent) mortality rate is .26 (.10).

Given the uncertainties involved, we consider both of these parameterizations: a "low-risk" pandemic scenario roughly matching the past century, i.e., with $\lambda_d = .02$ and $\beta_d = 24$, and a "high-risk" scenario with $\lambda_d = .04$ and $\beta_d = 24$.

---

[19]For Poisson arrivals with mean a arrival rate λ and exponentially distributed impact with parameter β, the probability of at least one event with loss $\geq L$ over the horizon T is

$$Pr(T, L) = 1 - \exp[-\lambda T((1 - L)^\beta],$$

and the expected total loss over the horizon T is $\mathbb{E}(L; T) = 1 - \exp[-\lambda T/(\beta + 1)]$.

[20]Also, in part because of the overuse of antibiotics, strains of bacteria have evolved that are resistant to most or all antibiotics. See Byrne (2008), Kilbourne (2008) and Enserink (2004), and for a detailed discussion of the increased likelihood and severity of pandemics, Harvard Global Health Institute (2018). Fan, Jamison and Summers (2018) point out that most estimates of the costs of pandemics ignore the cost of lost lives. They estimate that cost using age-specific mortality rates following the 1918 Spanish Flu. However, they consider a known pandemic as opposed to the random arrival of pandemics, with an increasing arrival rate.

5.2 Consumption Catastrophes

A major recession can be viewed as a pure consumption catastrophe. We want to estimate the frequency of such events (i.e., the Poisson arrival rate λ_c) and their average intensity (i.e., the power distribution parameter β_c). This can be done in two different ways.

The most common approach, used by Barro (2006, 2009), Barro and Ursúa (2008), Nakamura et al. (2013) and others, is to rely on historical data, e.g., a panel of countries over the past century. For example, Barro and Ursúa (2008), whose results we will use, estimate arrival rates and intensity from a panel of 36 countries (for GDP) from 1870 to 2006. One problem with this approach is that these catastrophes are, almost by definition, "rare events," limiting what we can learn from even a century or more of data. Also, many of the catastrophes that these and other authors have studied are manifestations of just three global events—World Wars I and II and the Great Depression. Nonetheless, given that this approach is so widely used, we will draw upon these studies to calibrate λ_c and β_c.

A second approach is to ask what arrival rate and impact distribution are implied by the behavior of economic and financial variables. Martin (2008) shows, for example, that the economy's consumption-wealth ratio along with the preference parameters δ and η are sufficient to estimate the welfare effects of eliminating or reducing consumption uncertainty, without knowing the stochastic process for consumption. Pindyck and Wang (2013) develop an equilibrium model of the economy with shocks to the capital stock, which, as here, are Poisson events with known arrival rate and exponential impact distribution. Using data on consumption, investment, output, the risk-free rate, and the first four moments of stock returns, they determine λ_c and β_c, as well as η and δ, as calibration outputs. They find $\eta \approx 3.1$, $\delta \approx .05$, $\lambda = .73$ and $\beta = 23.2$, implying that catastrophes are frequent (occurring about every 1.4 years on average) but with a mean loss of only about 4% (so most shocks are small and can be viewed as part of the normal fluctuations in the economy).[21] Backus, Chernov and Martin (2011) use the prices of equity index options to infer the characteristics of consumption contractions; they also find that shocks are more frequent but smaller on average than suggested by estimates obtained from historical consumption data.

To calibrate our model, we will draw upon Barro and Ursúa (2008). Using world GDP data, but only including declines of 10-percent or greater, they estimate an annual arrival rate of $\lambda_c' = .036$ and average decline of 20.8 percent. Thus for *all* declines, the parameters

