Estimating future pandemic harms and gains from preparedness investments

Christopher M Snyder, Catherine Che, Sarrin Chethik, Rachel Glennerster, Ganqi Li, and Sebastian Quaade

August 2026

Department of Economics, Dartmouth College, Hanover, NH, USA (Prof. C M Snyder PhD); Division of Social Science, University of Chicago, Chicago, IL, USA (Prof. R Glennerster PhD); Department of Economics, University of California, Berkeley, Berkeley, CA, USA (C Che PhD); Market Shaping Accelerator, Washington, DC, USA (C M Snyder, S Chethik BA, R Glennerster); Institute for Progress (C M Snyder, S Quaade MSc); Tuck School of Business, Dartmouth College, Hanover, NH, USA (G Li BA); Center for Global Development, Washington, DC, USA (C M Snyder, R Glennerster)

Correspondence to: Christopher Snyder, Department of Economics, Dartmouth College, Hanover NH 03755, USA chris.snyder@dartmouth.edu

Summary

Background **Evidence on the expected benefits and costs of many pandemic preparedness measures remains limited. This study estimates expected future losses from pandemics and the expected value of a suite of preparedness investments that enhance vaccine response.**

Methods **We forecast pandemic risk using a Poisson-generalized Pareto distribution fit to historical outbreak data. We convert mortality, economic, and learning losses to monetary terms, and use Monte Carlo simulation to estimate expected losses absent vaccine response. We then introduce a vaccine response model calibrated using an expert survey and precedent from COVID-19 and estimate the net present value (NPV) of status quo vaccine response and the incremental NPV of four preparedness investments: advance manufacturing capacity, prototype vaccine R&D, universal flu vaccine R&D, and improved early warning systems.**

Findings **We estimate $4 trillion in expected annual social losses from pandemics. Status quo vaccine response on average offsets 18% of that harm ($17.6 trillion NPV). Among preparedness investments, universal influenza vaccine R&D delivers the largest incremental NPV ($2.4 trillion), followed by improved early warning ($436 billion), advance capacity ($269 billion), and prototype vaccine R&D ($256 billion). The combination of the four investments has an NPV of $3.3 trillion.**

Interpretation **All four preparedness programs generated benefits exceeding their cost. Universal flu vaccine R&D in particular stands out. The high expected value of status quo vaccine response highlights the importance of preserving the capabilities that enabled rapid COVID-19 vaccine rollout.**

Funding **Coalition of Epidemic Preparedness Innovations (CEPI); Coefficient Giving.**

Research in context

Evidence before this study

We identified academic and grey literature on pandemic risk forecasting and benefit-cost analysis of pandemic preparedness investments via Google Scholar searches and reference-list snowballing from identified papers. This process uncovered two major clusters of work. The first were studies forecasting the frequency and severity of future pandemics without detailed evaluation of alternative preparedness scenarios. The second were studies evaluating the benefits and/or costs of various pandemic preparedness interventions conditional on a pandemic having emerged. We found no study integrating a forecast of future pandemic risk with an evaluation of multiple preparedness investments such that their relative cost-effectiveness and potential complementarities could be examined.

Added value of this study

This study integrates a forecast of future pandemic risk with a detailed model of vaccine response, allowing us to separate the benefit of status quo vaccine response from the incremental benefit of preparedness investments that build on it, and to quantify net complementarities among preparedness investments. We find that all four preparedness investments we model have positive net present value, with universal flu vaccine R&D providing the largest net benefit at $2.4 trillion. Net complementarities between the considered preparedness investments are small relative to their standalone contributions. We also find that the value of status quo vaccine response is more than five times greater than that of our combined preparedness investments.

Implications of all the available evidence

Our pandemic risk forecast corroborates other studies that find large expected losses from future pandemics. Our findings suggest that the modeled interventions are strong investments, but that there are notable differences in their benefits and costs. Universal flu vaccine R&D in particular stands out for its high net present value and benefit-cost ratio. Furthermore, the large returns to status quo vaccine response underscore the importance of maintaining and fully utilizing current vaccine response capabilities.

Introduction

How should the world prepare for the possibility of pandemics as or more devastating than COVID-19? History documents how disruptive and lethal pandemics can be. The 1918 flu, for example, is estimated to have killed between 1% and 5.4% of the global population.[1] However, severe pandemics have been relatively rare: since 1800, only two outbreaks have caused more deaths than COVID-19.

The devastating impact of COVID-19 renewed calls for improved pandemic preparedness,[2] yet credible estimates of the benefits and costs of such measures remain limited. A central challenge is to incorporate uncertainty about future pandemic risk. Many studies sidestep this challenge by taking the pandemic's arrival as given when evaluating pandemic policies.[3–5] However, estimating the *expected* benefits of preparedness investments requires pandemic forecasting. Virtually all articles providing pandemic forecasts focus on the forecasting exercise rather than detailed policy evaluation.[6–10] Previous work by our coauthor team in Glennerster, et al. is an exception, where we estimated the expected benefits and costs to advance manufacturing capacity for vaccine response.[11] Left unanswered is whether that investment is cost-effective compared to other interventions and whether the combination of different investments has positive or negative synergies. We are also unaware of any study answering the basic question of what the social return from the status quo vaccine response is.

This paper estimates the expected social losses from future viral pandemics and evaluates the costs and benefits of four preparedness investments that enhance vaccine response: advance vaccine manufacturing capacity, research and development (R&D) for prototype vaccines targeting neglected pathogens, R&D for a universal influenza vaccine, and an improved early-warning system. To do so, we integrate a forecast of future pandemic risk with a stylized model of vaccine response, which we calibrate using an original expert survey and patterns observed during the COVID-19 pandemic.

Methods

Overview

The sections below describe the four methodological components that comprise our work: a forecast of future pandemic risk, a procedure for estimating pandemic mortality, economic, and learning losses, a stylized model of vaccine response, and the definition of four preparedness programs that enhance that response. To derive several inputs into these components, we also conducted an original survey of vaccine, manufacturing, biosurveillance, and epidemiology experts that we describe in the appendix. Further methodological detail can also be found in that material.

Pandemic-risk forecast

We model pandemic risk using techniques from extreme-value statistics commonly used to model rare, high-consequence events across many domains.[12] Specifically, we model outbreak severity above a given threshold as following a generalized Pareto distribution, which is the limiting distribution for exceedances over high thresholds. We treat outbreaks exceeding that threshold as arriving according to a Poisson process, which implies the distribution of annual maximum severities follows a generalized extreme value (GEV) distribution, the formula for which is provided in equation (A3) in the appendix.

To estimate this distribution, we apply maximum likelihood estimation (MLE) to data on the frequency and severity of historical outbreaks from Marani, et al.[6] We curate this dataset in several ways. Since our approach overlays a model of vaccine response on estimates of pandemic severity in the "state of nature" in the absence of such response, we need to replace the observed severity for pandemics in the dataset with estimates of what the severity would have been in the absence of vaccines (requiring an inversion of our vaccine-response model, described in the appendix). We also limit the sample in our primary specification to novel viral outbreaks (since these are of most

relevance for our vaccine-response model) and restrict attention to those above the data source's threshold annual mortality to be not missing at random. To avoid extrapolating our forecast far beyond the support of our sample, we truncate the fitted distribution at a maximum severity of 200 deaths per 10,000 population—approximately the severity of the 1918 influenza pandemic.

