NBER WORKING PAPER SERIES

CALCULATING THE COSTS AND BENEFITS OF ADVANCE PREPARATIONS FOR FUTURE PANDEMICS

Rachel Glennerster
Christopher M. Snyder
Brandon Joel Tan

Working Paper 30565
http://www.nber.org/papers/w30565

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022, Revised June 2023

The paper, previously circulated as National Bureau of Economic Research working paper no. 30565, was substantially improved following the advice of two anonymous reviewers and the editor, Kerem Co ar. The authors are grateful to them as well as to Nicole Bouvier, Michael Kremer, Josh Morrison, Eric Mannes, Rui Mano, Tristan Reed, and participants at the IMF/BAM/ IMFER Conference on Transformational Recovery held in Morocco and the Infectious Diseases/ Global Health and Emerging Pathogens Institute (GHEPI) Seminar Series at Icahn School of Medicine at Mount Sinai for insightful comments. Yijing Zhang provided excellent research assistance. The views expressed are those of the authors and do not reflect those of the IMF, its Executive Board, or the National Bureau of Economic Research.

Calculating the Costs and Benefits of Advance Preparations for Future Pandemics
Rachel Glennerster, Christopher M. Snyder, and Brandon Joel Tan
NBER Working Paper No. 30565
October 2022, Revised June 2023
JEL No. H44,I15,I18,L65

ABSTRACT

While Covid-19 vaccines were developed and deployed with unprecedented speed, their widespread introduction could have been accelerated—saving millions of lives and trillions of dollars—had more vaccine capacity been available prior to the pandemic. Combining estimates of the frequency and intensity of pandemics with estimates of mortality, economic-output, and human-capital losses from pandemics of varying severities, we calculate that the present value of global social losses from the stream of future pandemics can be expected to be nearly $18 trillion —over $700 billion each year going forward. According to our model, a program spending $60 billion up front to expand production capacity and supply-chain inputs for vaccines and $5 billion annually thereafter would be sufficient to ensure production capacity to vaccinate 70% of the global population against a new virus within six months. The program would generate an expected net present value (NPV) gain of more than $500 billion over the status quo of delaying investment until a pandemic arrives. A program undertaken by the United States alone would generate an expected NPV gain of over $60 billion.

Rachel Glennerster
Saieh Hall for Economics, Office 432
University of Chicago
5757 S University Avenue
Chicago, IL 60637
and NBER
rglennerster@uchicago.edu

Christopher M. Snyder
Department of Economics
Dartmouth College
301 Rockefeller Hall
Hanover, NH 03755
and NBER
chris.snyder@dartmouth.edu

Brandon Joel Tan
International Monetary Fund
btan2@imf.org

1 Introduction

By 2024, it is estimated that the Covid-19 pandemic will have reduced economic output by $13.8 trillion relative to pre-pandemic forecasts (International Monetary Fund 2022). The pandemic resulted in an estimated 7–13 million excess deaths (*Economist* 2022) and an estimated $10–$17 trillion loss of future productivity and earnings from school disruption (Azevedo et al. 2021). Such devastating losses from a pandemic are not new: some sources estimate that the 1918 flu killed 2% of the world's population and reduced GDP by 6% (Barro, Ursúa, and Weng 2020) and that the Black Death killed 30% of Europe's population (Alfani 2022).

Vaccines against Covid-19 were developed, approved, and distributed at record speed, mitigating both the economic and social losses. Yet many countries waited years for sufficient supply, resulting in millions of deaths and trillions in economic damage that could have been averted. Estimates suggest that accelerating the capacity to produce 1.5 billion courses of vaccine annually by just three months would have had a social value of $1.3 trillion (Castillo et al. 2021). This large social value eclipsed the revenue that pharmaceutical companies were earning from Covid-19 vaccines. The gap between social and private returns suggests private companies will have inadequate incentives to invest in the capability to produce pandemic vaccines both before and during pandemics, calling for deliberate public policies to improve those incentives.

This paper makes the economic case for public investments in vaccine-manufacturing capacity in advance of the next pandemic. The first step in estimating the social return from such a program is to estimate the arrival rate of pandemics of different intensities. We follow (and extend) Marani et al. (2021, corrected in 2023), who find that the distribution of pandemic intensity conditional on one arriving has been fairly stable over the long time series they study (1600 to present), closely following a generalized Pareto distribution. They then use more recent data to estimate the arrival rate of pandemics of whatever intensity, allowing recent medical advances, climate change, and other factors to shift the arrival rate. We combine these projections with estimates of the economic costs of pandemics of different intensities from the economic literature (Barro, Ursúa, and Weng 2020; Keogh-Brown et al. 2009; Huber, Finelli, and Stevens 2018; United Nations Development Programme 2017; International Monetary Fund 2022), the value of lives lost, and losses from school closures (Azevedo et al. 2021).

We estimate that the expected present value of the stream of global social losses from future pandemics amounts to $17.8 trillion. Assuming a 4% social discount rate, this is equivalent to $712

billion each year going forward, more than the budget of U.S. Department of Defense in the baseline year (U.S. Department of Defense 2020). To account for the uncertainties in predicting pandemics, we estimate pandemic losses under a range of scenarios: under the most optimistic scenarios, expected losses exceed $350 billion annually, rising to over $2.1 trillion annually under some more pessimistic but still plausible scenarios.

With these estimates of large expected losses in hand, we proceed to estimate the return to a combination of up-front and continuing investments that reduce the damage caused by pandemics. We evaluate the benefits of this program compared to the status-quo policy investing in vaccine capacity at the maximum scale and speed that can be achieved if investment is delayed to the start of a pandemic. We estimate that expanding production capacity for vaccines and supply-chain inputs so that there is sufficient production capacity to vaccinate 70% of the global population against a new virus within six months would generate an expected net present value (NPV) of $735 billion, a gain of more than $500 billion over the status quo of delaying investment until the start of the pandemic, achieved by cutting the time to complete that vaccination campaign by more than half. These estimates account for the risk that vaccines might fail and that vaccine hesitancy might be high. According to our model, the program would require an up-front investment on the order of $60 billion and $5 billion to be spent each year thereafter.

Our analysis focuses on a global program since this would extend the benefits of accelerated and expanded capacity to the most people. However, full participation among all countries may not be achievable in equilibrium absent an international agreement. The reason is that advance investment by one country reduces its demand for in-pandemic capacity, generating positive spillovers for others, possibly leading to free riding. If countries fail to strike an international agreement, it is not an equilibrium for none of them to invest in advance. High-income countries would reap large net benefits from going it alone and investing in advance capacity. For example, we calculate that if the United States alone undertook an advance-investment program that enabled it to vaccinate 70% of its population within six months, this would generate an NPV of $77 billion, a gain of $61 billion ($183 per capita) over the status quo. The benefits extend to middle- and low-income countries as well. For example, we calculate that an advance-investment program undertaken by Brazil alone would have an NPV of $19 billion, a $15 billon ($70 per capita) gain over the status quo.

While our main focus is on preparatory investments in vaccine production capacity, the

logic of our arguments extends to a broader set of investments that could mitigate pandemic harm including research, development, and production of a universal coronavirus vaccine, development of broad-spectrum antivirals, and new antibiotics and investments that would streamline vaccine approval during pandemics (such as the rules under which human challenge trials would be appropriate).

Our paper builds on work in the scientific literature on the frequency of pandemics including Marani et al. (2021), Carlson et al. (2021), and Bernstein et al. (2022) as well as an economics literature on the economic costs of specific past pandemics (cited above) and the expected costs of future pandemics (Fan et al. 2018; Nakamura et al. 2013, Keogh-Brown et al. 2009, Martin and Pindyck 2015, Martin and Pindyck 2017, Jaimeson and Summers 2015). The epidemiological component of our analysis draws on methods that account for progress in the world's ability to moderate the impact of pandemics in modern times. We expand on previous literature on the cost-effectiveness of pandemic preparedness including increased surveillance (Bernstein et al. 2022), research and development (Crank et al. 2019), vaccines including building up stockpiles (Meltzer, Cox, and Fukuda 1990; Prager, Wei, and Rose 2017; Schoenbaum 1987), and all of the above (Yamey et al. 2017). Our model of the benefits of accelerating the pace of vaccination during a pandemic and the net benefits of investing in supply capacity draws on Castillo et al. (2021), Ahuja et al. (2021), and Athey et al. (2022).

The paper is outlined as follows. Section 2 reviews and adapts analysis on the frequency and epidemics of varying intensities. After synthesizing estimates of various losses from epidemics, Section 3 derives a relationship between total losses and mortality intensity. Section 4 provides a conceptual discussion of a program to accelerate widespread vaccine availability by investing in pre-pandemic preparedness. Section 5 sets out the formal model for evaluating the program's expected net benefits, and Section 6 presents the results. Section 7 discusses whether international cooperation is needed to achieve these gains, and Section 8 concludes. Details behind some derivations are provided in the appendix.

2 Probability of Future Pandemics

The challenge in estimating the arrival rate of future pandemics is that large pandemics only happen rarely, requiring a long time series to observe enough events to accurately estimate their frequency. However, over time the arrival rate changes due, for example, to medical advances

(e.g., the arrival of antibiotics) and environmental changes. To address this challenge, we follow the methodology of Marani et al. (2021, revised in 2023). They use the long-run history of epidemics (from 1600 to present) to estimate the distribution of relative pandemic intensity (i.e., when a pandemic arrives, how much more likely it is to be small than large). They use recent data (last 20 years) to estimate whether pandemics of whatever size have become more or less frequent recently.

Marani et al. (2021) document 476 significant epidemics since the year 1600, of which 271 have data on duration and deaths, forming the main basis of their estimations. Defining the intensity i of an epidemic in a year to be the associated mortality expressed in terms of deaths per thousand population, the authors show that the distribution of epidemic intensity, conditional on one arriving in a year, is stable over the long time series, well described by a generalized Pareto distribution having cumulative distribution function

$$\Phi_0(i) = \begin{cases} \alpha & i = [0, \mu') \\ 1 - (1-\alpha)\left[1 + \dfrac{\xi(i-\mu')}{\sigma}\right]^{-1/\xi} & i \in [\mu', \mu''] \\ 1 & i \in [\mu'', 1000], \end{cases} \quad (1)$$

where $\mu' = 10^{-3}$ is the threshold below which epidemics are too small to leave a detectable record, $\alpha = 0.62$ is the probability that the epidemic is below the threshold of detectability, and $\sigma = 0.0113$ and $\xi = 1.41$ are shape parameters estimated by the authors via maximum likelihood.[1] The complement to the cumulative distribution function, $\bar{\Phi}_0(\mathrm{i}) = 1 - \Phi_0(i)$, sometimes called the exceedance probability, has the useful interpretation as the annual probability that an epidemic with at least intensity i occurs.

Marani et al. (2021) do not provide guidance on how to extrapolate their estimates beyond the support of their data. Since we will be integrating over the distribution of pandemics of any conceivable size to compute expected pandemic losses, our approach requires such extrapolation. Our extrapolation strategy is embedded in equation (1). The support of the distribution has a natural upper bound at 1,000 deaths per thousand, representing human extinction. It remains to specify the mass in the tail of the distribution below this natural upper bound and above the support of their data. Small changes in the mass of its fat tail can have a large influence on the expected value of a Pareto random variable, but extrapolation in this interval is challenging given the expanding

[1] These and other parameters introduced throughout this section are collected in Table 3 for reference.

confidence intervals there and the inevitably growing inaccuracy of the Pareto law as intensity approaches population size. The conservative approach that we adopt to address this issue truncates maximum epidemic intensity at the highest credible reports of deaths for a pandemic in their data. Based on estimates that the 1918 flu killed 100 million people (Sanford 1979, Burnet 1979), dividing by global population at that time and dividing by the three-year duration of that pandemic yields a truncation threshold of $\mu'' = 17.8$ deaths per thousand per year.[2,3] We perform various sensitivity analyses for alternative truncation thresholds, one halving the truncation threshold, another doubling it, another removing it.

Starting with the basic distribution in (1), Marani et al. (2021) transform it in a way that maintains a constant distribution of relative intensities but allows the arrival rate of epidemics to vary over time. This approach allows them to exploit the long historical record to precisely estimate the distribution of relative intensities, overcoming the challenge that large pandemics are a "black swan" event, requiring a long time series to achieve a reasonable sample of them. The authors then use the frequency of recent pandemics to estimate the general arrival rate of any epidemic above a threshold size under modern conditions. Most pandemics within a given time band will be small but occur with enough frequency to provide a good estimate of an overall arrival rate, which can be extrapolated to pandemics of any size under the assumption that the historical distribution of relative intensities has remained constant.