---

[21]But as explained in Pindyck and Wang (2013), pp. 320–321, these results are similar to those in Barro and Jin (2011), and those in Barro and Ursúa (2008) using U.S. GDP data. Barro and Jin (2011) and Barro and Ursúa (2008) only include contractions greater than .095 and .10 respectively. Transforming their estimates to account for *all* contractions yields relatively large numbers for λ_c and β_c, as shown below.
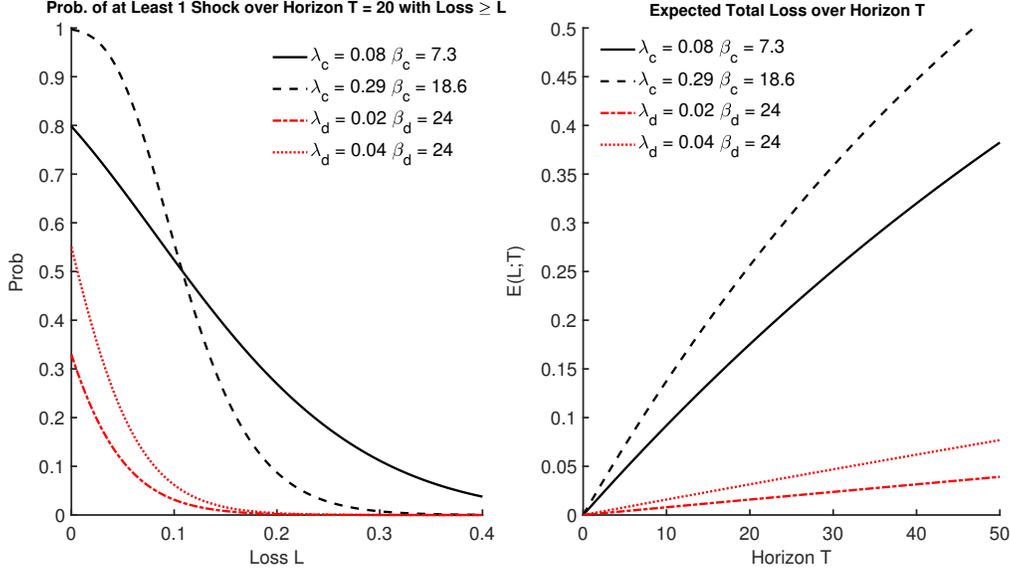
Figure 1: Probability of shock of size L or greater during next 20 years, and expected cumulative loss, for each of four catastrophes.

are $\lambda_c = .079$ and $\beta_c = 7.3$.[22] These numbers imply a contraction every 13 years on average, with an average contraction size of 12 percent. They also imply that the probability of a contraction occurring in the next 20 years of at least 10 percent (25 percent) in magnitude is .52 (.18). We refer to this calibration as BU1.

As an alternative, we also use the Barro and Ursúa (2008) data for GDP contractions of 10-percent or more in the U.S. They document 5 such contractions over 137 years with an average contraction size of 15.5 percent.[23] Adding the 2008 recession and its estimated contraction size of 10 percent[24] yields 6 contractions over 149 years, so $\lambda'_c = .040$, and an average contraction size of 14.6 percent. Adjusting to allow for contractions of *any* size yields $\lambda_c = .29$ and $\beta_c = 18.6$, which implies much more frequent (every 3.4 years on average) but smaller contractions (5 percent on average), and are closer to the results in Pindyck and Wang (2013) and Backus, Chernov and Martin (2011). This calibration (which we refer to as BU2) implies that the probability of a contraction in the next 20 years of at least 10 percent (25 percent) is .56 (.03).

The left-hand panel of Figure 1 plots the probability of a contraction of size L or greater

---

[22]$\mathbb{E}(L|L \geq .10) = .208$, so $\beta_c = 7.3$ is the solution to $1 - .9\beta_c/(\beta_c + 1) = .208$. Likewise, $\lambda_c = .079$ is the solution to $1 - \exp(-\lambda_c(.9)^{7.3}) = \lambda'_c = .036 =$ annual probability of 1 event with loss $L \geq .10$.

[23]From page 270 and Table C2, the contractions were "10 percent in 1908 and 1914 ..., 12 percent in 1921, 29 percent in 1933, and 16 percent in 1947."

[24]See, e.g., Fernald et al. (2017) and Barnichon, Matthes and Ziegenbein (2018).

occurring during the next 20 years for both calibrations, and illustrates how they differ. The average contraction size in BU2 is much smaller than in BU1, but the greater frequency of contractions results in a greater expected cumulative loss over time, as shown by the right-hand panel of Figure 1.[25] Thus we treat BU2 as the more "severe" parameterization.

Figure 1 also includes the pandemic calibrations. It shows the probability of a loss L or greater during the next 20 years and the expected total loss for both the "low-risk" and "high-risk" calibrations, i.e., $(\lambda_d, \beta_d) = (.02, 24)$ and $(\lambda_d, \beta_d) = (.04, 24)$. Here the loss is the fraction of the population that dies, so although the expected total loss is much lower than it is for consumption contractions, that does not make the corresponding WTP lower.