We model pandemic durations as following a lognormal distribution fit to the same sample. We further estimate pathogen-specific arrival by surveying experts in public health and pandemic preparedness on the expected share of future pandemics attributable to each of nine viral families prioritized by the Coalition for Epidemic Preparedness Innovations (CEPI) and the World Health Organization (WHO), along with catch-all categories for other known viruses and as-yet-unknown pathogens.

Social-loss estimation

Pandemics cause harm beyond their direct mortality. We therefore model mortality, economic, and educational losses from pandemics and combine them into a single monetized measure of social harm. All costs and benefits in the paper are expressed in 2024 U.S. dollars.

Mortality losses Pandemic mortality in any given year is calculated from the outbreak's severity and duration, as well as the size of the global population. We hold the global population fixed at its 2024 level of 8.14 billion throughout our simulations given uncertainty about long-run populations trends.[13] We monetize mortality losses using an estimate of the global value of statistical life (VSL) of $1.3 million.[14] The VSL reflects individuals' revealed willingness to pay to reduce their risk of death. The VSL is scaled by a projected long-run growth rate to reflect that willingness to pay for mortality risk reduction increases as incomes grow.

Output losses Economic disruption caused by a pandemic naturally rises with its lethality, as higher deaths, morbidity, and risk-mitigating measures obstruct more economic activity. We model

the share of global GDP lost during a pandemic as a power-law function of its severity, and estimate the parameters of this function using Poisson regression on a subsample of historical pandemics with documented economic losses (results reported in appendix Table A3).

Learning losses We calibrate learning losses starting with the assumption they are proportional to output losses. We estimate the ratio of learning to economic output losses using observations from COVID-19, dividing the World Bank's estimated $10 trillion in lifetime learning losses caused by the pandemic[15] by the IMF's estimate of $13.8 trillion in economic losses from COVID-19.[16]

Total social losses Social losses are the sum of mortality, output, and learning losses. We discount future losses 4% annually starting from the first simulation year. To express social losses in terms comparable to annual GDP figures, we convert the discounted present value of losses over the simulation horizon into an equivalent constant annual loss.

Model of status quo vaccine response

We develop a stylized model of vaccine response incorporating outbreak detection, response vaccine R&D and manufacturing capacity investments, vaccine production, and harm reduction from vaccine uptake (Figure 1). We calibrate the status quo vaccine program using our expert survey together with precedent from the COVID-19 vaccine response.

Outbreak detection and response initiation We assume outbreaks are detected three months after onset in the status quo program, a delay calibrated against observed detection lags in past outbreaks (see the appendix). Once an outbreak is detected, decision-makers may initiate response vaccine R&D and manufacturing capacity investments. We assume they do so only when an outbreak's severity exceeds a given threshold, by default set at half our estimate of COVID-19 severity absent vaccine response.

Response vaccine R&D Once response is triggered, the world invests in response R&D for both "traditional" and mRNA vaccines, which have independent success probabilities and require distinct manufacturing capacity to produce. Based on COVID-19 vaccine timelines and published results from an expert survey,[17] we assume vaccine development takes 11 months to complete if a prototype vaccine already exists for the causal pathogen, and 18 months otherwise. Response R&D succeeds or fails stochastically. Platform and pathogen-specific success probabilities are estimated from our own expert survey.

Manufacturing capacity We assume the world has a baseline level of vaccine capacity from routine immunization demand. When an outbreak is detected, the world further invests in surge capacity. Half of the surge capacity installed for each outbreak is retained after the outbreak ends. Once vaccines are approved, only part (30%) of the available capacity—that which was dedicated to successful vaccine candidates—can immediately produce them; the remainder is either repurposed after a delay (40%) or cannot be repurposed at all (30%).[11]

Vaccine administration and harm reduction Vaccines are administered at the pace allowed by available manufacturing capacity, up to a maximum coverage level of 70%, consistent with herd immunity threshold for pathogens of plausible transmissibility[18] and observed limits on global vaccine acceptance.[19] Following previous evaluations of programs accelerating vaccine availability,[4,11] we model the share of social losses averted as a concave function of coverage, reflecting higher marginal returns to the early vaccination of vulnerable populations. As a final step, we apply a "haircut" factor, capturing a variety of reasons vaccines may fail to achieve their full potential for harm reduction including imperfect efficacy or the emergence of other interventions obviating the benefit of vaccination. The "haircut" factor scales down the concave function relating coverage to social losses so that the maximum coverage level can mitigate at most 50% of social losses.

Program costs We assume the world spends $6.9 billion developing response vaccines, as it did during COVID-19.[20] Surge capacity unit costs rise with the level of investment, reflecting an inelastic supply of vaccine capacity in the short run. We calibrate the supply elasticity using responses to our expert survey. Surge capacity retained after a pandemic incurs a 9% annual maintenance cost relative to its capital value.

Preparedness investment programs

We model four preparedness investments that enhance the status quo vaccine response: advance vaccine manufacturing capacity, advance R&D for prototype vaccines, universal flu vaccine R&D, and improved early warning systems. Figure 2 summarizes each program's key features.

Advance capacity We model this program as building sufficient production capacity over 30 years to vaccinate 70% of the global population within six months. Advance capacity is built at constant unit cost and is subsequently maintained at the same annual maintenance cost as retained surge capacity. Advance capacity is a lower-cost substitute for surge capacity. With advance capacity available, the world can lower its per-pandemic surge capacity investment by up to 25%, depending on how much advance capacity has been installed when the outbreak emerges.

Prototype vaccine R&D We model this program as investing in developing prototype vaccines for pathogens with pandemic potential that currently lack them (listed in appendix Table A4). Prototype vaccines reduce the vaccine development timeline from 18 to 11 months for the targeted pathogen, and increase the probability of successful response R&D by 11% for both technology platforms. Using data on the costs, success probabilities, and development timelines of past preventive vaccine candidates,[21] we estimate that a program achieving a 90% probability that at least one vaccine candidate succeeds will cost $1.06 billion per pathogen and take 11 years to complete.

Universal flu vaccine R&D This program is modeled as funding the development of a universal flu vaccine protecting against both current and future flu strains. The vaccine confers partial population immunity before a flu pandemic begins (as it substitutes for seasonal flu vaccines) and can be produced and administered as soon as an outbreak is detected. However, it is only half as effective as a strain-matched (response) vaccine. We calibrate the baseline uptake of universal flu vaccines at 10%, based on 2024 seasonal flu vaccination coverage.[22] Given the greater scientific difficulty of a universal vaccine, we assume it costs twice as much as prototype vaccine R&D to reach the same 90% probability of success. The world still undertakes response vaccine R&D when universal vaccines are available, and switches to administering those when available on account of their higher efficacy.

Improved early warning This program implements a global biosurveillance infrastructure enabling one month earlier outbreak detection, broadly consistent with simulated performance of wastewater surveillance systems.[23] However, the speedup only succeeds 70% of the time, and also introduces a 55% false-alarm rate, generating wasted response spending. We put the program's cost at $490 million annually by scaling up the cost of a proposed US wastewater and nasal-swab surveillance system[24] to the rest of the world.