Formally, Marani et al. (2021) transform equation (1) via the metastatistical extreme value distribution (MEVD), averaging the distribution of the maximum order statistic from n_t draws corresponding to the number of epidemics in year t. The resulting formula is

$$\Phi(i) \approx \frac{1}{w}\sum_{t=1}^{w} \Phi_0(i)^{n_t}, \tag{2}$$

where w is the width of the window of years under consideration. Equation (2) has a particularly

[2] Marani et al. (2021) record the midpoint of the range of estimates from Patterson and Pyle (1991), 32 million, as mortality from the 1918 flu in their data. Patterson and Pyle's (1991) themselves report an upper estimate is 39.3 million deaths. The debate about global mortality from the 1918 flu continues in the literature (Spreeuwenberg, Kroneman, and Paget 2018; Chandra and Christensen 2019). There is no debate that the 1918 flu was extremely intense in certain regions. For example, Patterson and Pyle (1991) estimate that deaths from a single wave in Asia could have been as high as 33 million. Using this estimate would double the truncation point we are using to $\mu'' = 34.2$ deaths per thousand.

[3] Our adjustment requires an atom of mass $\bar{\Phi}_0(\mu'')$ to be added at $i = \mu''$. Marani et al. (2021) leave the distribution of intensity unspecified for $i < \mu'$. The specification in equation (1) fills this gap in a conservative way by adding an atom of mass a at $i = 0$ and positing zero mass for $i \in (0, \mu')$.

simple form if, following Marani et al. (2021), we take the window to be the most recent 20 years in their dataset, during which, according to their Supplementary Figure S1(a), there were 13 years without a detectable epidemic and seven years with one. Substituting those numbers, (2) becomes

$$\Phi(i) = \frac{1}{20}[13 + 7\Phi_0(i)]. \tag{3}$$

Figure 1 graphs the exceedance probability, $1 - \Phi(i)$, associated with equation (3). The annual probability of a pandemic at least as severe as Covid-19 is about 1%, implying that such a pandemic is a one in 105-year event.

The rate of epidemics over the last 20 years, which factors into $\Phi(i)$ as we have just seen, turns out to be historically low, reflecting two opposing forces operating recently.[4] Modern technology has allowed society to mitigate the death toll from pandemics. The invention of antibiotics sharply reduced the occurrence of the plague and other bacterial outbreaks. Better hospitals and medical care have also helped cut mortality. Working in the opposite direction, models of the effect of climate change suggest increasing zoonotic spillovers (the transmission of viruses and other parasites from animals to humans), increasing the frequency of future epidemics (Carlson et al. 2021).

The use of the most recent 20 years of data for our baseline forecast allows it to reflect current conditions. It is challenging to forecast how the rate of pandemic arrival will continue to change in the future. We simply assume that the current arrival rate will persist. To account for unknown changes in the arrival rate among other uncertainties inherent in forecasting future pandemics, we analyze the sensitivity of our estimates to changes in a variety of assumptions and parameters.

3 Social Losses in a Pandemic

3.1 Approach

Having estimated the arrival rate of pandemics of various intensities, we next need to pair that with estimates of the losses to society conditional on experiencing a pandemic of a given intensity to

[4] Previous estimates of the expected economic losses from future pandemics have tended to use a longer time horizon to estimate the intensity and frequency of pandemics, which may overestimate expected losses by underweighting progress made in combatting pandemics. This helps explain why our estimates are lower than Fan et al. (2018) despite their estimates covering losses from influenza epidemics alone.

calculate expected harm from the next pandemic. The literature suggests that expected losses from epidemics are dominated by high-intensity pandemics that come along only rarely. In this section, we seek to refine existing estimates of the possibly nonlinear relationship between the intensity i of an epidemic (measured by relative mortality) and associated social losses.

Some of the literature focuses on a particular category of social loss, say just deaths or just the shortfall in economic output.[5] Here, we seek a comprehensive measure that, in addition to these two categories, includes longer-term losses from the decline in human capital associated with closures of school and training programs. Our approach will be to use the best available information from the literature to map pandemic intensity into each category of loss and then sum the categories to obtain total losses to society. Our total measure will still be conservative since it will not include difficult-to-estimate categories such as the disutility of social distancing and pain and suffering from sickness.

Conditional on an epidemic of intensity i arriving in year t, let $ML_t(i)$ denote mortality losses from that pandemic, $OL_t(i)$ denote economic-output losses, and $LL_t(i)$ denote learning losses experienced in that year. Let $ML_0(i)$, $OL_0(i)$, and $LL_0(i)$ denote the analogous expressions for the base year $t = 0$. We will discuss the estimation of each loss category in turn, starting with $ML_t(i)$.

3.2 Mortality Losses

Mapping intensity i into mortality losses $ML_t(i)$ requires two steps. First, intensity i, which is a proportion, needs to be converted into expected deaths in year t, denoted d_t, which is a level. Marani et al. (2021) define i as deaths per 1,000 population, i.e., $i = d_t/(N_t/1000)$, where N_t denotes the global population in year t. Inverting, $d_t = iN_t/1000$.

Next, expected deaths d_t needs to be translated into a monetary value. We use Sweis's (2022) estimate, $1.3 million, of the value of a statistical life (VSL) for the world population in the base year 2021.[6] Mortality losses in the base year are thus

$$ML_0(i) = \left(\frac{\$1.3 \text{ million} \cdot N_0}{1000}\right) i. \tag{4}$$

[5] Fan et al. (2018) combine value of lives lost with falls in economic output while Keogh-Brown et al. (2009) also add losses due to school closures.

[6] This estimate has several advantages for our purposes. First, it is a global measure. Most available VSL estimates are for individual countries, with a focus on high-income countries, which would overstate the global VSL since VSL

Mortality losses in year t, $ML_t(i)$, can be derived from $ML_0(i)$ under assumptions on global annual growth rates for GDP per capita and population. Assume GDP per capita grows at a constant rate y. Our calculations set $y = 1.6\%$, the long run global rate projected through 2060 by OECD (2022). Population growth is a more complex issue. If the current slowdown in population growth continues, the world will eventually experience population declines, though presumably these declines would slow before they lead the population to disappear. An additional complexity raised by a changing population is that it changes the nature of the optimal vaccine program. It is challenging enough to model the costs of a program targeting a fixed population let alone one that adjusts dynamically to a first growing and then shrinking population. We finesse these complexities by fixing the population size at the baseline level, N_0, for the analysis. Our results will then apply to a program optimized for current citizens, understating that installing even more capacity to accommodate larger future generations would have positive option value. Under the preceding assumptions and parameter values,

$$ML_t(i) = (1 + y)^t ML_0(i). \tag{5}$$

3.3 Economic-output Losses

Turn next to the estimation of economic-output losses $OL_t(i)$ conditional on the arrival of an epidemic of intensity i in the base year. We include in this category only short-run deviations in economic output from trend caused by pandemics, deliberately excluding longer-term losses such as reduction in future wages due to declines in human-capital, covered by a later calculation.

Table 1 shows the five major pandemics over the previous century for which we could find a credible estimate of economic-output losses. All but one of the estimates come from studies in the literature that use deviations from global GDP trend as the main determinant of economic losses.[7] The last column puts all the economic-loss estimates quoted by the indicated studies on the same metric, an annual global loss in percentage terms, denoted Δ. The table also shows the

estimates increase with income. For example, Sweis's (2022) estimate of the VSL in the United States is $7.2 million, over five times higher than her global estimate. Second, the author applies a rigorous methodology due to Becker (2007) to data on health risks during the Covid-19 pandemic, a relevant domain to our analysis, and time period close to our 2021 base year. Third, for comparable countries, Sweis's (2022) VSL estimates are close to those in other well-cited studies (e.g., Viscusi 2020) but typically slightly lower, suiting our approach of maintaining conservative assumptions. Finally, other policy studies adopt a similar estimate; for example, Ahuja et al. (2021) take the global VSL to be $1.2 million.

[7] We have adjusted for the fact that some studies include the value of lives lost and others do not by taking out the value of lives lost from those that include it and adding it separately using our own valuation later in the calculation.

estimate of intensity taken from sources on which Marani et al. (2021) drew with the exception of Covid-19, which is not in their data.

The five rows in Table 1 provide a sample that can be used to estimate the relationship between Δ and i. Figure 2 plots those variables using log scales on the axes along with a regression line estimated via ordinary least squares,

$$\underset{}{\ln \Delta =} \underset{(0.56)}{0.74} + \underset{(0.08)}{0.46} \ln i, \tag{6}$$

where standard errors are reported in parentheses below coefficient estimates. The regression exhibits remarkably good fit, with $R^2 = 0.92$.

The regression can be paired with the distribution of pandemic intensity from equation (3) to compute expected annual economic output losses from pandemics in the base year:

$$OL_0(i) = Y_0 N_0 \frac{\Delta(i)}{100}, \tag{7}$$

where, denoting per-capita GDP in year t by Y_t, Y_0 is per-capita GDP in the base year. Exponentiating both sides of regression (6) yields $\Delta(i) = 2.09 i^{0.46}$. Dividing by 100 converts the percentage into a proportion. Multiplying by $Y_0 N_0$, which equals GDP in the base year, converts a proportional loss into a loss in levels. $OL_t(i)$ can be derived from $OL_0(i)$ by analogy to equation (5). We use 2021 as the base year, the most recent year that World Bank data are available for Y_0 and N_0, and the same year used for the VSL estimate used to value mortality losses.

3.4 Learning Losses

The final category of losses we consider is $LL_t(i)$, learning losses from the arrival of an epidemic of intensity i in year t. Under the assumption that school disruption moves in line with the disruption to other economic behavior, we take $LL_t(i)$ to be proportional to $OL_t(i)$. We derive the proportionality constant by examining the ratio of economic-output losses and learning losses for Covid-19 for which we have good estimates of both loss categories. We take the conservative end of the World Bank's estimated range for learning losses from Covid-19 at an aggregate $10 trillion in lifetime earnings in present value (Azevedo et al. 2021).[8] For economic-output losses from

[8] Azevedo et al. (2021) use the correlation between years of schooling and wages to calculate the return to an additional year of schooling and hence the cost of closed schools. If wages reflect the worker's marginal product and private returns to education reflect social returns, then future wage losses will be a good measure of future GDP losses. While

Covid-19, we take the International Monetary Fund's (2022) estimate of a $13.8 trillion shortfall relative to pre-pandemic forecasts due to the pandemic.[9] Taking the ratio of the two estimates, we have

$$LL_t(i) = \frac{10}{13.8} OL_t(i). \tag{8}$$

3.5 Total Losses

Let $TL_t(i) = ML_t(i) + OL_t(i) + LL_t(i)$ denote total social losses from all three categories conditional on an epidemic of intensity i arriving in year t. For program evaluation, this conditional loss measure needs to be converted into an unconditional one reflecting the distribution of epidemic intensity estimated in Section 2 since the intensity of future pandemics is uncertain when advance investment is undertaken. We will build up to the unconditional loss measure we ultimately use in a series of steps.

A straightforward measure of unconditional losses is the present value of expected losses from the stream of pandemics into the future,

$$PV(\overline{TL}) = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^{t+1} E(TL_t(i)), \tag{9}$$

where r denotes the social discount factor, which we set at $r = 4\%$ in the baseline scenario.[10] The factor $E(TL_t(i))$ denotes expected pandemic losses in year t:

$$E(TL_t(i)) = \int_0^{\infty} TL_t(i)\, d\Phi(i). \tag{10}$$

Mincer (1974) equations do not measure the causal effect of education on earnings, Duflo (2001) concludes that causal estimates of the income benefits of education are close to Mincer-regression estimates where the two can be compared. In a Spence (1973) model, education can be rewarded with higher wages even if it does not increase productivity, leading private returns to exceed social returns to education. Positive spillovers from education would lead social returns to exceed private returns. We proceed by assuming that Mincer regressions give an adequate estimate of the private returns to education and that social returns weakly exceed private returns on average. If social returns to education strictly exceed private returns on average, our measure of learning losses will be conservative.

[9] We take these loss estimates for Covid-19 losses as they are reported by our sources: not annualized figures but accumulated losses over a multiyear pandemic. Since both sources use this same accounting frame, taking the ratio of them produces the proper proportionality constant.