5.3 WTPs

The WTPs also depend on η, δ, g, n and the VSL parameter s. We use values for η and δ broadly consistent with both the macro/finance and "disaster" literatures. Pindyck and Wang (2013) obtain $\eta \approx 3.1$ and $\delta \approx .05$ as a calibration output of their general equilibrium model. Using estimates of λ and β for historical GDP contractions, Barro and Ursúa (2008) and Barro (2009) find that $\eta \approx 3.5$ or 4 and $\delta \approx .05$ are consistent with an unlevered equity premium of .05. We will set $\delta = .05$ and $\eta = 3$, and examine how the WTPs vary with η. Historical data for the U.S. puts $g = .02$ and $n = .01$. We set $s = 7$, which is the in the mid-range of most estimates, so $\varepsilon^{1-\eta} = s(\eta - 1) + 1 = 7\eta - 6$. However, estimates of the VSL (and thus s) vary widely, so we also examine how the WTPs vary with s.

We calculate the WTP to avert each catastrophe and the WTP to avert both, i.e., w_c, w_d, and $w_{c,d}$. To gauge the interdependence of the two catastrophes, we also calculate w_c' and w_d', the WTP to avert one but ignoring the presence of the other. We calculate these WTPs for two sets of calibrations: (1) $(\lambda_c, \beta_c) = (.08, 7.3)$ and $(\lambda_d, \beta_d) = (.02, 24)$; and (2) $(\lambda_c, \beta_c) = (.29, 18.6)$ and $(\lambda_d, \beta_d) = (.04, 24)$. Set (1) can be viewed as relatively "low risk," in that the total expected losses for both catastrophes are smaller than for set (2).

Figure 2 shows the WTPs as functions of η (for $s = 7$) and s (for $\eta = 3$) for the "low-risk" set of calibrations, and Figure 3 for the "high-risk" set. Three things stand out. First, the WTPs are large, even for the "low-risk" calibration. Note from Figure 2 that the WTP to avert a pandemic like the Spanish Flu — with an annual likelihood of 2 percent — is around 10% of consumption. The WTP to avert consumption contractions is around 23% of consumption for $\eta = 3$, and higher for $\eta > 3$. This may seem high, but is within the range

---

[25]For any horizon T, the expected total loss, $\mathbb{E}(L; T) = 1 - \exp[-\lambda T/(\beta + 1)]$, is increasing in $\lambda/(\beta + 1)$, which is .0095 for BC1 and .0145 for BC2. We are assuming that contractions are permanent, although Nakamura et al. (2013) have shown that on average about half of the lost GDP is regained during recoveries.
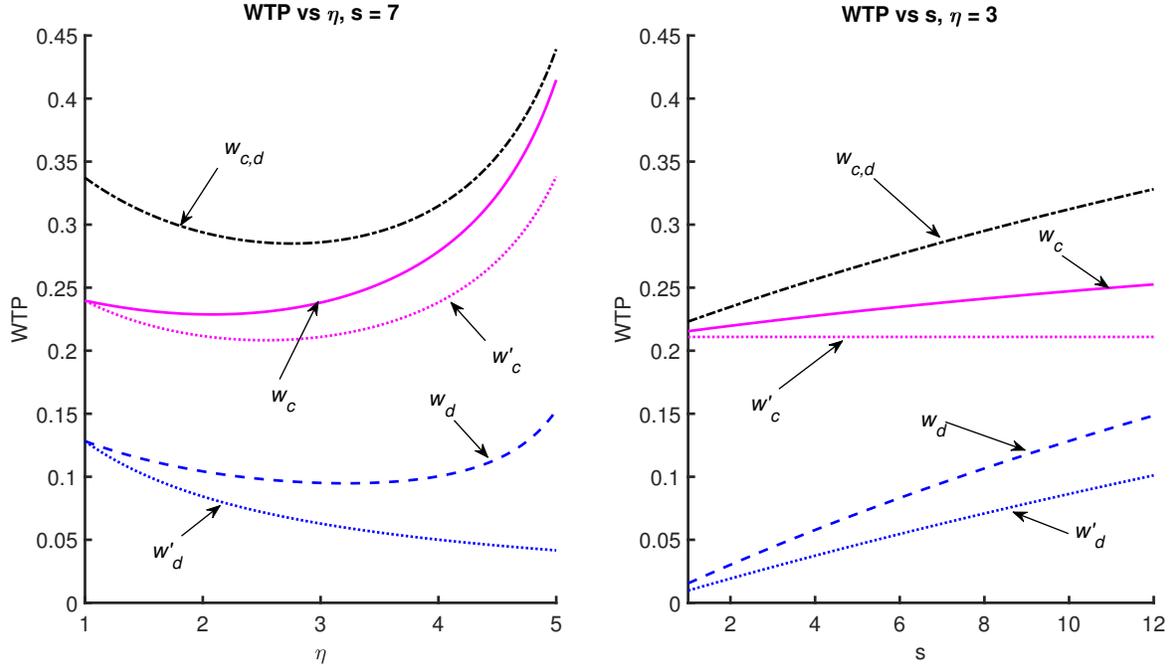
Figure 2: WTP vs η and s, $(\lambda_c, \beta_c) = (.08, 7.3)$ and $(\lambda_d, \beta_d) = (.02, 24)$.