Simulation and sensitivity analysis

We estimate future pandemic losses, as well as the gains from status quo and enhanced vaccine response, using Monte Carlo methods, running one million 200-year simulations. In each simulation, we first draw a full set of pandemic arrival, severity, and duration parameters from their estimated sampling distributions, so that our results reflect uncertainty in the underlying pandemic risk forecast. We then simulate a sequence of pandemics over each 200-year horizon and, where relevant, false alarms, before applying the vaccine response model described above to determine the harm mitigated and costs incurred by each intervention. To minimize the influence of simulation noise

on the comparison of programs, we use the same sequence of simulated pandemics across all of the scenarios.

The appendix reports analyses of the sensitivity of our estimates of pandemic loss and damage mitigation to each of our parameters, varying them one at a time from their default value to gauge their influence on the headline results.

Role of the funding source

Staff at CEPI assisted with the development of the survey instrument and helped identify respondents for our expert survey. They also provided feedback on early versions of the model. CEPI and Coefficient Giving had no role other role in this study beyond providing funding.

Results

Pandemic arrival and severity

Figure 3 presents our estimates of the annual risk of pandemics exceeding the severity indicated on the horizontal axis. The blue curve plots our primary specification using our curated dataset. A pandemic at least as severe as we estimate COVID-19 would have been absent a vaccine response (12.3 deaths per 10,000) has an annual 2% chance of occurring, corresponding to a 50-year expected recurrence rate. A pandemic at least as severe as the 1918 flu (171 deaths per 10,000)[6]—the worst experienced since 1900—has an annual 0.43% chance of occurring, equivalent to an expected 230-year recurrence rate. The red curve in the figure presents estimated annual risk using the full set of historical epidemics since 1900 with severities above the recording threshold. That the blue curve lies substantially below the red indicates that the curated data behind our primary specification was not "cherry picked" to exaggerate pandemic harm, instead producing more conservative estimates of pandemic risk.

Pandemic losses absent vaccine mitigation

In our primary specification, the expected present discounted value of social losses from pandemics is $101 trillion, equivalent to $4 trillion in losses annually (Table 1). Mortality losses at $2.5 trillion comprise 61% of total social losses, exceeding annual economic-output losses ($920 billion) and learning losses ($660 billion) combined. These estimates are sensitive to several parameter assumptions, most notably the assumed ceiling on pandemic severity, discount rate, and the dataset used for risk estimation. For instance, raising the severity ceiling from 200 to 1,000 deaths per 10,000 raises expected annual losses to $7.4 trillion, and to $21.2 trillion if raised further to 10,000 deaths per 10,000—effectively allowing outbreaks to threaten human extinction. All remaining individual parameters change estimated losses by 15% or less (appendix Figure A8).

Figure 4 highlights the outsized importance of the most severe pandemics. In our primary specification, the top severity decile accounts for more than half of expected social losses before economic growth and social discount factors are applied.

Status quo vaccine response

We estimate that status quo vaccine response provides $17.7 trillion in present discounted social benefits and saves an expected 38 million lives at a present discounted cost of $116 billion (Table 2). Since benefits are orders of magnitude larger than costs, the program's NPV at slightly under $17.6 trillion is hardly less than its gross benefits.

The high mean NPV is driven by a minority of simulations with quite severe pandemics: the median NPV across simulations of $6.3 trillion is well below the mean (Figure 5). The dispersion in NPV across simulations is correspondingly wide. At the 10th percentile, no pandemic arrives over the 200-year horizon and the program produces neither benefits nor costs; at the 90th percentile, NPV reaches $51 trillion.

While substantial in absolute terms, the harm mitigated by the status quo vaccine program

amounts to only 18% of total social losses on average. This reflects several frictions built into our model of vaccine response. The response is only initiated for outbreaks at least half as severe as "state of nature" COVID-19; response R&D has a risk of failing; some manufacturing capacity is assumed not to suit successful vaccines or can only be repurposed with delay; some outbreaks subside before vaccine contribute much benefit. The model also adds the conservative 50% "haircut" factor reflecting a range of contingencies dampening the ability of vaccines to mitigate harm.

As indicated by the gap between the blue curves in Figure 6, the status quo program mitigates an increasing share of harm at the high end of the severity distribution. Substantially more harm could be mitigated is if frictions to vaccine development and efficacy were eliminated (green curve).

Sensitivity analysis provided in the appendix (Table A6) suggest that our estimate of the NPV from the status quo vaccine response is robust to most parameter changes, although some have considerable effect on the results. The largest changes come from varying the "haircut" parameter taken off the top of program benefits and raising the severity ceiling on the underlying pandemic forecast.

Preparedness investments

All four standalone preparedness programs we model deliver benefits exceeding their costs (Table 2). However, their benefits and costs vary considerably. Universal flu vaccine R&D generates the largest incremental NPV, providing $2.4 trillion beyond the status quo program and saving 7.9 million additional lives. At a cost of $1.7 billion, the program achieves an exceptionally high benefit–cost ratio of 1,414.

Improved early warning provides the second highest NPV at $436 billion. Results reported in appendix Table A7 can be used to decompose program gains, showing that 82% comes from the straightforward benefit of accelerating vaccine rollout, but a small share (12%) stems from surge capacity installed in response to false pandemic alarms that is retained and thus accelerates vaccine

rollout in subsequent outbreaks. The improved early warning system costs a present discounted value of $12 billion, leading to a benefit–cost ratio of 37.

Advance capacity and prototype vaccine R&D generate similar NPVs at $269 billion and $256 billion, respectively. The cost of the advance-capacity program is higher, leading to a lower benefit–cost ratio of 15 compared to 102. The nature of both of these programs allows them to be flexibly scaled up or down. Results repored in appendix Table A8 show that prototype vaccine R&D has uniformly decreasing returns to additional spending, whereas advance capacity displays a more complicated pattern of initially increasing but eventually decreasing returns to scale.

The combination of all four preparedness programs generates an NPV of $3.3 trillion. This is within 1% of the sum of the standalone investment programs' NPVs. Our systematic analysis of the outcomes from implementing the programs in pairs reported in appendix Figure A7 reveals modest, positive synergies between most program pairs, with small, negative synergies for two pairs.

As with our model of the status quo vaccine response program, our models of the advance investment programs also rely on parametric assumptions to which our results may be sensitive. Appendix Table A7 reports how each standalone program's NPV changes in response to changes to key parameters.

Discussion

In contrast to our previous work in Glennerster et al.,[11] which analyzed an isolated pandemic-preparedness investment. the present paper is among the first to provide a comprehensive evaluation of a portfolio of pandemic preparedness investments (advance vaccine capacity, prototype vaccine R&D for neglected pathogens, universal flu vaccine R&D, and improved early-warning systems) within a common framework. This enables systematic comparison of the returns of different investments as well as exploration of potential synergies and substitution effects. We were also able to make wholesale improvements to the richness of the model and rigor of the methodology as

detailed above.