[10] Gollier and Hammitt (2014) discuss the debate in the economic literature on the appropriate value to assume for the social discount factor r. Our baseline value $r = 4\%$ is within the range of the literature but toward the upper end of the interval between 1% and 4% recommended by Gollier and Hammitt (2014) and thus conservative because higher social discount rates generate lower present discounted values.

Table 2 reports a convenient rescaling of $PV(\overline{TL})$, interpreted as expected annual pandemic losses. Formally, let $AV(\overline{TL})$ denote the constant expected loss that if experienced in perpetuity would generate the present value in equation (9), i.e., $AV(\overline{TL})/r = PV(\overline{TL})$. We show in the appendix

$$AV(\overline{TL}) = rPV(\overline{TL}) = \frac{r}{r-y}E\big(TL_0(i)\big). \tag{11}$$

Table 2 reports the expected annual pandemic losses in total, $AV(\overline{TL})$, as well as for the component losses $AV(\overline{ML})$, $AV(\overline{OL})$, and $AV(\overline{LL})$, defined analogously to $AV(\overline{TL})$. In the baseline scenario, expected annual pandemic losses are $AV(\overline{TL}) =$ \$712 billion annually. In other words, the world can expect to lose \$712 billion every year to pandemics going forward. Mortality losses account for 73% of the total, followed by economic-output losses (16% of the total), followed by learning losses (11% of the total).

3.6 Sensitivity of Pandemic-Loss Estimates

We already discussed the first row of Table 2, which provides our baseline results for annual pandemic losses. Subsequent rows analyze the sensitivity of the results to changes to parameters and changes to the assumptions behind the distribution of pandemic intensities. The first alternative scenario cuts the probability of pandemic arrival in half, which cuts the expected loss estimates in half. Still, total expected annual pandemic losses are a substantial \$356 billion. Doubling the probability of pandemic arrival increases expected losses to \$1.2 trillion.

The next set of alternative scenarios revisit the challenge of extrapolating the distribution of pandemic intensities outside of the data that Marani et al. (2021) used to estimate their power-law parameters. Mathematically, the estimated power law must break down for extreme intensities i approaching 1,000 (corresponding human extinction), but it is hard to know where beyond the range of the data this breakdown occurs. In the baseline, we adopt the conservative approach of setting the truncation point μ'' on intensity at the highest credible report of intensity for pandemics in their data ($\mu'' = 17.8$, based on estimates for the 1918 flu cited above). One alternative halves that truncation point, another doubles it, another eliminates truncation (apart from the natural maximum of human extinction). Social losses are reduced when the truncation point is tightened and increased when the truncation point is relaxed. Eliminating truncation below human extinction

more than doubles $AV(\overline{TL})$ to \$1.5 trillion.[11]

The last set of scenarios analyze the sensitivity of the loss estimates to changes to the VSL, growth rate of GDP per capita y, and social discount rate r. Perhaps the most consequential change is that $AV(\overline{TL})$ increases to \$2.1 trillion when r is reduced to 2%.[12]

Several broad observations about our baseline estimates can be drawn from the sensitivity analyses. First, they are robust. In no row does $AV(\overline{TL})$ fall below \$356 billion, which is still half of a substantial baseline. Second, our baseline estimates may be quite conservative. We truncated the distribution of epidemic intensity out of an excess of caution about extrapolating the estimates of Marani et al. (2021) too far beyond the range of their data. Removing this truncation increases $AV(\overline{TL})$ substantially. We also erred on the conservative side by setting the social discount rate to $r = 4\%$. It is not uncommon for studies to set $r = 2\%$; at this rate $AV(\overline{TL})$ rises substantially to \$2.1 trillion.

3.8 Expected Losses from Next Pandemic

The criterion we ultimately use for program evaluation is more complicated than $PV(\overline{TL})$ and $AV(\overline{TL})$, requiring some discussion. The criterion involves the expected present value of social losses—not from the stream of all future pandemics—but just from the next significant epidemic. Denote this concept by $PV(\overline{TLN}^*)$, where the N suffix stands for next pandemic and the star superscript indicates that to qualify for the next significant pandemic, the epidemic must exceed some threshold intensity i^*.

As will be seen, using $PV(\overline{TLN}^*)$ as the loss measure rather than $PV(\overline{TL})$ or $AV(\overline{TL})$ facilitates modeling the status-quo program to which our proposed advance-investment program will be compared. We presume that this status-quo policy would not roll out a global vaccination campaign for a minor epidemic but only one of significance. We also presume that at least some of the capacity that this program installs to mitigate harm from the next pandemic would be

[11] The sensitivity analyses with respect to where the intensity distribution is truncated highlight the implications of a power-law distribution's fat tails. Varying the truncation point for a normal distribution (which has thin tails) hardly matters if the truncation point is extremely high. With a power-law distribution, by contrast, varying a truncation point in an extreme range beyond the data can have a substantial effect on expected values.

[12] It is unsurprising that present values grow large when the discount rate shrinks. Here, however, $AV(\overline{TL})$ is not a present value of a stream of losses but an annualized loss. The annualized loss still grows large because future losses—reflecting growing GDP per capita with time—weigh more heavily in the annualized figure.

retained to use in pandemics after that. To absolve ourselves from having to guess how much capacity would be retained in the absence of a coordinated program to do that, we effectively cut the future off after the arrival of the next pandemic. We argue that that modeling device leads to a conservative evaluation of the benefits of our proposed advance-investment program since the proposed program generates more capacity in the next pandemic, leading to weakly more capacity available in epidemics after that.

To derive an expression for $PV(\overline{TLN}^*)$, let $h_t(i^*)$ denote the hazard of a significant epidemic, i.e., the probability that an epidemic of at least intensity i^* arrives in year t conditional on no epidemic of at least that intensity having yet arrived by then since base year 0. We have

$$h_t(i^*) = \Phi(i^*)^t[1 - \Phi(i^*)]. \tag{12}$$

Then

$$PV(\overline{TLN}^*) = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^{t+1} h_t(i^*) E(TL_t(i)|i \geq i^*), \tag{13}$$

The last factor is the expectation of social losses from an epidemic of at least intensity i^* certainly arriving in year t:

$$E(TL_t(i)|i \geq i^*) = \int_{i^*}^{\infty} \frac{TL_t(i)}{1 - \Phi(i^*)} d\Phi(i). \tag{14}$$

The baseline scenario sets the threshold intensity i^* for a significant pandemic to be half the intensity of Covid-19 or worse.

To compare the related loss measures, in the baseline, the expected present value of losses from the whole stream of future pandemics is $PV(\overline{TL}) = \$17.8$ trillion. The expected present value of losses from next pandemic at least half as intense as Covid-19 is $PV(\overline{TLN}^*) = \$8.4$ trillion. Neither $PV(\overline{TL})$ nor $PV(\overline{TLN}^*)$ average over years, so both are orders of magnitude larger than $AV(\overline{TL})$, which does average over years to provide an annual measure.

4 Conceptual Discussion of Advance Investment Program

This section motivates and outlines a program of advance investment to accelerate the availability of vaccine capacity in a pandemic. The discussion in this section is conceptual; a formal model of

the program is deferred to Section 5.

4.1 Program Motivation and Basic Design

The estimates from Section 3 suggest that future pandemics are too large a problem to ignore. The expected present value of the future stream of social losses from pandemics is nearly $18 trillion in our baseline estimates, averaging over $700 billion annually. There is a strong case for any program that can mitigate some of the enormous losses from the next significant pandemic at reasonable cost.

The experience from the Covid-19 pandemic highlights a promising possibility. Vaccines were developed and deployed with unprecedented speed. The reduction in mortality bolstered countries' confidence to reopen their economies, reducing economic-output losses. Yet the enormous social losses that still mounted each month beg the question of whether there was room to further accelerate vaccinations. There was roughly a two-year lag between the date Covid-19 was detected and the production of enough vaccine to fully immunize 70% of the world's population. Most of this lag was not the time required to discover or approve effective vaccines but rather the time needed to scale up production after approval. We therefore focus most of our attention in this paper on policies to accelerate production scale up. This focus is partially motivated by our belief that investing in accelerating vaccine production is one of the most cost-effective ways to prepare for future pandemics. This focus does not sacrifice much generality because the analysis of this specific investment program readily applies to other promising pandemic preparations, discussed in the conclusion.

Athey et al. (2022) explain how accelerating and expanding vaccine capacity can mitigate pandemic losses. The existence of long lags between when a facility starts to add production lines for a new vaccine and when those lines are producing at full capacity[13] provide an opportunity for accelerating the availability of capacity. Presuming that every reasonable technological avenue for shortening the time to capacity availability would be exploited in a pandemic, there remains another possibility for accelerating capacity, using a strategy that Athey et al. (2022) and others call "at-risk investment." Production capacity can be expanded before regulatory approval, in

[13] Production lines are technologically complex, requiring months to set up. Obtaining reasonable yields requires skilled technicians to learn by doing in a specific production facility. Each facility must receive independent regulatory approval. All these factors result in long lags before a facility can start producing and additional lags before production fully spins up.

parallel with clinical trials, rather than sequentially, after clinical trials have succeeded and regulatory approval gained. Ordinarily, this strategy would be socially wasteful. Vaccine programs exhibited a 70% failure rate in clinical trials over the past two decades (Lo, Siah, and Wong 2020). If a vaccine fails to be approved, any of the at-risk investment that is difficult to repurpose for other uses is wasted. However, the benefit of speed in a pandemic may justify this high risk of wastage. Athey et al. (2022) calculate that if the capacity for Covid-19 vaccines available by April 2021 were available three months earlier, that would have had a global benefit of $3 trillion.

Not just early capacity but expanded capacity can also accelerate a vaccination program. To the extent that scarce supply is the limiting factor in vaccine distribution, and scarce capacity is the rate-limiting step in supply, doubling capacity can double the rate at which vaccine is rolled out to the population.

There are limits to how much and how quickly capacity can be installed during a pandemic. Short-run supply curves for necessary inputs can be sharply upward sloping and may hit hard constraints. Supply curves tend to be more elastic and constraints more relaxed over the longer term. Thus, there are potential returns from investing in vaccine capacity in advance of the next pandemic rather than waiting until the pandemic arrives.

Here we propose and analyze a program that secures vaccine production capability before a new pathogen emerges. Up-front investment expands total general-purpose vaccine capacity. The investment can fund an expansion of plant and equipment for vaccine manufacture and the expert labor force to manage and run the manufacturing process. The investment could also fund expanded capacity to produce inputs that otherwise would be bottlenecks in the supply chain. An annual fee is paid to reserve some of that vaccine capacity to be quickly switched to the production of a vaccine when a pandemic threat emerges. For concreteness, we will analyze a program that installs enough advance capacity that, when topped up with additional investment undertaken in a pandemic, ensures the world would have enough capacity to fully vaccinate 70% of the population in six months. We factor into our calculations the risk that not all targeted individuals may wish to be vaccinated and thus only a fraction of the benefit of that target coverage may be realized for this or other reasons. The quantitative targets matter less than the key qualitative point that the program we have in mind is ambitious, involving expansive capacity. To evaluate the benefit of this advance-investment program, we compare it—not to the absence of any vaccination—but to a status-quo program that also runs a vaccination campaign against the next significant pandemic

also employing the strategy of at-risk investing but without the benefit of the extra capacity coming from advance investment.

4.2 Gap in Commercial Versus Social Incentives

Can the world rely on commercial markets to provide adequate advance investment in the absence of a dedicated public program? We would argue no for several reasons. In the absence of a dedicated program, commercial incentives for advance investment would presumably come from the promise of high returns from vaccine sales in the next pandemic, possibly decades in the future. However, social and political pressure to keep prices of vaccines low during a pandemic can create an enormous divergence between commercial and social returns to a vaccine. Covid-19 vaccines sold for between $6 and $40, much less than the $5,800 social value of a course of annual capacity in early 2021 as estimated by Castillo et al. (2021). Few incentives were provided for speed. Even a dose price much higher than $40, if it is fixed independent of delivery date, provides little incentive to supply those doses sooner rather than later. But getting a vaccine earlier can mitigate enormous social harm. Thus, we will analyze in effect a procurement program using government or donor-organization resources to procure more capacity in advance and reserve some for pandemic preparedness.