of estimates by others of the welfare costs of consumption uncertainty. For example, Barro (2009) estimates the welfare cost at about 20% of GDP, Martin (2008) gets an estimate of about 14% (for $\eta = 4$ and $\delta = .03$, but about 25% for $\eta = 3$ and $\delta = .05$), and Pindyck and Wang (2013) get a WTP of 30% (15%) of consumption to eliminate all contractions greater than 5% (10%). For the "high-risk" parameterizations in Figure 3, the WTPs are much higher: 18% (31%) of consumption to eliminate pandemic (consumption) risk.

Second, the interaction between the two catastrophes is strong, although it depends on η. For $\eta \geq 1.5$, a good part of each WTP results from the "background risk" coming from the threat of the other catastrophe. For example, if $\eta = 3$ and $s = 7$, for the "low-risk" ("high-risk") parameterizations, consumption risk accounts for about 40 percent (35 percent) of the 10% (18%) WTP to avert pandemics, and pandemic risk accounts for about 13 percent (17 percent) of the 23% (31%) WTP to eliminate consumption risk. This interdependence is even greater for larger values of η, as illustrated in Figure 4, which shows the ratios w'_c/w_c and w'_d/w_d as functions of η for the two calibrations. The dependence on η comes about because the threat of Catastrophe i reduces expected future consumption, which increases the future marginal utility of consumption by an amount increasing in η, and this raises the value of averting Catastrophe j, which if it occurs would further reduce consumption.

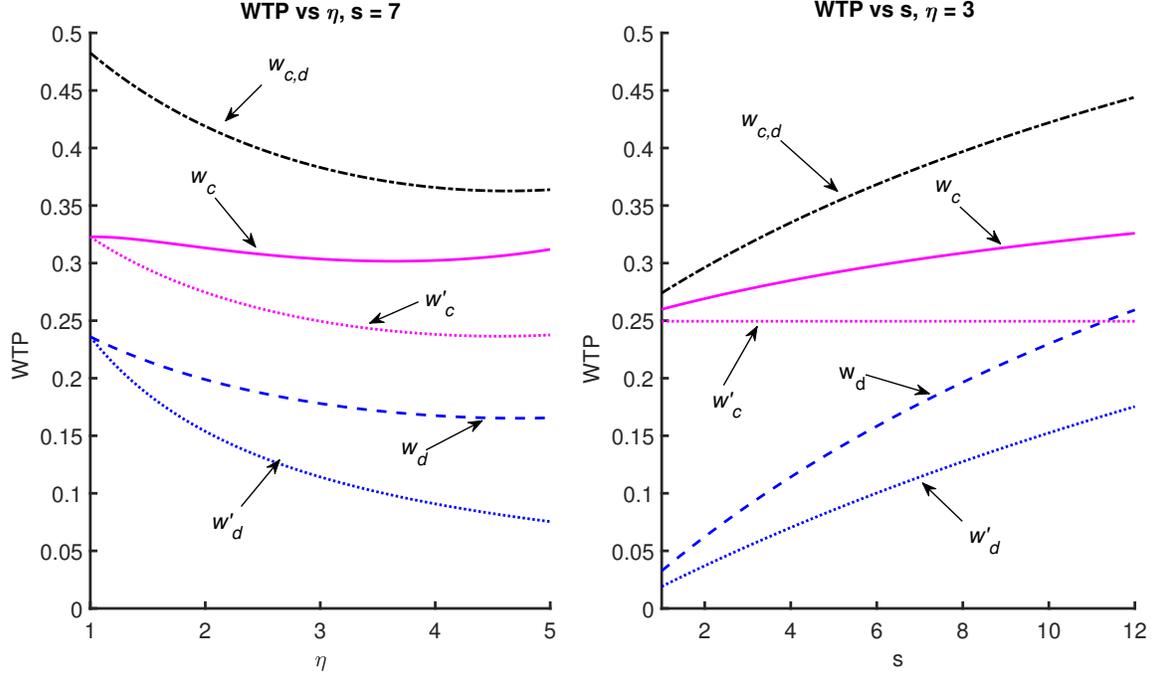Third, the dependence of the WTPs on η is complicated. In Figure 2 both w_d and w_c