All four preparedness investments we analyze yield strongly positive NPVs and high benefit–cost ratios. Universal flu vaccine R&D stands out among the programs. According to both our historical data and the expert survey, flu may cause the majority of future pandemics. A universal vaccine against this common pathogen is especially valuable because it can provide protection from the first day of a pandemic rather than having to wait for a vaccine matched to the circulating strain to be developed. The program's NPV is high despite the baseline assumption that a universal vaccine is only half as effective as a strain-matched vaccine. In our simulations, the earlier availability of protection more than offsets this efficacy penalty. Our finding supports calls from the WHO, NIAID, and National Academies, among others, for research into a more broadly protective flu vaccine.[25–27]

The combined effect of the four preparedness investments is roughly additive. The positive and negative synergies we estimate (appendix Figure A7) are modest and largely offset each other. Together, the four programs have a combined NPV of $3.3 trillion.

Our layered approach to modeling the returns to preparedness—estimating harm absent a vaccine response, then the status quo response, then each enhancement—yields an important by-product: an estimate of the status quo response's own value. We find this to be $17.6 trillion, more than five times the NPV of the four enhancement programs combined. This has important practical implications. The status quo response is automatic in the model but not so in practice and cannot be guaranteed. The COVID-19 vaccine response required substantial public spending, policy coordination, and logistical capabilities, and may also have benefited from favorable scientific circumstances. The development of mRNA vaccines only weeks after the virus causing COVID-19 was sequenced was only possible due to years of prior research on mRNA vaccines and structural biology.[28] Our estimate of the value of the status quo response therefore underscores the importance of preserving the institutional, scientific, and logistical foundations of the existing response system.

The large returns to the status quo response and preparedness investments are perhaps un-

surprising against the backdrop of the enormous pandemic losses they address. In our primary specification, which embeds a variety of conservative assumptions, we estimate expected social losses from pandemics at $4 trillion *per year*. Less conservative specifications only strengthen the case for preparedness. For instance, relaxing the cap on the severity distribution from 200 deaths per 10,000 (approximately that of the 1918 influenza pandemic) to 10,000 deaths per 10,000 raises expected annual losses more than five-fold to $21.2 trillion, and nearly doubles the NPV of the status quo vaccine response. Additionally, switching from our curated subsample to the fuller sample of historical outbreaks from Marani et al.[6] raises expected annual harm from $4 trillion to $7.3 trillion.

Three further considerations support the claim that our forecast of pandemic risk is conservative. First, major determinants of pandemic risk (including population density, agricultural practices, global connectivity, and biotechnology) have changed substantially over time.[29] The available evidence suggests that pandemic risk has increased over our historical observation period,[30] suggesting that our approach of projecting stationary pandemic risk into the future may be conservative. Second, a leading pandemic-risk forecast based on epidemiological modeling[10] yields substantially higher pandemic risk than ours estimated from historical outbreak data (see appendix Figure A9 for a comparison). Third, our analysis is restricted to naturally emerging outbreaks. To the extent that engineered pathogens or other biosecurity threats raise pandemic risk, they would further increase the value of the status quo vaccine response and the enhancement programs studied here.

Several limitations of our study should be acknowledged. Even granting that our pandemic-risk forecast is conservative, the extremely wide confidence bands around the forecast (Figure 3) imply substantial uncertainty about realized future pandemic severity, including a meaningful possibility that it will fall well below the mean forecast. This uncertainty carries directly over for the NPVs of preparedness and response programs: their realized values could likewise fall well below our mean estimates. The wide interquartile and 10th–90th percentile ranges in Figure 5 reflect precisely this possibility. Our models of vaccine preparedness programs are highly stylized and thus no substitute

for the careful design of programs rolled out in the future, which will undoubtedly differ in many details. Our hope is to provide preliminary NPV estimates to inform decisions to proceed to design and implement prioritized programs. A number of key model parameters were set using responses to an expert survey. The difficulty in securing the participation of the leading experts we targeted led to a small survey sample of 17 participants, further reduced by science and manufacturing experts responded only to questions relevant to their domains. Still, we regard the expert survey as more reliable than isolated expert interviews, which in turn would be superior to setting the parameters to convenient round numbers or according to subjective author judgment, as done in previous work.[11]

We focused on investments that enhance vaccine response. Future work could use our framework to evaluate additional preparedness investments such as R&D for antivirals and rapidly configurable diagnostics, stockpiling of personal protective equipment, and installation of improved indoor air quality systems.