4.3 Coordinated International Program

Our analysis will focus on the total net benefits assuming all countries undertake advance preparedness action. Below we discuss how positive spillovers from advance investment may lead to suboptimal investment (particularly for smaller countries), suggesting some benefits from coordination. However, if no other countries are making advanced investments, individual countries actually have stronger incentives to undertake advance investment unilaterally than what our computations show the average country has in the coordinated global program. Thus, our analysis of a coordinated global program provides a benchmark that is useful for several exercises, setting a goal for countries to coordinate on and showing that advance investment can be incentive compatible for countries acting alone if coordination fails.

The positive spillovers from advance investment are in contrast to the often-expressed concern that investment in vaccine capacity has negative spillovers to other countries. Such concerns loom particularly large if one's mental model is that vaccine production capability is

fixed, constrained by the fixed supply of key inputs. However, one lesson from Covid-19 is that additional investments were able to expand total vaccine supply in a pandemic, and there is reason to think that additional at-risk investment would have expanded it further (Ahuja et al. 2021). As supply is even more elastic in the long than short run, negative spillovers are even less of a concern for the pre-pandemic investments proposed in this paper. Our analysis instead points to the possibility of positive spillovers: countries coming into the next pandemic with advance capacity will compete less hard for new capacity being installed, leaving more supplies for others. Once a country's dedicated capacity has served its population (in the modeled program, taking six months), that capacity can be used to serve other countries. The latter effect is a positive externality from unilateral country investment not picked up by our analysis of a coordinated global program. Hence, our analysis should be regarded as a lower bound on internal and external benefits from a single country's investment.

4.4 Did Covid-19 Already Prepare the World?

Does the world already have enough vaccine capacity as a result of the investments undertaken in the Covid-19 pandemic? We would argue no for several reasons. At its peak, the world was producing 580 million doses of mRNA Covid-19 vaccine per month. At that rate, it would take over a year and a half to produce enough vaccine to cover 70% of the world population with a full course of an mRNA vaccine. To reach this coverage in six months during the next pandemic would require an expansion of mRNA capacity. Instead, in response to the falling demand for Covid-19 vaccines, some existing mRNA capacity is reportedly being shut down (Kay, Makol, and Paton 2022). Our analysis will show that paying the owners of existing capacity an annual fee to keep their existing capacity in place and paying others to build new mRNA capacity would have a high expected return.

The world has considerably more capacity for traditional than mRNA vaccines. However, some of the traditional vaccines that are regularly administered are of sufficiently high value that health authorities might not want facilities producing these to switch to pandemic vaccines even during a pandemic. For example, during Covid-19, very few facilities producing childhood vaccines switched to producing Covid-19 vaccines. Thus, despite the larger production capacity for traditional vaccines, reaching the desired level of reserve capacity will require new capacity to be built.

The mRNA technology is particularly useful for pandemic preparedness as it appears to be easier to scale rapidly than many other vaccine technologies. However, mRNA is still too new a technology to know whether it can replace traditional vaccines for all diseases. There is no guarantee any single technology will provide the most effective vaccines for all viruses. Thus, the proposed program takes a portfolio approach, involving investing in capacity for several technologies to combat future pandemics.

4.5 Further Design Features

Contracts to ensure sufficient vaccine capacity was in place to vaccinate the world or an individual country for the next pandemic would need to guarantee that such capacity was functional and up to date. (During Covid-19 some reserve vaccine capacity failed.)[14] Given the billions of dollars at stake, appropriate monitoring systems could be devised. One way to do this is to allow—indeed encourage—contracted reserve capacity to be used for the production of other vaccines.

We prefer the proposed program contract on advance capacity rather than contracting in advance on doses to be supplied in the next pandemic and leaving it to the producer to arrange requisite capacity. As explained in Section 4.2, producers do not internalize the full social benefit of speed, and so may underinvest in capacity, fulfilling their supply obligations but more slowly than would be socially desirable. Contracts on doses may generate negative spillovers for other countries, pushing other countries down the queue waiting for scarce supplies. Contracts on capacity, on the other hand, can generate positive spillovers for other countries, increasing the rate at which the queue is served.

For concreteness, the program will focus on advance vaccine capacity, but the logic applies to other advance preparations for the next pandemic. Consider advance investments in drug treatments. More research and development could lead to the discovery of new antivirals to mitigate harm in the next pandemic. Securing raw materials has been a constraint in scaling Paxlovid, an effective treatment for Covid-19, a constraint which interim supply-chain investments can relax in the next pandemic. Similar logic applies for investments increasing supply-chain resilience for other products such as personal protective equipment, testing equipment, and so forth. One reason for focusing on vaccines rather than drugs or these other products is that they

[14] Mole (2021) reports on the failure of Emergent BioSolutions to fulfill its reserve contracts to supply Johnson & Johnson vaccine because of cross-contamination problems.

are simpler to produce and easier to scale up than vaccines, so maintaining spare capacity for them is valuable, but not as valuable as for vaccines.

5 Model

This section formally models a program of advance investment in production capacity to accelerate vaccinations in the next pandemic.

5.1 Setup

Starting from base year $t = 0$, the next epidemic of at least intensity i^* arrives with hazard rate $\lambda_t(i^*)$ in year t, where i^* is the minimum intensity evoking a global vaccine response. We consider two basic options for a coordinated global vaccine program: a program of advance preparatory investment or the status quo without. Distinguish variables associated with the status-quo program with a single prime and variables associated with the advance-preparation program with two primes. Omitted primes refer to a generic program option.

Even in the status quo in the absence of advance preparations, if a significant epidemic arrives in year t, as we saw with the Covid-19 pandemic, the world will seek to obtain vaccines as quickly as possible. Conditional on the arrival of an epidemic, let x' denote the capacity (measured in annual courses) for vaccines against that epidemic that the world can possibly obtain without advance preparations. The model will allow more spending to buy more capacity, but disruptions to supply chains and limited input supplies will mean that no more capacity than x' can be installed in the short span of an epidemic. The model takes x' to be an exogenous parameter ($x' = 4.5$ billion annual courses).

5.2 Role of Advance Investment

Advance investment—undertaken when supply chains are more fluid and input supplies are more elastic—can generate more capacity. Indeed, we will suppose that the program specifies sufficient capacity be installed in advance that, in expectation, it achieves an ambitious target of 70% coverage of the population within six months when combined with the maximum capacity that can be installed during an epidemic. Formally, letting x'' denote the ultimate capacity available with advance investment, we have $x''/2 = 0.7N_0$, where the left-hand side is multiplied by one half

since x'' is measured in annual courses, but the target coverage is sought in half that time (six months). Since $x' = 4.5 \text{ billion} < 15.7 \text{ billion} = 1.4N_0$, we have $x'' > x'$.

> **Advantage 1.** The fact that $x'' > x'$ is the main advantage of advance investment in the model: without advance investment, the world simply cannot obtain the capacity it needs a vaccination campaign of the desired scale and speed.

Let z'' denote the amount of advance capacity that supports the higher capacity level x'' desired during a pandemic. Rather than fully endogenizing z'', we will take a more reduced-form approach, specifying an exogenous parameter θ measuring how much of the maximum capacity x' that could possibly be installed in-pandemic that the world avoids having to undertake then to attain the target 70% coverage in six months. When $\theta = 0$, the world leaves itself the task of installing the entirety of x' in-pandemic. When $\theta > 0$, it leaves itself a smaller in-pandemic investment. Either way, the world achieves the target coverage with advance investment. But if $\theta > 0$, more of this investment is undertaken earlier, when the cost of capacity installation is lower, as the cost model will make clear.

> **Advantage 2.** The second advantage of advance investment, which emerges from the model when $\theta > 0$, is that it can save on some of the cost of capacity installation.

Formally, we have $x'' = x' + (1 - \theta)z''$. In the absence of advance investment, by definition, $z' = 0$.

5.3 Vaccine Technology Platforms

Vaccines differ on many dimensions including how much of the virus is presented to the body to stimulate an immune response (some vaccines only introduce a spike protein or sugar while others introduce the whole virus) and how this antigen is introduced into the body. They also differ in how the antigen is cultured. Most vaccines are cultured in eggs or cells (whether mammalian or insect). While production processes are not identical across various cell-cultured vaccines, they share some commonalities. By contrast, mRNA vaccines are produced quite differently, in a cell-

free system. Whether any specific platform can produce a successful vaccine, and if so, which can be developed most quickly or generate the greatest efficacy, will likely vary from disease to disease and involve a considerable element of luck. Repurposing an mRNA vaccine production line from one disease to another may be a simple programming exercise, holding out the promise that it will be much quicker to scale in response to a pandemic. Since the production process for mRNA vaccines is so different from the process for other vaccines, it is likely to be difficult to quickly repurpose mRNA production lines to manufacture other vaccines and vice versa. Switching between production lines that use egg or cell cultures is not trivial—we specify it takes several months—but we do assume different types of cell-cultured vaccines share enough commonalities in their production processes that repurposing capacity from one type to another is eventually possible.[15]

For purposes of the model, we will collect the various vaccine technology platforms into two: $v \in \{m, o\}$, where $v = m$ denotes the mRNA platform and $v = o$ denotes other (mostly cell-cultured) platforms. Each platform could contain many subtechnologies, and each subtechnology could contain several vaccine candidates. The model reflects the presence of multiple candidates within platform by allowing some repurposing of capacity dedicated to unsuccessful candidates for production of successful candidates within that same platform. Let z_m and x_m denote advance and ultimate global capacity installed for mRNA vaccines and z_o and x_o denote those capacities installed for other vaccines. Although the ratios could be endogenized, to simplify modeling we fix the ratio of existing and future vaccine capacity so that 1/3 is devoted to mRNA vaccines and the rest across other vaccine platforms.

Let p denote the probability of that some vaccine is successfully approved as safe and effective for use against disease arising that year. Let p_m denote the probability that only mRNA vaccines succeed, p_o denote the probability that only vaccines using another platform are successful, and p_b denote the probability that both vaccines using mRNA and those using other platforms are successful. Conditional on some vaccine succeeding, these are exhaustive and mutually exclusive events, implying $p_m + p_o + p_b = p$.

Reflecting the advantages discussed in Athey et al. (2022) of at-risk investment—the strategy of allocating capacity to a portfolio of vaccine candidates before any are approved to

[15] The discussion has omitted the DNA technology, as it has yet to yield an approved vaccine. For modeling purposes, one could consider it as being combined into the mRNA technology category.

reduce the lag in scaling up capacity for candidates that turn out to be successful—we make the extreme assumption that the at-risk strategy is used to install all capacity under both the advance-investment program and the status-quo program. While the status-quo program does involve delay, the delay is not so severe as to preclude at-risk investment entirely but severe enough to reduce the amount of at-risk investment that can be undertaken in that short span. Conditional on some mRNA vaccine succeeding, assume a fraction a_m of the capacity was dedicated to successful mRNA candidates and can begin producing right away after approval date τ_A. A fraction b_m is dedicated to unsuccessful mRNA candidates but can be repurposed to produce successful mRNA candidates after a lag of τ_m. The remaining fraction $1 - a_m - l_m$ of capacity cannot be repurposed and is therefore useless during the pandemic. Vaccines using other platforms are modeled similarly, with a fraction a_o of capacity being immediately available upon approval at date τ_A, fraction l_o being available with lag τ_o, and the remaining fraction useless. We assume capacity is not fungible between mRNA vaccines and other vaccines during the short epidemic span.[16]

Let $\tau \in [0, T]$ index continuous time within the span of the epidemic (measured a finer scale—say days or months—than the yearly scale of epidemic arrival indexed by t). Production capacity totaled across all vaccines available at time τ under both platforms is

$$x(\tau) = \begin{cases} 0 & \tau \in [0, \tau_A] \\ s_m x_m a_m + s_o x_o a_o & \tau \in (\tau_A, \tau_A + \tau_m] \\ s_m x_m (a_m + l_m) + s_o x_o a_o & \tau \in (\tau_A + \tau_m, \tau_A + \tau_o] \\ s_m x_m (a_m + l_m) + s_o x_o (a_o + l_o) & \tau \in (\tau_A + \tau_o, T], \end{cases} \tag{15}$$

where T is the duration of the epidemic and s_m and s_o are indicators for whether mRNA and other vaccines are successful, respectively. To avoid a proliferation of cases, we have written equation (15) under the assumption that $\tau_m \leq \tau_o$, reflecting the relationship between the parameter values chosen for program evaluation below. The formula when $\tau_m > \tau_o$ is obvious by analogy. The total number of people vaccinated by time τ^* equals $\int_0^{\tau^*} x(\tau) d\tau$.