Figure 3: WTP vs η and s, $(\lambda_c, \beta_c) = (.29, 18.6)$ and $(\lambda_d, \beta_d) = (.04, 24)$.

decrease and then increase with η, and in Figure 3, w_d is always decreasing with η and w_c is roughly constant. Increasing η has two off-setting effects. First, it increases the risk- (and impact-) adjusted arrival rate for the consumption catastrophe, λ_c' (but has no effect on the impact-adjusted arrival rate for the death catastrophe, λ_d'), which increases both w_d and w_c. 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.

5.4 Implications and Caveats

The high values of w_c and w_c' support earlier studies that found the welfare costs of (un-predictable) macroeconomic contractions to be large, but they add an important element: These costs are partly due to the threat of *other* catastrophes, in this case pandemics. Pandemic risk, a problem in its own right, adds to the welfare cost of consumption risk. And pandemics are just one source of "background risk;" here we are ignoring the other potential catastrophes (nuclear and bioterrorism, climate, earthquakes, etc.) examined in Martin and Pindyck (2015). Likewise, macroeconomic risk is an important component of the welfare costs of pandemics. The WTP to avoid future pandemics would be about a third lower if we could somehow eliminate macroeconomic fluctuations.

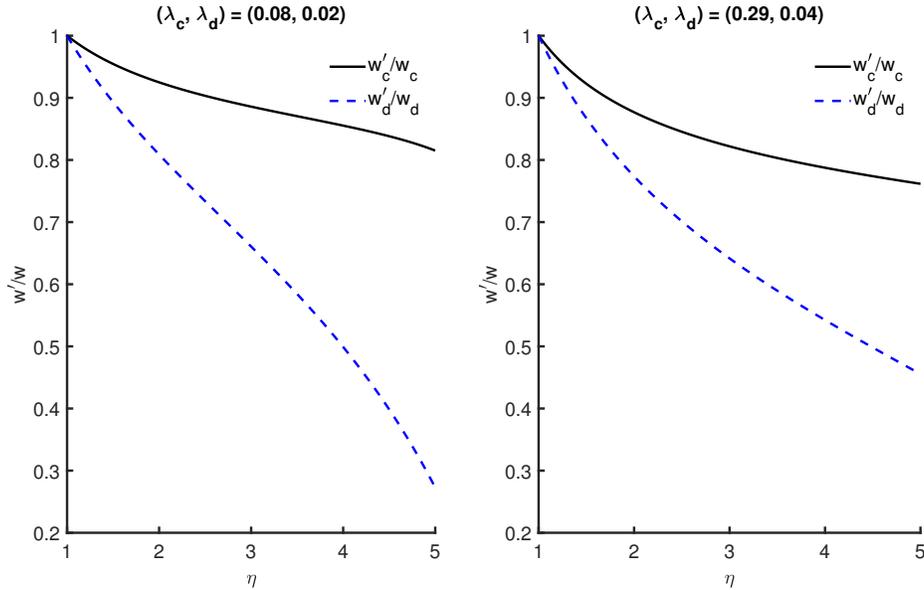But even without "background risk" and even for the conservative calibration of Fig-

Figure 4: Interdependence of WTPs: w'_i is the WTP to avert i but ignoring the threat of j.

ure 2, the WTP to avoid pandemics is substantial — depending on η, at least 5 to 10% of consumption. For the U.S. in 2018, this translates to around $1 trillion per year, and would be much higher if we account for consumption risk and use the calibration of Figure 3 (which assumes that pandemics have become more likely and more virulent).