Exhibits

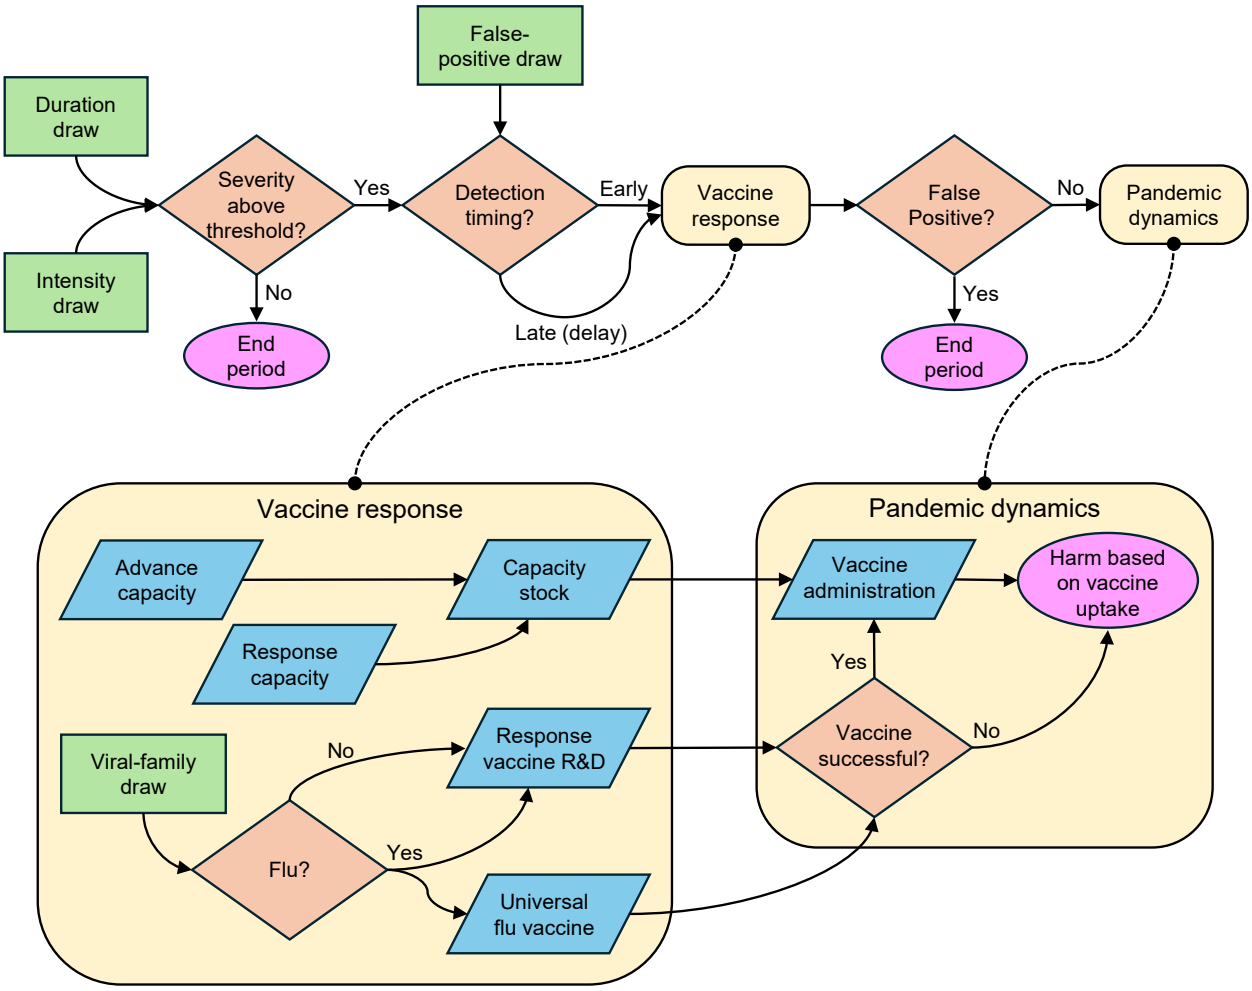


Figure 1: Flowchart of vaccine response model. Rectangles represent random draws, diamonds branch points, parallelograms process inputs and outputs, and ovals terminal nodes. Rounded beige rectangles are subroutines expanded at the bottom of the figure.

Preparedness investment	Example	Program highlights	Program spending
Advance capacity 	Expand vaccine-manufacturing capacity in reserve	Builds 8 bil. courses of reserve capacity over 30 years	$512 mil. annually for 3(years, afterwards $3 bil annually to maintain
Prototype vaccine R&D 	Develop prototypes for Lassa, RVF	90% chance prototype developed for three additional pathogen families. Prototype speeds development of matched vaccine in pandemic by 7 months and raises its probability of success by 14 pp.	$3.18 bil. annually over 11 year development period
Universal flu vaccine R&D 	Develop vaccine covering current and new flu strains	90% chance of developing universal vaccine. 50% effective against pandemic flus with no R&D lag. 10% of population pre-vaccinated.	$2.12 bil. annually over 11 year development period
Early warning 	Global wastewater and nasal-swab surveillance	60% chance of one month earlier warning, 30% false-positive rate	$490 mil. annually
Combined 	All of the above	All features of Investments above	Sum of costs from moderate program versions

Figure 2: Summary of preparedness investments analyzed. Abbreviation "bil." stands for billion, "mil." for million, and "pp." for percentage points.

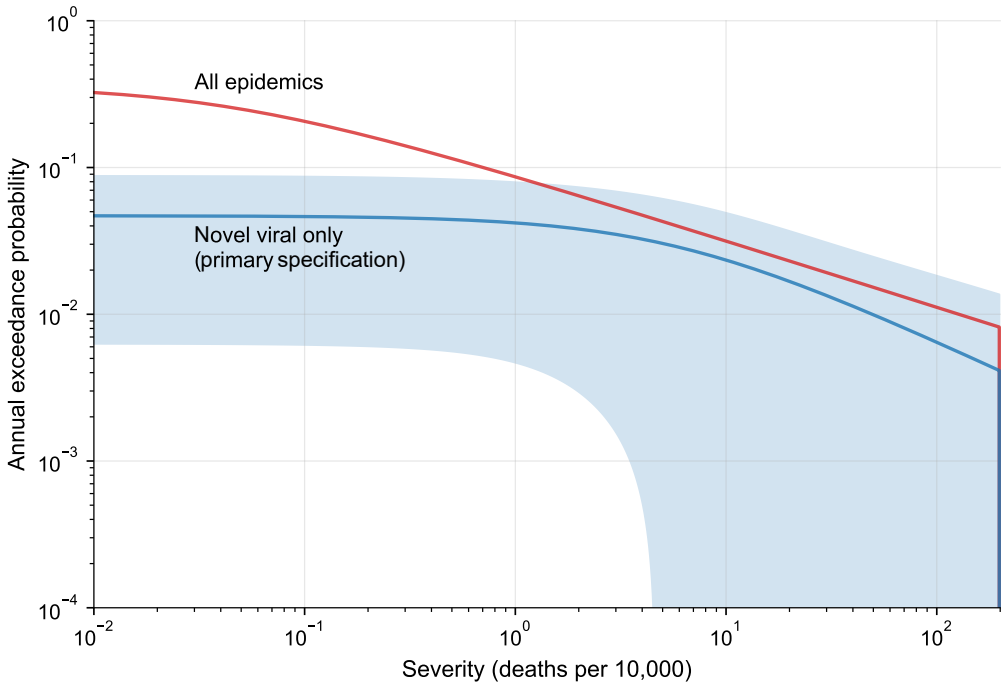


Figure 3: Estimates of future pandemic risk. Curves plot our estimates of the annual probability of an epidemic of at least the severity indicated on the horizontal axis. Blue curve is the primary specification, fit using curated subsample of novel viral outbreaks since 1900. Red curve presents alternative estimates using the full sample of outbreaks since 1900 from Marani et al.[6] Shaded regions are 95% confidence intervals around MLE estimates. Axes use log scales. Specification caps the maximum possible severity at 200 deaths per 10,000, reflected in the truncation of the curves at that point.

Specification	Expected annual harm				
	Annual deaths (million)	Annual social loss (trillion $)			
		Mortality	Economic	Learning	Total
Primary specification	1.2	2.5	0.9	0.7	4.0
Change severity ceiling					
• Increase to 1,000 deaths per 10,000	2.7	5.5	1.1	0.8	7.4
• Increase to 10,000 deaths per 10,000	9.3	18.9	1.3	1.0	21.2
Change social discount rate					
• Reduce to 2%	1.2	4.3	1.7	1.2	7.3
• Increase to 6%	1.2	1.9	0.7	0.5	3.1
Change pathogen data					
• Airborne novel viral outbreaks	0.9	1.9	0.4	0.3	2.6
• Novel viral outbreaks since 1950	1.1	2.2	1.0	0.7	3.9
• All outbreaks since 1900	2.3	4.8	1.4	1.0	7.3

Table 1: Estimated annual pandemic losses. Means across one million simulation runs. First row reports estimates under our baseline assumptions; subsequent rows vary one parameter at a time from its baseline value, holding other parameters fixed. The "Change severity ceiling" rows raise the maximum severity (deaths per 10,000) permitted in the underlying pandemic forecast from its baseline value of 200 deaths per 10,000. The "Change social discount rate" rows vary the annual rate used to discount future losses from its baseline value of 4%. The "Change pathogen data" rows recalibrate the model using sub- or supersets of the default pathogen data used to estimate the pandemic-risk forecast.