[16]Silver (2021) provides accounts of the transfer of Pfizer's mRNA technology to contract manufacturers to scale up Covid-19 vaccine production. Modern technologies for vaccine manufacturing allow for increasing flexibility and can work especially well in repurposing capacity within a vaccine technology platform (Sell et al. 2021). Fill and finish and warehouse capacity is particularly fungible (Mirasol 2021).

5.4 Vaccination Costs

Consider global expenditures on one of the vaccine technologies, $v \in \{m, o\}$. There are four categories of expenditure for this technology. First, advance investment costs k_v per course, for a total sunk expenditure of $k_v z_v$. Second, this investment depreciates and needs to be replenished at rate δ. We assume that during years without a pandemic, such capacity can be rented out to pharmaceutical firms for routine vaccine production to recapture a fraction ϕ of the yearly cost. The expected present value of the effective expenditures (net of rental income) from these two channels through the end of the next pandemic equals

$$(1-\phi)\left[k_v z_v + \sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^{t+1}\Phi(i^*)^t \delta k_v z_v\right]. \tag{16}$$

A third category of vaccine expenditure is the cost of installing capacity during the pandemic needed to bridge the gap between the advance capacity z_v and the capacity x_v ultimately used during the pandemic. Let $K_v(x_v - z_v)$ denote the investment cost, specified following Castillo et al. (2021):

$$K_v(q_v) = \begin{cases} k_v q_v & q_v \leq \beta \\ k_v q_v \left\{\frac{1}{1+\varepsilon}\left[\varepsilon \cdot \frac{\beta}{q_v} + \left(\frac{q_v}{\beta}\right)^{\varepsilon}\right]\right\} & q_v > \beta, \end{cases} \tag{17}$$

where $q_v = x_v - z_v$. According to this specification, the marginal cost of capacity installed during a pandemic is the same k_v as that installed in advance for levels of capacity below a kink point β (taken to be 100 million annual courses in the baseline); but above that kink point, K_v exhibits decreasing returns to scale, which are more severe the higher is $\varepsilon > 0$.

A fourth category of vaccine expenditure is the marginal cost of production c_v per course administered during a pandemic. Since a target 70% of the population is assumed to be covered by the pandemic's end under either program (albeit more slowly without advance investment), the expense from this threshold conditional on being in a pandemic is

$$c_v \cdot \frac{x_v}{x} \cdot 0.7 N_0. \tag{18}$$

As the last two categories are only expended conditional on a pandemic, the expected present value of these expenditures equals

$$\sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^{t+1} h_t(i^*)\left[K_v(x_v - z_v) + c_v \cdot \frac{x_v}{x} \cdot 0.7N_0\right]. \tag{19}$$

Combining the four categories of expenditures and substituting simplified expressions for infinite series, as shown in the appendix, the expected present value of effective (net of rental income) program expenditures on vaccines using technology v can be written

$$(1-\phi)k_v z_v + \frac{(1-\phi)\delta k_v z_v}{r+1-\Phi(i^*)} + \frac{1-\Phi(i^*)}{r+1-\Phi(i^*)}\left[K_v(x_v - z_v) + c_v \cdot \frac{x_v}{x} \cdot 0.7N_0\right]. \tag{20}$$

Total program expenditures can be found by summing (20) over the two technologies $v \in \{m, o\}$. The baseline model assumes that every dollar of program expenditures costs the funder a dollar of surplus. We examine alternative specifications in which the funder's surplus loss is more than dollar for dollar, either because spending is paid for with distortionary taxation or because expenditures are diverted from other worthy projects with high opportunity costs. We will consider extensions in which each dollar of expenditures reduces surplus by $1 + \lambda$ dollars, where λ denotes the so-called social cost of public funds.

Values of the cost-function parameters are drawn from Camillo et al. (2021), who in turn based their parameter choices on news reports and interviews with industry experts. Parameters related to capacity repurposing are drawn from the G20 High Level Independent Panel (2021).

5.5 Vaccination Benefits

We adopt the model (with some simplifications) of the benefits of vaccination that Castillo et al. (2021) used in their analysis of the Covid-19 pandemic. Vaccination mitigates the harm from the epidemic experienced by the population at time τ by the proportion $H(\hat{x}(\tau))$, a function of the fraction $\hat{x}(\tau) = x(\tau)/P_0$ of the world population that has been vaccinated. Following Castillo et al. (2021), assume H is a continuous, concave, piecewise-linear function such that $H(0) = 0$ and $H(\hat{x}(\tau)) = 1$ for $\hat{x}(\tau) > 0.7$, which can be interpreted as the threshold for herd immunity, above which the epidemic is quelled and all harm relieved.[17] The supplemental appendix to Castillo et al. (2021) details why this functional form is a good approximation for vaccination benefits during

[17] Based on data on the proportion of high-risk individuals and the differential burden of the disease on them versus others, Castillo et al. (2021) specify two additional kinks at 0.13 (the fraction of high-risk population) and 0.5, setting $H(0.13) = 0.395$ and $H(0.5) = 0.816$.

the Covid-19 pandemic. The function's concavity reflects the larger benefits from initial courses administered to more vulnerable populations (such as frontline workers and the elderly in the case of Covid-19). Other diseases might have different demographic patterns, but benefits are still likely to be concave given any heterogeneity in harm across groups. As explained in the supplementary appendix to Castillo et al. (2021), a 70% threshold for herd immunity leaves room for an imperfectly effective vaccine to achieve the threshold at plausible transmission rates as indexed by the basic reproductive number $\mathcal{R}_0$.

The harm mitigated by vaccines is scaled by $\gamma \in [0,1]$ to allow for the possibility that some of the harm would have been mitigated by other measures such as improved treatments, better contact tracing, and so forth even without a vaccine, or that a superior vaccine technology comes along that renders advance capacity for existing technologies obsolete. That $\gamma < 1$ might also reflect imperfect vaccine efficacy or unwillingness among a segment of the population to be vaccinated. Recalling the definition of $PV(\overline{TLN}^*)$ as the expected present value of total social losses from the next pandemic evoking a vaccine response, the expected present value of vaccination benefits in the next pandemic equals

$$\gamma PV(\overline{TLN}^*) \int_0^T H(\hat{x}(\tau))\, d\tau. \tag{21}$$

The pandemic's duration T is endogenous, given by the implicit solution to

$$\int_0^T \hat{x}(\tau) d\tau = 0.7. \tag{22}$$

This completes the specification of the model. Table 3 summarizes definitions of parameters used in the model and their baseline values.

6 Results for Program Evaluation

Table 4 presents the baseline results for outcomes from the program of advance investment evaluated relative to the status-quo program. Recall that the status quo we consider is relatively sophisticated: it also runs a vaccination campaign in the next significant pandemic, installs the maximum capacity supply constraints allow, and puts as much of that capacity as possible to work

at-risk to reduce production delays. The only difference is that in the status quo, the world comes into the pandemic without the extra capacity installed in advance.

We consider a baseline design of the advance-investment program that installs an extra 24 billion annual courses of advance capacity for the next pandemic above and beyond that needed to meet typical vaccination demands. We intentionally specify an exogenous capacity level in the baseline design rather than optimizing over this capacity. A much larger program might run into constraints outside of the model—pushing the limits of standard government budgeting and international cooperation or running into the increasing portion of the marginal-cost curve, straining our linear-capacity-cost specification. At the end of the section, we address the question of whether the program could do even better with more or less than the baseline capacity according to the model.

Assuming two doses are required for full vaccination course, and substituting the values for the cost parameters in Table 3 into the cost function from Section 5.4, one can calculate that the baseline design spends $60 billion up front to expand production capacity for vaccines and supply chain inputs. To offset the modeled depreciation, one can calculate that the baseline design requires $5 billion to be spent each year to maintain capacity.

By shifting some expenditures ex ante, the baseline advance-investment program actually reduces expenditures that have to be made ex post during the next pandemic by $32 billion: less capacity is left to be installed during the pandemic, when the inelastic short-run supply entails quite high unit cost for capacity. Netting out rental income earned deploying the advance capacity for vaccines for other diseases between pandemics, the expected present value of the stream of expenditures through the next pandemic are $27 billion higher under the advance program. Compared to the counterfactual of delayed investment, the expected present value of gross benefits is $539 billion higher under advance investment. (To emphasize, this is the expected present value of the additional social losses mitigated just in the next pandemic, not considering pandemics after that, absolving us from having to model how much capacity remains available under alternative programs.) The gross benefits come from achieving the target 70% vaccination coverage about seven months earlier. Advance investment yields an NPV gain of $504 over the status quo.

Table 5 shows the sensitivity of these results to selected assumptions and parameters. For space considerations, parameters that changed the results less than those displayed are omitted (see table notes for complete list). The NPV gain from advance investing always greater than $288

billion over the status quo for any single parameter change considered, whether the probability of epidemic arrival is cut in half, the proportion of losses left to be mitigated by vaccines (so not already mitigated, say, by effective treatments) is reduced to $\gamma = 30\%$, the proportion of at-risk capacity assumed to be devoted to successful vaccines is reduced to $a_m = a_o = 20\%$, VSL is reduced to $1 million, the social discount rate is raised to $r = 6\%$, or the social cost of public funds is raised to $\lambda = 30\%$. The net gain from advance investment is large in any case and thus evidently robust.

Some of the robustness exercises indicate just how conservative some model assumptions are. If we relax the truncation of pandemic intensity beyond the range of historically observed epidemics, expected net benefits from the advance program can be nearly $1.3 trillion greater than the counterfactual. Reducing the social discount rate to $r = 2\%$, a value commonly used in the literature for program evaluation, the advance program results in more than a $1 trillion gain over the counterfactual.[18]

As we saw, the baseline design of the advance program spends $60 billion up front to install 24 billion annual courses of capacity. Supposing more spending were not prevented by political constraints or steeply convex capacity costs, our model would suggest that even larger up-front investments would be worthwhile. Installing 45 billion courses of annual capacity, which one can compute entails an upfront expenditure of $113 billion, increases the NPV gain over the counterfactual program from $504 billion to $705 billion. Smaller upfront investments than the baseline design would have lower expected net benefits but would still be worthwhile. For example, installing half the advance capacity, 12 billion annual courses, entailing an upfront investment of $30 billion, would generate an NPV of $534 billion, $303 billion more than the counterfactual.

[18] As another gauge of the robustness of the results, for all parameters in Table 5, we engaged in a "stress test" that substituted increasingly extreme parameter values until a value was found making the advanced-capacity program worse than the counterfactual. No positive value of VSL does so; even with mortality losses zeroed out, the remaining output and learning losses are substantial enough that the program continues to dominate the counterfactual. For other parameters, while values can be found for which advance investment is worse than the counterfactual, extreme values are required. The probability of pandemic arrival must be cut by a factor of 20. The probability that the need for the vaccine is not obviated must be cut by a factor of 17 from the baseline $\gamma = 50\%$ to $\gamma = 3\%$. The fraction of advance capacity that can be immediately employed for successful candidates must be cut by a factor of 25 from $a_v = 30\%$ to $a_v = 1.2\%$ for both technologies $v \in \{m, o\}$. An eight-fold increase in the social discount factor is required, from $r = 4\%$ to $r = 114\%$. The social cost of public funds must be increased from $\lambda = 0\%$ to $\lambda = 1{,}440\%$.

7 International Versus National Programs

Our analysis thus far has focused on a global program. This focus allows for convenient exposition, enabling us to report a single world number rather than a set of numbers for individual countries. More importantly, a global program extends the net benefits from advance preparations, which we found to be very large, to the greatest number of individuals possible.

Because advance investment in vaccine supply by one country generates positive externalities to other countries, advance investment is likely to be suboptimal unless countries strike a cooperative agreement. Countries investing in advance reduce their demand for in-pandemic capacity, lowering the price for those who have engaged in advance investment and relaxing the constraint that at most capacity x' can be installed in the short run. Technically, if all functions in the model were linear, there would be no spillovers between countries and thus no benefit from coordination. But the concavity of the benefit function and the convexity of the short-run capacity-cost function leads countries' investments to be strategic substitutes.