As the right-hand panels of Figures 2 and 3 make clear, these high values for w_d and w'_d are dependent on the VSL parameter s. Reducing s from its base value of 7 to, say, 3 cuts w_d roughly in half in both calibrations. So is $s = 7$ the "correct" value to use for this parameter? It is in the mid-range of estimates and consistent with the VSL values of $9 to $10 million used by U.S. regulatory agencies, but it values the lives of the entire U.S. population at about $3,000 trillion, which is some 4 times greater than the roughly $800 trillion present value of U.S. GDP over the next 40 years. When large numbers of lives are involved — as is the case with pandemics — aggregating individuals' VSL values (based on small reductions in the risk of death) will overstate society's WTP to save lives.

However, even if we reduce s to 2 or 3, the WTP to avert pandemics is high — some 4% to 9% of consumption ($0.5 to $1.3 trillion for the U.S. in 2018), depending on the calibration and on η. Is this completely unrealistic? Remember that the WTP is a reservation price, i.e., the *maximum* fraction of consumption society would be willing to sacrifice to achieve a policy objective. The cost of achieving this objective might be much lower, and indeed for pandemics that appears to be the case. Several studies have concluded that pandemic risk could be greatly reduced (but not eliminated entirely) at a cost well below our WTP numbers. But these studies also make clear that we should be spending far more than we

23

currently are on pandemics, a policy position strongly supported by our results.[26]

6   Conclusions

In our earlier work we showed that decisions to avert (fully or partially) major catastrophes are interdependent, even if the catastrophes themselves are independent, and we showed how to determine which ones to avert. But as with most (if not all) of the literature, we assumed that catastrophes only cause a reduction in the flow of consumption. Here we have shown how catastrophes that cause death can be incorporated into our framework. In so doing, we link the VSL and consumption disaster literatures.

Because we work with CRRA utility and a constant rate of time preference, consumption and wealth are proportional to each other, and we have shown that the social welfare loss from a death is equivalent to the loss that would occur from a drop in wealth (or consumption) from w to εw (or C to εC) for a living person, where ε is tied to the VSL via eqn. (4). We use estimates of the VSL to determine the value of ε and thus the utility loss from death. We can then find the WTPs to avert a "death" catastrophe (w_d), a consumption catastrophe (w_c), and both ($w_{c,d}$). Our general results hold for any probability distributions for the random reductions ϕ and ψ of $\log C$ and $\log N$, but if ϕ and ψ are exponentially distributed, these WTPs are very easy to calculate.

We used our model to explore the welfare implications of the threat of pandemics, a threat that epidemiologists argue is substantial and growing. We did this in the context of macroeconomic uncertainty, and in particular the ongoing threat of sharp consumption contractions. There is a large literature on consumption disasters, which we used to calibrate the parameters for the stochastic process for consumption. The literature on pandemics is more speculative, but sufficient to reasonably calibrate the corresponding parameters. We find the utility loss from a "death" catastrophe is large, as is the WTP to avert it. Using middle-of-the-road estimates of the VSL and conservative estimates of the likelihood and impact distribution, we find the WTP to avert future pandemics to be well over 10% of consumption. We also found that a good portion of the WTP to avert pandemics comes from "background risk," here the threat of a consumption catastrophe. Likewise, the threat of pandemics substantially affects the WTP to avert consumption catastrophes. This is another example of the policy interdependence of major catastrophic threats.

One could argue that middle-of-the-road estimates of the VSL are unrealistically large

---

[26]See, for example, National Academy of Medicine (2016) and Harvard Global Health Institute (2018). The entire FY 2018 budget for the Centers for Disease Control and Prevention was about $8 billion.

when applied to catastrophes that could kill many people. But even a much smaller value for the VSL results in a WTP to avert pandemics that is far large than current expenditures.

An important caveat to this paper is our starting assumption in eqn. (1) that social welfare is proportional to the size of the population N. How should we think about the welfare effects of changes in population? This question has received considerable attention by economists and philosophers, but there is no consensus regarding the social value of more or fewer people.[27] As a practical matter, society spends considerable resources to save lives, including the lives of the very old and very young, and prevent life-threatening disasters large and small. Consistent with this, a catastrophic loss of life reduces social welfare in our model. On the other hand, an increase in the natural rate of population growth reduces welfare in our model, consistent with the views of some environmentalists who are concerned about congestion and environmental stress.

References

Aldy, Joseph E. 2019. "Birds of a Feather: Estimating the Value of a Statistical Life from Dual-Earner Families." National Bureau of Economic Research Working Paper 25708.

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 Priebsch. 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.