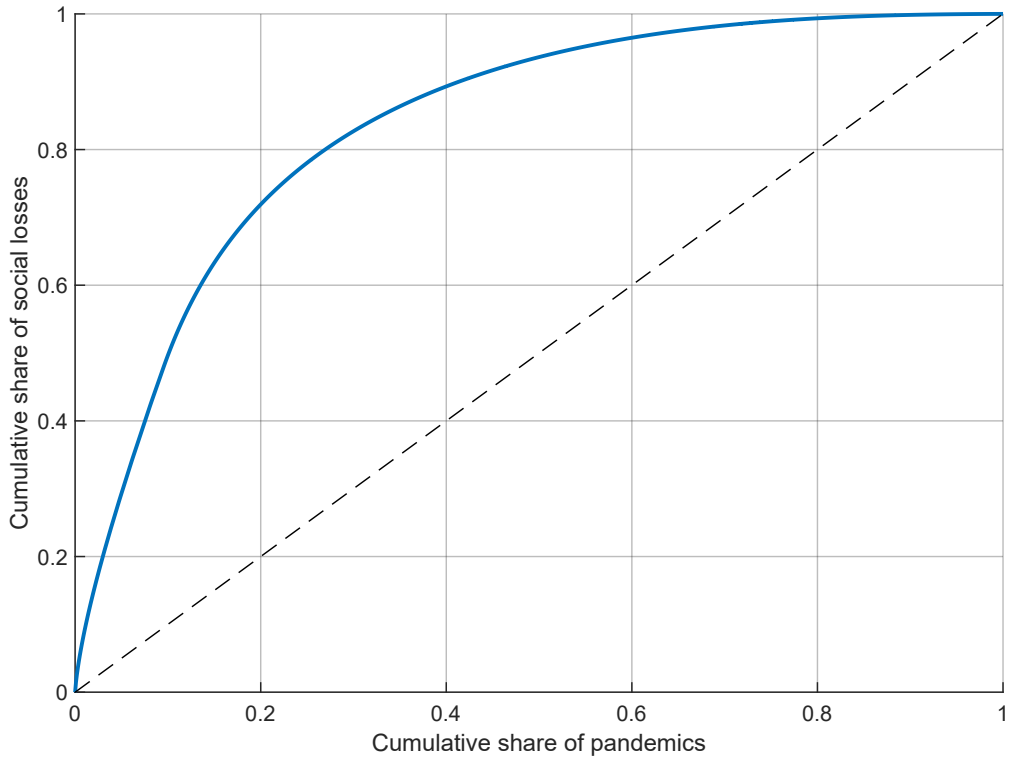


Figure 4: Concentration of losses in severest pandemics. Blue curve is an inverse Lorenz curve, constructed by ranking all the outbreaks in the million simulations in decreasing order of the total social losses they generate and calculating the cumulative share of losses across all the pandemics up to that point. Social losses in this curve do not account for time discounting or future growth; all pandemic losses are weighed as if they occur in the first year. For reference, dashed black 45° line corresponds to a uniform distribution of losses across outbreaks.

Program	Present discounted value ($ billion)			Benefit-cost ratio	Lives saved (million)
	Costs	Benefits	NPV		
Status quo response	116	17,699	17,583	152	38
Preparedness investments					
• Advance capacity	19	288	269	15	0.6
• Prototype vaccine R&D	2.5	258	256	102	0.7
• Universal flu vaccine R&D	1.7	2,387	2,385	1,414	7.9
• Improved early warning	12	448	436	37	1.3
• Combined	35	3,378	3,343	96	11

Table 2: Estimates of costs and benefits of vaccine programs. Means over the million simulation runs for each program. First row shows outcome from standard vaccine response. Subsequent rows report incremental values when preparedness investment is added to standard response.

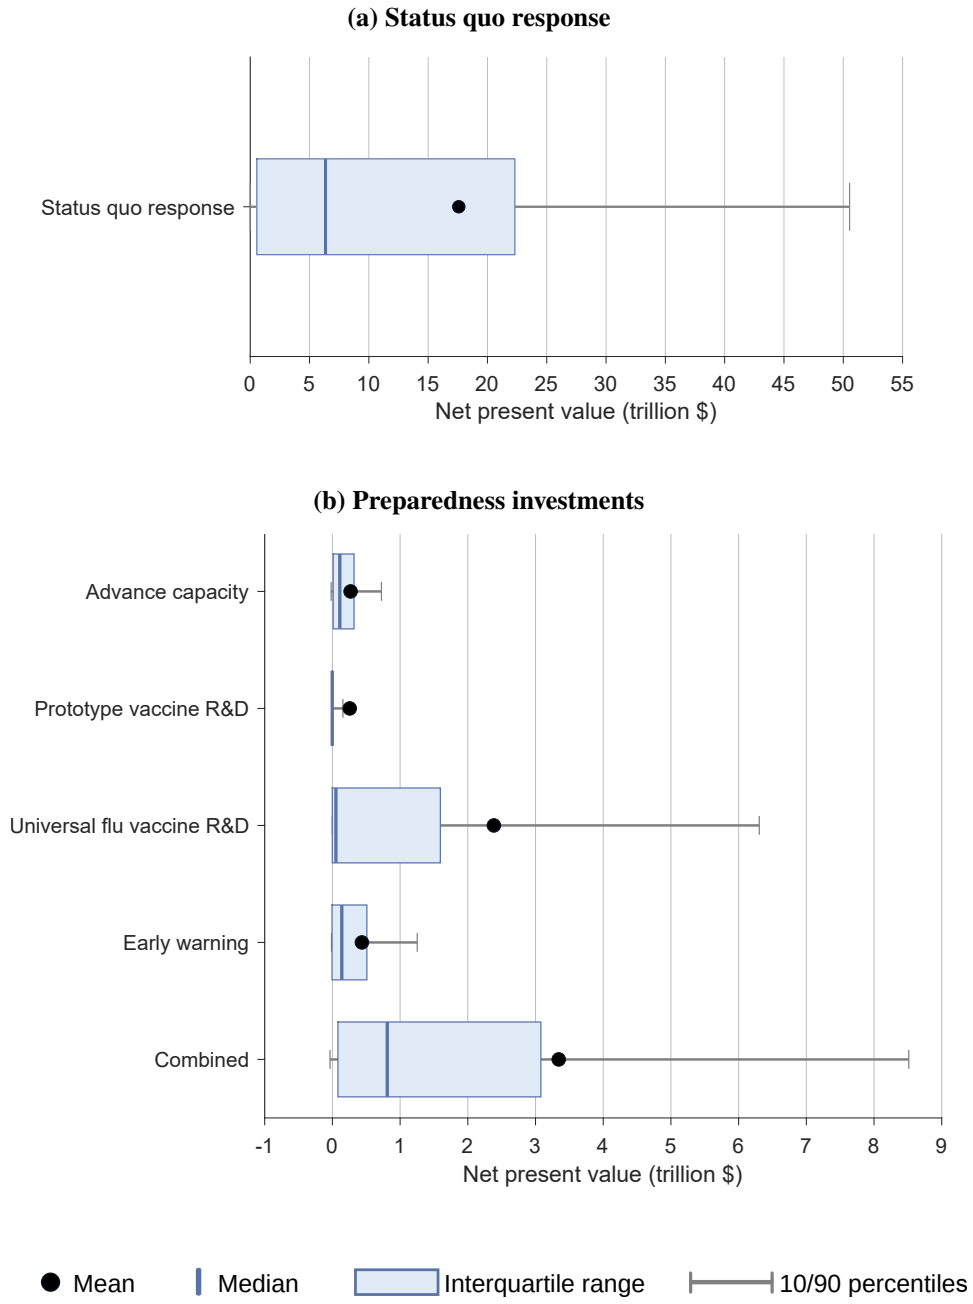


Figure 5: Distribution of vaccine programs' NPV. Summary statistics for the million simulations run for each program. Interquartile range and 10/90 percentile whiskers are not visible for the prototype vaccine R&D program since the 10th percentile outcome is the same as the 90th: in the simulation runs for both of those percentiles, the program involves a cost but no benefit as none of the pathogens covered by the new prototypes cause outbreaks over the simulation run.