An unmodeled benefit of a cooperative international agreement is that could provide insurance against variation in the severity of the pandemic across countries. The model does not allow for heterogeneity in intensity across countries, so leaves no role for such insurance. To address real-world heterogeneity across countries, both in the average intensity of the pandemic there and the timing of when the waves hit it, the international program could pool advance investment but then allocate vaccines based on the current case rate or death rate in individual countries.[19] As with other forms of insurance, such a program carries a risk of moral hazard, i.e., that countries might take less stringent control measures knowing this will increase their vaccine supply. However, given the high costs of the pandemic even with vaccine access, moral hazard is unlikely to be a significant problem in this case.[20]

Participating countries in an international program would have to agree on how program expenditures should be allocated among them. It might be natural to ask participating countries to

[19] Even after the start of the pandemic, considerable uncertainty remains about both the relative severity and timing of waves in different parts of the world suggesting insurance even during a pandemic could be beneficial. For example, India's mortality from Covid-19 was initially low, only to be badly hit by the Delta wave.

[20] The perception that allocating vaccine based on local cases rewards bad performers might still undermine efforts to include insurance-type provisions in an international program. COVAX, a large, coordinated vaccine purchase mechanism for Covid-19 initially allocated vaccines without regard to cases, mortality rates, or even demand (proxied by utilization of previous shipments been utilized), possibly indicating the political challenges of building insurance into an international program.

pay into the program according to GDP and receive vaccines according to population (or number of high-risk individuals). However, such a mechanism would involve substantial redistribution from richer to poorer countries. There may be a limit to how much redistribution can be supported before it is no longer incentive compatible for high-income countries to participate.[21] What that limit is may depend in part on the nature of the bargaining process, whether negotiators commit to abandoning the whole agreement if pivotal countries do not participate or whether the agreement forges on without those countries. Perhaps paradoxically, countries are more likely to participate if they believe themselves to be pivotal, since declining to participate destroys the program they would otherwise free ride on. Making a country pivotal relaxes its incentive-compatibility constraint and allows the program to support more redistribution from it. If high-income countries do not believe themselves to be pivotal, it will be harder to engender the optimal level of their participation even in a program without any redistribution.

In the absence of an international agreement, if no other country engages in advance investment, individual countries would have strong unilateral incentives to be the one to do so. With few countries investing in advance, competition for capacity installed during the pandemic can be expected to be intense, leading to high capacity prices and capacity shortages. An individual country has an incentive to install capacity in advance to avoid this competition.

For example, if the United States undertook an advance-investment program (proportional to the size of the global program analyzed above), we find that the program would generate an expected present value of $77 billion in benefits net of program costs, a gain of $61 billion ($183 per capita) over the counterfactual program. Not just high-income countries but middle- and lower-income countries could also benefit from unilateral advance investment. For Brazil, for example, we find that advance investment would generate an expected present value of $19 billion in net benefits over program costs, a gain of $15 billion ($70 per capita) over the counterfactual.[22]

Investing countries could generate some revenue for themselves and social value for others by signing bilateral agreements with non-investing countries to use their facilities to produce vaccines for non-investing countries while the pandemic is severe there but not domestically. The

[21] International discussions during the early stages of Covid-19 to develop a coordinated vaccine purchase arrangement across countries of very different income levels ran into some of these issues.

[22] Since our analysis of individual-country programs holds mortality losses per capita constant at the global average, the difference in benefits between the United States and Brazil is mainly driven by the greater economic-output losses suffered in the country with higher GDP per capita and are only partially offset by longer school closures observed in MICs and LICs.

investing country would retain priority over courses if the pandemic rises there but could benefit other countries meanwhile.

8 Conclusion

While the world would like to move on from worrying about pandemics, expected losses from the next pandemic are too high to ignore. Combining data on the probability distribution of epidemics of different severity and estimates of the relationship between epidemic severity (measured in deaths) and mortality, economic-output, and education losses from epidemics, we conservatively estimated that the present discounted value of the stream of expected social losses from future pandemics is $17.8 trillion. This is equivalent to losing $712 billion every year going forward, more than the budget of the U.S. Department of Defense in the baseline year. Investments that reduce the cost of the next pandemic, even if they are relatively ambitious and expensive, can generate very high expected returns.

Investing now in building the capacity to rapidly vaccinate a large percentage of the population against a new virus can be a key way to dramatically reduce the cost of future pandemics. Specifically, we calculate that $60 billion in upfront investment and $5 billion in annual expenditure would be sufficient to fund capacity to produce 24 billion vaccine courses per year, which in turn would be sufficient to vaccinate 70% of the world's population in six months in a pandemic. The NPV gain from this program over the status quo in which capacity investment is delayed until the start of the next pandemic would be over $500 billion. Under reasonable assumptions, even larger capacity would have a high social return. Social returns are high even if we factor in a risk that vaccines will not work against the next virus, that there will be a high degree of vaccine hesitancy, that an effective antiviral will reduce the benefit of a vaccine, or that new vaccine technology, superior to mRNA, renders a greater fraction of capacity sunk now obsolete later. Social returns remain high even if the funder faces a substantial social cost of public funds owing to distortionary taxation or opportunity cost of diverting resources from other worthy projects.

In contrast to the recommendations in this paper, valuable mRNA vaccine capacity is currently being, or about to be, converted to other uses. Allowing this capacity to be dismantled suggests we are failing to learn the lesson of Covid-19.

Our proposal is broadly consistent with recommendations in high-profile reports on

pandemic preparedness including those issued by the G7, G20, and McKinsey. The G7 set out a road map for accelerating the development and distribution of vaccines in a future pandemic that included greatly expanding vaccine manufacturing capacity and ensuring it is "kept warm" (G7 Pandemic Preparedness Partnership 2021). Different elements of the international system are tasked with implementing the recommendations. The G20 make similar recommendations (G20 High Level Independent Panel on Financing the Global Commons for Pandemic Preparedness 2021). A McKinsey report (Craven et al. 2021) proposes a five-prong approach to improving global pandemic preparedness, one of which calls for advance investment in vaccine capacity, the focus of our paper. Our contribution beyond these reports is to estimate benefits and costs based on a formal model of the program's design.

While this paper has focused on the benefits of maintaining vaccine supply capacity, the expected large losses from future pandemics imply that other investments that reduce their costs are also likely to be good investments. Perhaps most closely related to this paper are the proposals to develop and produce in large quantities a universal coronavirus vaccine that would protect against a wide variety of coronaviruses including ones that are yet to emerge. A new World Bank fund (Financial Intermediary Fund for Pandemic Prevention) plans to invest in, among other things, increased surveillance for the emergence of new pathogens. Research and development into new mRNA vaccines will help us improve this still relatively new technology and understand which type of viruses it is best suited to combatting. Berry et al. (2020) have suggested that putting in place a framework for when human challenge trials could accelerate testing of new vaccines and save millions of lives in a future pandemic. Investments in health systems to store and transport vaccines and the health staff to administer them would be helpful to reduce the lag in rolling out available supplies. Finally, while much of the analysis and recommendations in this paper are specific to reducing the cost of future viral pandemics, multidrug resistant bacteria remain a threat. Development of new antibiotics to be kept in reserve for use only in combination therapy for multidrug-resistant strains would reduce the probability of a highly damaging bacterial pandemic. Fortunately, scaling up antibiotics tends to be easier, cheaper, and faster than scaling up vaccines, hence our focus on vaccines. A similar logic applies to the development of broad-spectrum antivirals. Given the large expected losses from future pandemics, investments in the development of broad-spectrum antivirals are likely to be highly cost effective even if there were no certainty they would work against the next pandemic.

Appendix A. Details on Derivations

This appendix provides details on several derivations omitted from the text for space considerations.

Derivation of Equation (11)

A preliminary step is to prove that the analogous formula for $ML_t(i)$ in equation (5) holds for $OL_t(i)$ and $LL_t(i)$. We have

$$\begin{aligned} OL_t(i) &= Y_t N_t \frac{\Delta(i)}{100} \\ &= Y_t N_0 \frac{\Delta(i)}{100} \\ &= \frac{Y_t}{Y_0} OL_0(i) \\ &= (1+y)^t OL_0(i). \end{aligned} \tag{23}$$

The first step is a generalization of the formula given for $OL_0(i)$ in (7). The second step follows from fixing the population at N_0 for our analysis. The third step follows from substituting from (7). The last step follows from the definition of y as the GDP growth rate, implying $Y_t = (1+y)^t Y_0$. We also have

$$\begin{aligned} LL_t(i) &= \frac{10}{13.8} OL_t(i) \\ &= \frac{10}{13.8} (1+y)^t OL_0(i) \\ &= (1+y)^t LL_0(i). \end{aligned} \tag{24}$$

The first step follows from (8), the second step follows from (23), and the third step follows again from (8).

Combining these formulas,

$$\begin{aligned} TL_t(i) &= ML_t(i) + OL_t(i) + LL_t(i) \\ &= (1+y)^t [ML_0(i) + OL_0(i) + LL_0(i)] \\ &= (1+y)^t TL_0(i). \end{aligned} \tag{25}$$

Substituting (25) into (10),

$$
\begin{aligned}
E\left(TL_t(i)\right) &= \int_0^{\infty} TL_t(i)\, d\Phi(i) \\
&= (1+y)^t \int_0^{\infty} TL_0(i)\, d\Phi(i) \\
&= (1+y)^t E\left(TL_0(i)\right).
\end{aligned} \tag{26}
$$

Substituting (26) into (9), rearranging, and summing the resulting series,

$$
\begin{aligned}
PV(\overline{TL}) &= \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho}\right)^{t+1} E\left(TL_t(i)\right) \\
&= E\left(TL_0(i)\right) \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^{t+1} (1+y)^t \\
&= \frac{E\left(TL_0(i)\right)}{1+y} \sum_{t=0}^{\infty} \left(\frac{1+y}{1+r}\right)^{t+1} \\
&= \frac{E\left(TL_0(i)\right)}{r-y}.
\end{aligned} \tag{27}
$$

Multiplying both sides by ρ yields equation (11).

Derivation of Equation (20)

Equation (20) is the sum of the two sources of costs given in equations (16) and (19). Simplifying the series in (16),

$$
\begin{aligned}
&(1-\phi)\left[k_v z_v + \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^{t+1} \Phi(i^*)^t \delta k_v z_v\right] \\
&= (1-\phi)k_v z_v + \frac{(1-\phi)\delta k_v z_v}{\Phi(i^*)} \sum_{t=0}^{\infty} \left(\frac{\Phi(i^*)}{1+r}\right)^{t+1} \\
&= (1-\phi)k_v z_v + \frac{(1-\phi)\delta k_v z_v}{1+r-\Phi(i^*)}.
\end{aligned} \tag{28}
$$

Simplifying the series in (19),

$$\sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^{t+1} h_t(i^*)\left[K_v(x_v - z_v) + c_v \cdot \frac{x_v}{x} \cdot 0.7N_0\right]$$

$$= \left[K_v(x_v - z_v) + c_v \cdot \frac{x_v}{x} \cdot 0.7N_0\right]\sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^{t+1} \Phi(i^*)^t[1 - \Phi(i^*)] \tag{29}$$

$$= \left[K_v(x_v - z_v) + c_v \cdot \frac{x_v}{x} \cdot 0.7N_0\right]\frac{1 - \Phi(i^*)}{\Phi(i^*)}\sum_{t=0}^{\infty}\left(\frac{\Phi(i^*)}{1+r}\right)^{t+1}$$

$$= \left[K_v(x_v - z_v) + c_v \cdot \frac{x_v}{x} \cdot 0.7N_0\right]\frac{1 - \Phi(i^*)}{1 + r - \Phi(i^*)}.$$

Summing (28) and (29) yields (20).

References

Ahuja, Amrita, Susan Athey, Arthur Baker, Eric Budish, Juan C. Castillo, Rachel Glennerster, Scott D. Kominers, Michael Kremer, Jean Lee, Canice Prendergast, et al. 2021. Preparing for a Pandemic: Accelerating Vaccine Availability. *American Economic Association Papers and Proceedings* 111: 331–335.

Alfani, Guido. 2022. Epidemics, Inequality, and Poverty in Preindustrial and Early Industrial Times. *Journal of Economic Literature* 60(1): 3–40.

Athey, Susan, Juan C. Castillo, Esha Chaudhuri, Michael Kremer, Alexandre Simoes Gomes, and Christopher M. Snyder. 2022. Expanding Capacity for Vaccines Against Covid-19 and Future Pandemics: A Review of Economic Issues. *Oxford Review of Economic Policy* 38(4): 742–770.

Azevedo, Joao P. Amer Hasan, Diana Goldemberg, Koen Geven, and Syedah A. Iqbal. 2021. Simulating the Potential Impacts of COVID-19 School Closures on Schooling and Learning Outcomes: A Set of Global Estimates. *World Bank Research Observer* 36(1): 1–40.