Ashenfelter, Orley. 2006. "Measuring the Value of a Statistical Life: Problems and Prospects." *The Economic Journal*, 116: C10–C23.

Backus, David, Mikhail Chernov, and Ian Martin. 2011. "Disasters Implied by Equity Index Options." *Journal of Finance*, 66(6): 1969–2012.

Barnichon, Regis, Christian Matthes, and Alexander Ziegenbein. 2018. "The Financial Crisis at 10: Will We Ever Recover?" *FRBSF Economic Letter*, 1–4.

---

[27]See, e.g., Razin and Sadka (1995), Golosov, Jones and Tertilt (2007), Broome (2004) and Blackorby, Bossert and Donaldson (2005).

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121: 823–866.

Barro, Robert J. 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review*, 99: 243–264.

Barro, Robert J. 2014. "Environmental Protection, Rare Disasters, and Discount Rates." *Economica*, 82(325): 1–23.

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.

Blackorby, Charles, Walter Bossert, and David Donaldson. 2005. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics.* Cambridge University Press.

Bohn, Henning, and Charles Stuart. 2015. "Calculation of a Population Externality." *American Economic Journal: Economic Policy*, 7(2): 61–87.

Broome, John. 2004. *Weighing Lives.* Oxford University Press.

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.

Enserink, M. 2004. "Avian Influenza: Catastrophe Waiting in the Wings?" *Science*, 306(5704): 2016.

Fan, Victoria Y., Dean T. Jamison, and Lawrence H. Summers. 2018. "The Loss from Pandemic Influenza Risk." In *Disease Control Priorities, 3rd ed., Volume 9.* , ed. D.T. Jamison, et al, 347–358. World Bank.

Fernald, John G., Robert E. Hall, James H. Stock, and Mark W. Watson. 2017. "The Disappointing Recovery of Output after 2009." *Brookings Papers on Economic Activity*, 1–58.

Golosov, Mikhail, Larry E. Jones, and Michèle Tertilt. 2007. "Efficiency with Endogenous Population Growth." *Econometrica*, 75(4): 1039–1071.

Greenstone, Michael, Stephen P. Ryan, and Michael Yankovich. 2012. "The Value of a Statistical Life: Evidence from Military Retention Incentives and Occupation-Specific Mortality Hazards." Unpublished manuscript.

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.

Harvard Global Health Institute. 2018. "Global Monitoring of Disease Outbreak Preparedness: Preventing the Next Pandemic." Harvard University Report.

Johnson, Niall P.A.S., and Juergen Mueller. 2002. "Updating the Accounts: Global Mortality of the 1918-1920 'Spanish' Influenza Pandemic." *Bulletin of the History of Medicine*, 76(1): 105–115.

Jones, Charles I. 2016. "Life and Growth." *Journal of Political Economy*, 124(2): 539–578.

Jones, Kate E., Nikkita G. Patel, Marc A. Levy, Adam Storeygard, Deborah Balk, John L. Gittleman, and Peter Daszak. 2008. "Global Trends in Emerging Infectious Diseases." *Nature*, 451(7181): 990–993.

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 Univ. 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.

Martin, Ian W.R., and Robert S. Pindyck. 2017. "Averting Catastrophes that Kill." National Bureau of Economic Research Working Paper 23346.

Nakamura, Emi, Jon Steinsson, Robert Barro, and José Ursúa. 2013. "Crises and Recoveries in an Empirical Model of Consumption Disasters." *American Economic Journal: Macroeconomics*, 5(3): 35–74.

National Academy of Medicine. 2016. *The Neglected Dimension of Global Security: A Framework to Counter Infectious Disease Crises.* National Academies Press.

Pindyck, Robert S., and Neng Wang. 2013. "The Economic and Policy Consequences of Catastrophes." *American Economic Journal: Economic Policy*, 5(4): 306–339.

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.

Razin, Assaf, and Efraim Sadka. 1995. *Population Economics.* MIT Press.

Rohlfs, 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.

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. 2018. *Pricing Lives: Guideposts for a Safer Society.* Princeton University Press.

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.

Voigtländer, Nico, and Hans-Joachim Voth. 2013. "The Three Horsemen of Riches: Plague, War, and Urbanization in Early Modern Europe." *Review of Economic Studies*, 80(2): 774–811.

Young, Alwyn. 2005. "The Gift of the Dying: The Tragedy of AIDS and the Welfare of Future African Generations." *Quarterly Journal of Economics*, 120: 423–466.