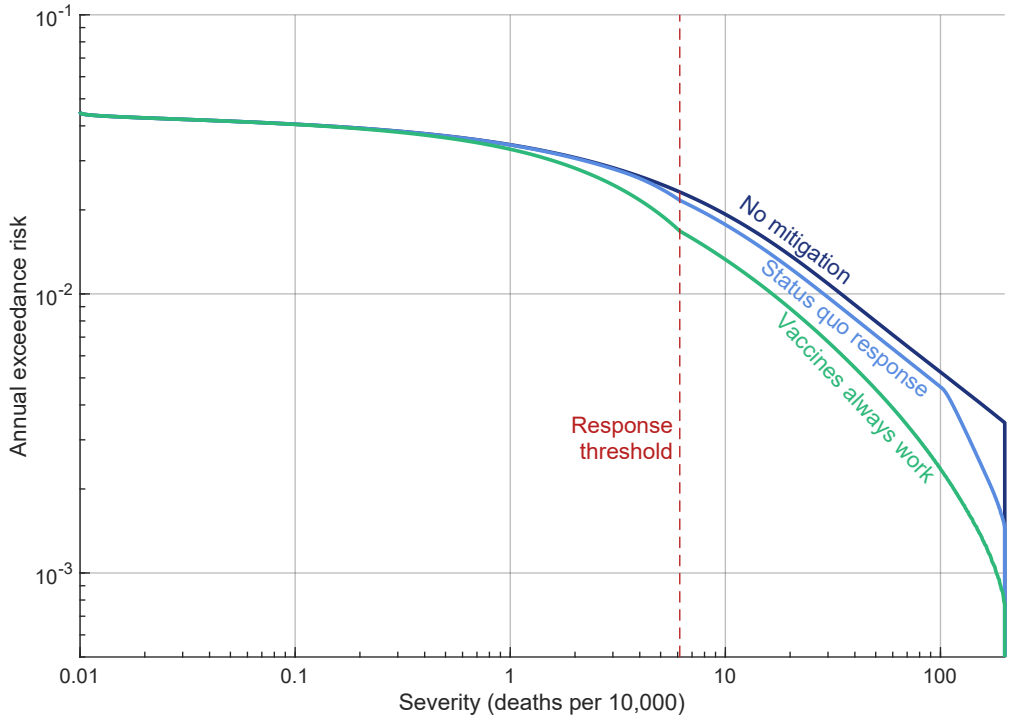


Figure 6: Mitigation of pandemic exceedance risk by status quo vaccine program. Dark-blue curve reprises the exceedance curve in the counterfactual absence of vaccine response from Figure 3. Light-blue curve shows the reduction in exceedance risk after mitigation by the modeled status quo vaccine response program. Green curve shows the exceedance curve under the status quo program if vaccines always work: i.e., setting the probability of technical and regulatory success for response R&D to 1 and the "haircut" taken off program benefits to 0. Curves drop vertically at the maximum-severity threshold of 200 deaths per 10,000. Dashed red line indicates the threshold severity for pandemic response.

Acknowledgments

We are grateful to Dimitrios Gouglas, Ragini Rangarajan, Alice Sabino, and Saul Walker at the Coaliation for Epidemic and Pandemic Preparedness Innovation for insightful comments on an early version of the model. We thank Andy Torkelson, Khatuna Giorgadze, Marina Theopold, and Peter Windus at the Linksbridge consultancy for helpful conversations and for providing data used to calibrate several model parameters. We thank Siddhartha Haria, Michael Kremer, Leah Rosenzweig, and Claire McMahon at the Market Shaping Accelerator for helpful feedback throughout the research process. Witold Wiecek provided invaluable feedback on an early set of methods and results.

This work was supported by the Coalition of Epidemic Preparedness Innovations (CEPI), which supported early work of CMS, CC, SC, and RG on the project. The analysis and conclusions are the authors' own and not subject to approval by CEPI. CMS is grateful for a gift from Coefficient Giving for support as the analysis and write-up were finalized. The content is solely the responsibility of the authors and does not necessarily represent the official views of CEPI or Coefficient Giving.

Author contributions

CMS contributed to conceptualization, methodology, validation, formal analysis, investigation, resources, writing the original draft, reviewing and editing the manuscript, visualization, supervision, and funding acquisition. CC contributed to conceptualization, methodology, software, formal analysis, investigation, and visualization. SC contributed to software, validation, investigation, data curation, project administration, and funding acquisition. RG contributed to conceptualization, resources, reviewing and editing the manuscript, supervision, and funding acquisition. GL contributed to methodology, software, validation, formal analysis, investigation, data curation, reviewing and editing the manuscript, and visualization. SQ contributed to methodology, software, validation, formal analysis, investigation, data curation, writing the original draft, reviewing and editing the

manuscript, visualization, and project administration.

Declaration of interests

Early stages of our work were supported by a contract from the Coalition of Epidemic Preparedness Innovations (CEPI). Our contracted deliverable was provided to CEPI on June 18, 2024. No financial relationship existed between CEPI and the authors beyond that point, and the work has been developed without their input since that time.

Data sharing

Data and code required to reproduce this work can be found at https://github.com/DevInnovationLab/pandemic_model_public. In line with participant confidentiality regulations, raw data from the expert survey are not available in the repository, but summary data in accordance with confidentiality requirements are represented.

AI use

Claude Code and Cursor were used to assist with writing and revising portions of our code. We further used Refine.ink to conduct a thorough review of a draft manuscript; we made several revisions to our work based on this feedback, primarily to remove typographical or recording errors. We also used ChatGPT to help generate icons used in Figure 2.

References

[1] Dattani S. What were the death tolls from pandemics in history? Our World in Data. 2023. Available from: https://ourworldindata.org/historical-pandemics.

[2] G20 High Level Independent Panel. Closing the deal: Financing our security against pandemic threats. Report on Financing the Global Commons for Pandemic Preparedness and Response. G20 High Level Independent Panel; 2025. Available from: https://nam.edu/pandemic-financing.

[3] Risko N, Werner K, Offorjebe OA, Vecino-Ortiz AI, Wallis LA, Razzak J. Cost-effectiveness and return on investment of protecting health workers in low- and middle-income countries during the COVID-19 pandemic. PLOS ONE. 2020 Oct;15(10):e0240503.

[4] Castillo JC, Ahuja A, Athey S, Baker A, Budish E, Chipty T, et al. Market design to accelerate COVID-19 vaccine supply. Science. 2021 Mar;371(6534):1107-9.

[5] Nascimento de Lima P, Karr S, Lim JZ, Vardavas R, Roberts D, Kessler A, et al. The value of environmental surveillance for pandemic response. Scientific Reports. 2024 Nov;14(1):28935.

[6] Marani M, Katul GG, Pan WK, Parolari AJ. Intensity and frequency of extreme novel epidemics. Proceedings of the National Academy of Sciences. 2021 Aug;118(35):e2105482118.

[7] Fan VY, Jamison DT, Summers LH. Pandemic risk: How large are the expected losses? Bulletin of the World Health Organization. 2018 Feb;96(2):129-34.