Barro, Robert J., José F. Ursúa, and Joanna Weng. 2020. The Coronavirus and the Great Influenza Pandemic: Lessons from the "Spanish Flu" for the Coronavirus's Potential Effects on Mortality and Economic Activity. National Bureau of Economic Research working paper no. 26866.

Becker, Gary S. 2007. Health as Human Capital: Synthesis and Extensions. *Oxford Economic Papers* 59(3): 379–410.

Bernstein, Aaron S., Amy W. Ando, Ted Lock-Temzelides, Mariana M. Vale, Binbin V. Li, Hongying Li, Jonah Busch, Colin A. Chapman, Margaret Kinnaird, Katarzyna Nowak, et al. The Costs and Benefits of Primary Prevention of Zoonotic Pandemics." *Science Advances* 8(5): eab14183.

Berry, Donald A., Scott Berry, Peter Hale, Leah Isakov, Andrew W. Lo, Kien W. Siah, and Chi H. Wong. 2020. A Cost/Benefit Analysis of Clinical Trial Designs for COVID-19 Vaccine Candidates. *PLoS One* 15(12): e0244418.

Burnett, F. M. 1979. Portraits of Viruses: Influenza Virus A. *Intervirology* 11(4): 201–214.

Carlson, Colin J., Maxwell J. Farrell, Zoe Grange, Barbara A. Han, Nardus Mollentze, Alexandra L. Phelan, Angela L. Rasmussen, Gregory F. Albery, Bernard Bett, David M. Brett-Major, et al. 2021. The Future of Zoonotic Risk Prediction. *Philosophical Transactions of the Royal Society B* 376(1837): 20200358.

Castillo, Juan C., Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Tasneem Chipty, Rachel Glennerster, Scott D. Kominers, Michael Kremer, Greg Larson, et al. 2021. Market Design to Accelerate COVID-19 Vaccine Supply. *Science* 371(6534): 1107–1109.

Centers for Disease Control and Prevention (CDC). 2019. 2014–2016 Ebola Outbreak in West Africa. [Internet.] Downloaded June 9, 2023 from https://www.cdc.gov/vhf/ebola/history/2014-2016-outbreak/index.html.

Chandra, Siddharth, and Julia Christensen. 2019. Re: "Reassessing the Global Mortality Burden of the 1918 Influenza Pandemic. *American Journal of Epidemiology* 188(7): 1404.

Crank, Michelle C., Tracy J. Ruckwardt, Man Chen, Kaitlyn M. Morabito, Emily Phung, Pamela J. Costner, Lasonji A. Holman, Somia P. Mickman, Nina M. Berkowitz, Ingelise J. Gordon, et al. 2019. A Proof of Concept for Structure-Based Vaccine Design Targeting RSV in Humans. *Science* 365(6452): 505–509.

Craven, Matt, Adam Sabow, Lieven Van der Veken, and Matt Wilson. 2021. Not the Last Pandemic: Investing Now to Reimagine Public-health Systems. McKinsey & Company report. [Internet.] Downloaded May 10, 2023 from https://www.mckinsey.com/industries/public-and-social-sector/our-insights/not-the-last-pandemic-investing-now-to-reimagine-public-health-systems.

Duflo, Ester. 2001. Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *American Economic Review* 91(4): 795–813.

Economist. 2022. Excess Deaths Tracker. [Internet.] Downloaded May 8, 2022 from https://www.economist.com/graphic-detail/coronavirus-excess-deaths-tracker.

Fan, Victoria Y., Dean T. Jamison, and Lawrence H. Summers. 2018. Pandemic Risk: How Large Are the Expected Losses? *Bulletin of the World Health Organization* 96(2); 129–134.

G7 Pandemic Preparedness Partnership. 2021. 100 Days Mission to Respond to Future Pandemic Threats. June 12. [Internet.] Downloaded May 18, 2023 from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/992762/100_Days_Mission_to_respond_to_future_pandemic_threats__3_.pdf.

G20 High Level Independent Panel. 2021. A Global Deal for our Pandemic Age: Report of the G20 High Level Independent Panel on Financing the Global Commons for Pandemic Preparedness and Response. [Internet.] Downloaded May 3, 2023 from https://pandemic-financing.org/report/foreword.

Gollier, Christian and James K. Hammitt. 2014. The Long-run Discount Rate Controversy. *Annual Review of Resource Economics* 6: 273–295.

Huber, Caroline, Lyn Finelli, and Warren Stevens. 2018. The Economic and Social Burden of the 2014 Ebola Outbreak in West Africa. *Journal of Infectious Diseases* 218(issue supplement 5): S698–S704.

Keogh-Brown, Marcus R. and Richard D. Smith. 2008. The Economic Impact of SARS: How Does the Reality Match the Predictions? *Health Policy* 88(1): 110–120.

Keogh-Brown, Marcus R., Simon Wren-Lewis, W. John Edmunds, Philippe Beutels, and Richard D. Smith. 2009. The Possible Macroeconomic Impact on the UK of an Influenza Pandemic. *Health Economics* 19(11): 1345–1360.

International Monetary Fund. 2022. World Economic Outlook Update: Rising Caseloads, a Disrupted Recovery, and Higher Inflation. January. [Internet.] Downloaded June 8, 2022 from https://www.imf.org/en/Publications/WEO/Issues/2022/01/25/world-economic-outlook-update-january-2022

Kay, Chris, Malavika K. Makol, and James Paton. 2022. World Moves from Shortages to Possible Glut of Covid-19 Vaccines. *Bloomberg*. March 28. [Internet.] Downloaded September 7, 2022 from https://www.bloomberg.com/news/articles/2022-03-28/world-moves-from-shortages-to-possible-glut-of-covid-19-vaccines#xj4y7vzkg.

Lo, Andrew W., Kien W. Siah, and Chi H. Wong. 2020. Estimating Probabilities of Success of Vaccine and Other Anti-Infective Therapeutic Development Programs. *Harvard Data Science Review* (Special Issue 1).

Marseille, Elliot, Bruce Larson, Dhruv S. Kazi, James G. Kahn, and Sydney Rosen. 2015. Thresholds for the Cost-effectiveness of Interventions: Alternative Approaches. *Bulletin of the World Health Organization* 93(2): 118–124.

Marani, Marco, Gabriel G. Katul, William K. Pan, and Anthony J. Parolari. 2021. Intensity and Frequency of Extreme Novel Epidemics. *Proceedings of the National Academy of Sciences* 118(35): e2105482118.

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 no. 23346.

Meltzer Martin I., Nancy J. Cox, and Keiki Fukuda. 1999. The Economic Impact of Pandemic Influenza in the United States: Priorities for Intervention. *Emerging Infectious Disease* 5(5): 659–671.

Mincer, Jacob. 1974. *Schooling, Experience and Earnings*. Columbia University Press.

Mirasol, Feliza. 2021. Meeting Fill/Finish Challenges for COVID-19 Vaccines. *BioPharm* 34(2): 40–43.

Mole, Beth. 2021. Before Ruining Millions of Vaccines, Emergent Failed Inspections, Raked in Cash. Arstechnica May 19. [Internet.] Downloaded September 9, 2022 from https://arstechnica.com/science/2021/05/emergent-got-27m-a-month-to-prep-vaccine-plant-then-ruined-15m-jj-doses/.

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.

OECD. 2022. The Long Game: Fiscal Outlooks to 2060 Underline Need for Structural Reform. [Internet.] Downloaded September 5, 2022 from https://www.oecd.org/economy/growth/scenarios-for-the-world-economy-to-2060.htm.

Patterson, K. David, and Gerald F. Pyle. 1991. The Geography and Mortality of the 1918 Influenza Pandemic. *Bulletin of the History of Medicine* 65(1): 4–21.

Prager, Fynnwin, Dan Wei, and Adam Rose. 2016. Total Economic Consequences of an Influenza Outbreak in the United States. *Risk Analysis* 37(1): 4–19.

Sanford, J.P. 1969. Influenza: Consideration of Pandemics. *Advances in Internal Medicine* 15: 419–453.

Schoenbaum, Stephen C. 1987. Economic Impact of Influenza. *American Journal of Medicine* 82(6): 26–30.

Sell, Tara K., Daniel Gastfriend, Matthew Watson, Crystal Watson, Lauren Richardsson, Anita Cicero, Tom Inglesby, and Nancy Connell. 2021. Building the Global Vaccine Manufacturing Capacity Needed to Respond to Pandemics. *Vaccine* 39(12): 1667–1669.

Silver, Kate. 2021. Shot of a Lifetime: How Pfizer is Partnering with CMOs to Increase COVID-19 Vaccine Production and Reach More People. Pfizer October 7 press release. [Internet.] Downloaded May 11, 2023 from https://www.pfizer.com/news/articles/shot_of_a_lifetime_how_pfizer_is_partnering_with_cmos_to_increase_covid_19_vaccine_production_and_reach_more_people

Spreeuwenberg, Peter, Madelon Kroneman, and John Paget. 2018. Reassessing the Global Mortality Burden of the 1918 Influenza Pandemic. *American Journal of Epidemiology* 187(12): 2561–2567.

Snow, Arthur, and Ronald Warren, Jr. 1995 The Marginal Welfare Cost of Public Funds: Theory and Estimates. *Journal of Public Economics* 61(2): 289–305.

Spence, Michael. 1973. Job Market Signaling. *Quarterly Journal of Economics* 87(3): 355–374.

Sweis, Nadia J. 2022. Revisiting the Value of a Statistical Life: An International Approach Duiring COVID-19. *Risk Management* 24: 259–272.

United Nations Development Programme. 2017. A Socio-economic Impact Assessment of the Zika Virus in Latin America and the Caribbean. [Internet.] Downloaded June 8, 2022 from https://global.undp.acsitefactory.com/sites/g/files/zskgke326/files/publications/UNDP-RBLAC-Zika-English-WEB.pdf.

U.S. Department of Defense. 2020. DOD Releases Fiscal Year 2021 Budget Proposal. [Internet.] Downloaded June 10, 2023 from
https://www.defense.gov/News/Releases/Release/Article/2079489/dod-releases-fiscal-year-2021-budget-proposal/.

Viscusi, W. Kip. 2020. Pricing the Global Health Risks of the COVID-19 Pandemic. *Journal of Risk and Uncertainty* 61 (2): 101–128.

Yamey, Gavin, Marco Schäferhoff, Ole K. Aars, Barry Bloom, Dennis Carroll, Mukesh Chawla, Victor Dzau, Ricardo Echalar, Indermit S Gill, Tore Godal, et al. 2017. Financing of International Collective Action for Epidemic and Pandemic Preparedness. *Lancet Global Health* 5(8): E742–E744.

Table 1 Mortality and Economic Losses from Selected Pandemics Over Last Century

	Deaths		Economic losses	
Epidemic (date)	Total deaths over pandemic (source)	Mortality intensity = annual deaths per thousand people in world	Economic loss over pandemic	Annual economic loss (% global GDP)
1918 flu (1918–20)	32.0 mil. (Patterson & Pyle 1991)	5.7×10^{1}	6% global GDP (Barro, Ursúa, & Weng 2020)	2.0
SARS (2002–03)	1,000 (Marani et al. 2021)	1.2×10^{-4}	0.1% global GDP (Keogh-Brown & Smith 2008)	2.5×10^{-2}
Ebola (2013–16)	11,325 (CDC 2019)	3.9×10^{-4}	0.06% global GDP (Huber, Finelli, & Stevens 2018)	8.3×10^{-2}
Zika (2015–17)	1,000 (Marani et al. 2021)	4.5×10^{-5}	0.05% Latin American and Caribbean GDP annually (UN Development Programme 2017)	1.7×10^{-2}
Covid-19 (2020-22)	21.3 mil. (*Economist* 2022)	3.2×10^{-1}	14.4% global GDP (IMF 2022)	3.6

Notes: For SARS and Zika, estimated deaths set to 1,000, the lower bound on observation threshold from Marani et al. (2021) power-law distribution.

Table 2 Expected Annual Global Pandemic Losses

	Expected annual pandemic losses (billion dollars)			
Scenario	Mortality $AV(\overline{ML})$	Economic output $AV(\overline{OL})$	Learning $AV(\overline{LL})$	Total $AV(\overline{TL})$
Baseline	519	112	81	712
Probability of pandemic arrival				
• Halved	259	56	41	356
• Doubled	849	183	133	1,165
Truncating intensity distribution				
• Halve upper truncation	339	90	65	494
• Double upper truncation	640	114	83	836
• Remove upper truncation	1,283	118	86	1,486
Value of statistical life (VSL)				
• Reduce to $1 million	399	112	81	592
• Increase to $1.6 million	638	112	81	832
GDP per capita growth rate y				
• Reduce to 1.4%	479	104	75	657
• Increase to 1.8%	566	122	89	777
Social discount rate r				
• Reduce to 2%	1,556	336	244	2,136
• Increase to 6%	424	92	66	582

Notes: Entries are expected annual global pandemic losses in each category and in total following equation (11) in billions of 2021 dollars rounded to the nearest billion. The baseline scenario in the first row of results assumes a distribution of pandemic intensity given by equation (3); truncates pandemic intensity at the highest figure for the 1918 flu reported in a scholarly source; uses the arrival rate calculated using the Marani et al. (2021) data; and sets the VSL at $1.3 million, social discount factor at $r = 4\%$, and GDP per capita growth rate at $y = 1.6\%$. Other scenarios change only the indicated feature from the baseline.

Table 3 Model Parameters

Notation	Definition	Baseline value	Range for sensitivity analysis
Pandemic-arrival parameters from Marani et al. (2021)			
α	Probability detectable epidemic	0.62	n.a.
σ	Pareto shape parameter	0.0113	n.a.
ξ	Pareto shape parameter	1.41	n.a.
μ'	Lower threshold for detectable epidemic	10^{-3}	n.a.
Other pandemic-arrival parameters			
μ''	Upper threshold for largest epidemic (authors)	17.8	8.9–35.6
Pandemic-loss parameters from various sources			
VSL	Value of statistical life (Sweis 2022)	$1.3 million	$1–1.6 mil.
y	GDP per capita growth rate (OECD 2022)	1.6%	1.4–1.8%
δ	Depreciation rate (authors)	8%	6–10%
r	Social discount rate (authors)	4%	2–6%
Vaccine-supply parameters from Castillo et al. (2021)			
γ	Fraction remaining harm mitigated by vaccine	50%	30–70%
θ	Reduction of pandemic-time investments	25%	0–50%
k_m	Unit cost of advance mRNA capacity	$1.50 per ann. course	$0.75–3.00
k_o	Unit cost of advance capacity for other vaccines	$3.00 per ann. course	$1.50–6.00
c_m	Marginal cost of producing mRNA vaccines	$6.00 per course	$3.00–12.00
c_o	Marginal cost of producing other vaccines	$3.00 per course	$1.50–6.00
ε	Decreasing returns to pandemic capacity	1	0.75–1.25
p_b	Probability both technologies successful	50%	30–70%
p_m, p_o	Probability technology alone successful	15%	10–20%
Vaccine-supply parameters from G20 High Level Independent Panel (2021)			
a_m, a_o	Fraction of at-risk capacity successful	30%	20–40%
b_m, b_o	Fraction of at-risk capacity repurposable	40%	20–60%
τ_m	Time to repurpose mRNA candidate	2 months	1–3 months
τ_o	Time to repurpose other candidate	6 months	3–9 months
ϕ	Rental recovery fraction	70%	50–90%
Other cost parameters			
λ	Social cost of public funds	0%	0-30%

Notes: For parameters for which the source is listed as "authors," Section 3 provides details on the calculations or judgment we used to select the baseline value of the parameter, often informed by other scholarly sources cited in the text. Sensitivity analysis only conducted for last of the pandemic-arrival parameters; rest are marked "not applicable." The 30% upper bound on parameter λ is based on an estimate of the deadweight loss of taxation in developed countries by Snow and Warren (1995).

Table 4 Baseline Results for Program Outcomes

	Costs and benefits of program to undertake vaccination campaign in next significant pandemic (billion $)		
	With advance investment	Without advance investment	Difference
Current value of expenditures			
• Initial advance investment	60	0	60
• Annual maintenance of advance capacity	5	0	5
• Additional expenditures in pandemic	22	53	-32
Present value of program outcomes			
• Expected program costs (net of rental income)	50	15	35
• Expected gross benefits	785	246	539
• Expected net benefits	735	231	504

Notes: All entries are in billions of 2021 dollars. First set of rows report current value of expenditures in year undertaken. These are actual, not effective, expenditures, so do not net out rental income. Second set of rows report present values (from the perspective of the base year) of costs and benefits of program leading to vaccination campaign in next pandemic of at least half the intensity of Covid-19.

Table 5 Sensitivity of Program Outcomes to Changes in Selected Parameters

Scenario	Present value of expected net benefits of vaccination campaign in next significant pandemic (billion $)		
	With advance investment	Without advance investment	Difference
Baseline	735	231	504
Probability of pandemic arrival	437	144	293
• Halved	1,079	329	749
• Doubled			
Truncating intensity distribution	586	184	402
• Halve upper truncation	910	286	624
• Double upper truncation	1,828	573	1,255
• Remove upper truncation			
γ (fraction of remaining harm mitigated by vaccine)	421	133	288
• Reduce to 30%	1049	329	719
• Increase to 70%			
Value of statistical life (VSL)	582	183	399
• Reduce to $1 million	887	279	609
• Increase to $1.6 million			
a_m, a_o (fraction of at-risk capacity successful)	596	171	425
• Reduce to 20%	848	284	563
• Increase to 40%			
r (social discount rate)	1,518	473	1,045
• Reduced to 2%	480	152	328
• Increased to 6%			
λ (social cost of public funds)	720	227	494
• Increased to 30%	735	231	504

Notes: All entries are expected present values in billions of 2021 dollars. For space considerations, sensitivity analyses have been omitted for other parameters that had smaller effect on difference in program outcomes than those displayed, for example, for f_m and f_o. Those parameters include ε, y, δ, ϕ, and the pairs (b_m, b_o), (k_m, k_o), and (τ_m, τ_o), where the two elements of each pair are changed together. See Table 3 for the baseline values of those parameters and changes considered.

Figure 1. Exceedance Probability Based on Marani et al. (2021) Estimates.

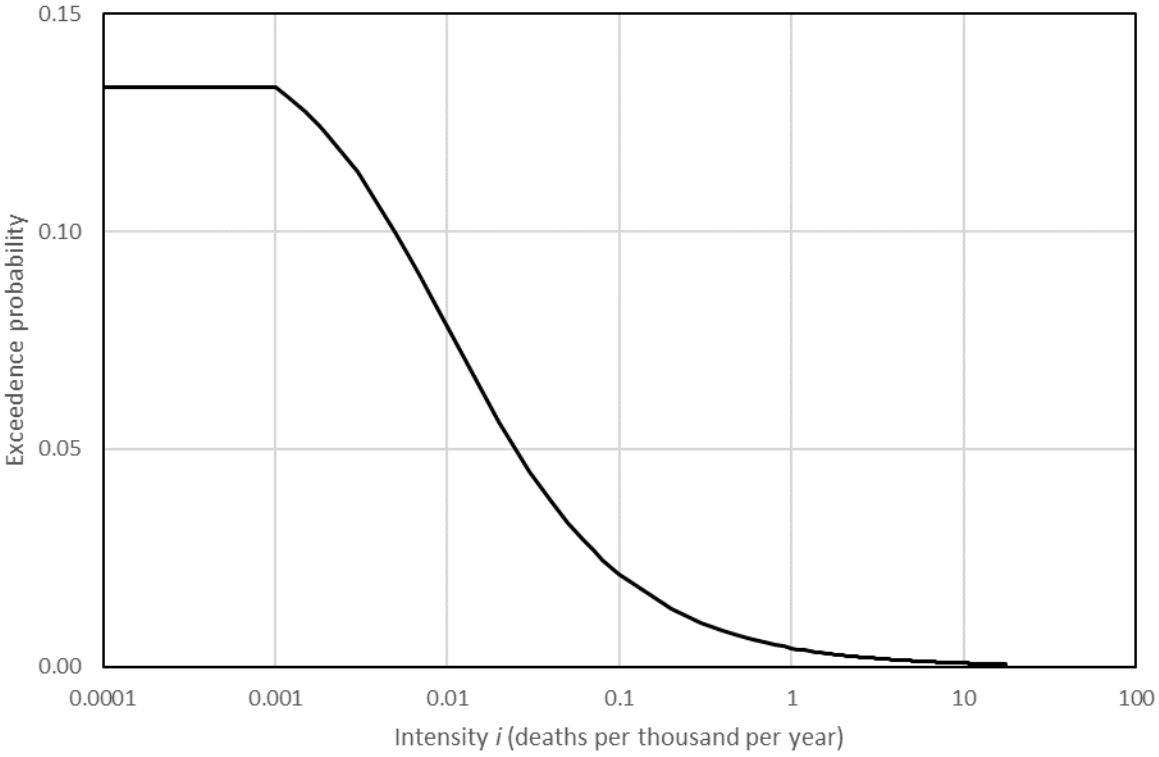

Notes: Graph of the exceedance probability, interpreted as the probability that an epidemic of at least that intensity arrives in a year. Formally, the exceedance probability equals $1 - \Phi(i)$, where $\Phi(i)$ is the cumulative distribution function in equation (3). Log scale used for horizontal axis. Above an intensity of $\mu'' = 17.5$, probability drops to zero in the baseline specification, reflecting the conservative baseline assumption truncating the intensity of pandemics at the highest report in a scholarly source for the most intense pandemic in the Marani et al. (2021) data (1918 Flu).

Figure 2. Relationship Between Epidemic Intensity Economic Losses in Historical Pandemics.

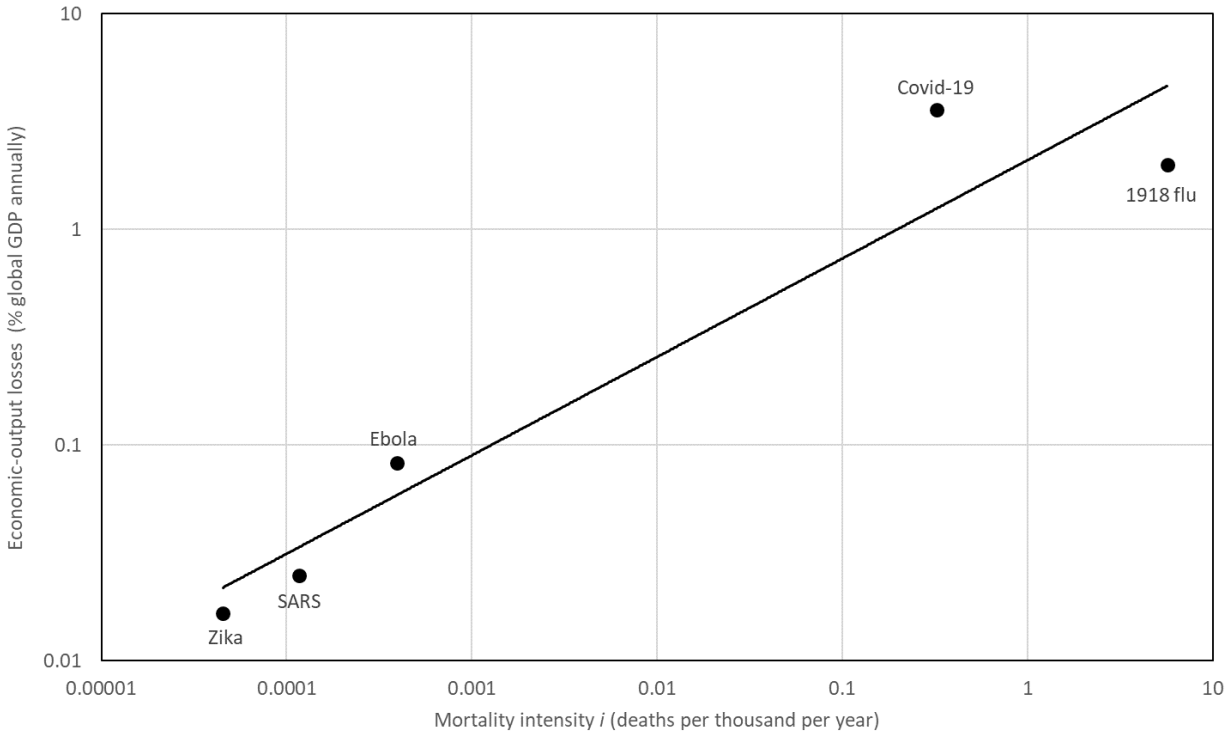

Notes: Data points from Table 1. Regression line given by equation (6). Log scales used for both axes.