[8] Bernstein AS, Ando AW, Loch-Temzelides T, Vale MM, Li BV, Li H, et al. The costs and benefits of primary prevention of zoonotic pandemics. Science Advances. 2022 Feb;8(5):eabl4183.

[9] Li G, Quaade S, Snyder CM. Forecasting Pandemic Tail Risk: Model Choice, Specification, and Inference. Dartmouth College; 2025.

[10] Madhav NK, Oppenheim B, Stephenson N, Badker R, Jamison DT, Lam C, et al. Estimated future mortality from pathogens of epidemic and pandemic potential. In: Sharma S, Bertozi SM, Fan VY, Jamison DT, Norheim OF, Oshitani H, et al., editors. Disease Control Priorities Fourth Edition (DCP4), Volume 2 Investing in Pandemic Prevention, Preparedness, and Response. 3rd ed. Washington DC: International Bank for Reconstruction and Development, World Bank; 2026. .

[11] Glennerster R, Snyder CM, Tan BJ. Calculating the costs and benefits of advance preparations for future pandemics. IMF Economic Review. 2023 Sep;71(3):611-48.

[12] Reiss RD, Thomas M. Statistical analysis of extreme values: with applications to insurance, finance, hydrology and other fields. Springer; 2007.

[13] United Nations. World population prospects 2024. New York; 2024. Available from: https://www.un.org/development/desa/pd/world-population-prospects-2024.

[14] Sweis NJ. Revisiting the value of a statistical life: An international approach during COVID-19. Risk Management. 2022 Sep;24(3):259-72.

[15] Azevedo JP, Hasan A, Goldemberg D, Geven K, Iqbal SA. Simulating the potential impacts of COVID-19 school closures on schooling and learning outcomes: A set of global estimates. World Bank Research Observer. 2021 Mar;36(1):1-40.

[16] International Monetary Fund. World Economic Outlook update: Rising caseloads, a disrupted recovery, and higher inflation. International Monetary Fund; 2022. Available from: https://www.imf.org/en/publications/weo/issues/2022/01/25/world-economic-outlook-update-january-2022.

[17] Jeyapragasan G, Graabak J, Luby S, Esvelt KM. Modeling the benefits of virus discovery and pandemic virus identification. bioRxiv. 2024.

[18] Ahuja A, Athey S, Baker A, Budish E, Castillo JC, Glennerster R, et al. Preparing for a pandemic: Accelerating vaccine availability. American Economic Association Papers and Proceedings. 2021;111:331-5.

[19] Lazarus JV, Wyka K, White TM, Picchio CA, Rabin K, Ratzan SC, et al. Revisiting COVID-19 vaccine hesitancy around the world using data from 23 countries in 2021. Nature Communications. 2022;13(1):3801.

[20] Global Health Centre. COVID-19 vaccine R&D investments. Graduate Institute of International and Development Studies; 2021. Available from: https://www.knowledgeportalia.org/covid-19-vaccine-r-d-funding.

[21] Sertkaya A, Beleche T, Jessup A. New estimates of the cost of preventive vaccine development and potential implications from the COVID-19 pandemic. Office of the Assistant Secretary for Planning and Evaluation (ASPE) brief. Office of the Assistant Secretary for Planning and Evaluation; 2024. Available from: https://www.ncbi.nlm.nih.gov/books/NBK611605.

[22] World Health Organization. WHO seasonal influenza vaccine global market study, January 2024. WHO technical document. World Health Organization; 2024. Available from: https://www.who.int/publications/m/item/who-seasonal-influenza-vaccine-global-market-study-january-2024.

[23] Liu AB, Lee D, Jalihal AP, Hanage WP, Springer M. Quantitatively assessing early detection strategies for mitigating COVID-19 and future pandemics. Nature Communications. 2023 Dec;14(1):8479.

[24] Grimm S, Justen L, Kaufman J. Scaling US pathogen detection [Text]. Nucleic Acid Observatory. 2025 Sep. Available from: https://naobservatory.org/blog/biothreat_radar/.

[25] Erbelding EJ, Post DJ, Stemmy EJ, Roberts PC, Deckhut Augustine A, Ferguson S, et al. A universal influenza vaccine: The strategic plan for the National Institute of Allergy and Infectious Diseases. Journal of Infectious Diseases. 2018;218:347-54.

[26] World Health Organization. Global influenza strategy, 2019–2030. WHO technical document. Geneva: World Health Organization; 2019. Available from: https://www.who.int/publications/i/item/9789241515320.

[27] National Academies of Sciences, Engineering, and Medicine. Exploring Lessons Learned from a Century of Outbreaks: Readiness for 2030. Washington, DC: National Academies Press; 2019. Available from: https://doi.org/10.17226/25391.

[28] Covid Crisis Group. Lessons from the Covid War: An Investigative Report. New York, NY: PublicAffairs (Hachett Book Group); 2023.

[29] Board GPM. The Changing Face of Pandemic Risk: 2024 GPMB Pandemic Risk Report. World Health Organization; 2024. Available from: https://gpmb.org/reports/m/item/the-changing-face-of-pandemic-risk-2024-report.

[30] Smith KF, Goldberg M, Rosenthal S, Carlson L, Chen J, Chen C, et al. Global rise in human infectious disease outbreaks. Journal of the Royal Society interface. 2014;11(101):20140950.

Appendix

Estimating future pandemic harms and the gains from preparedness investments

Christopher M. Snyder*, Catherine Che, Sarrin Chethik,
Rachel Glennerster, Ganqi Li, and Sebastian Quaade

*Corresponding author. Email: chris.snyder@dartmouth.edu

Included in this PDF file:

Details on methods
Expert survey
Complementarity analysis
Sensitivity analysis
Tables A1–A8
Figures A1–A7

Details on methods

This section provides details omitted for brevity from the Methods section of the main text. It describes the data and methods underlying the major components of our analysis: our pandemic-risk forecast, our model of the status quo vaccine response to a pandemic, and our models of alternative pandemic preparedness investments, which enhance the vaccine response. It also describes how various datasets used in the paper were assembled and analyzed, including the expert survey, pathogen dataset, and economic-losses dataset.

Forecasting pandemic risk

Dataset Our forecast of pandemic risk starts with the dataset constructed by Marani et al.[1] The authors document 476 significant epidemics since 1600, 271 of which have information on duration and total deaths. While our coauthor team has used this dataset without modification in related past work,[2] for this paper, we curated the data to suit our focused forecasting exercise. Including, for example, bacterial plagues and recurring smallpox outbreaks dotting the historical record could overstate the expected social harm from novel pandemics not easily quelled by modern treatments and vaccines.

Figure A1 documents the steps we took to curate the final dataset used in our estimation. In the first step, we excluded outbreaks starting before 1900. This cutoff date strikes a balance between selecting a sample that reflects modern epidemiological conditions and providing an adequate sample size for estimation. In the second step, to reduce bias from smaller epidemics missing not

at random,[3] we excluded outbreaks with fewer than 0.01 deaths per 10,000 people annually—the stated recording threshold in Marani et al.[1] In the third step, we exclude bacterial epidemics, as their contribution to pandemic risk may have changed dramatically with the introduction of antibiotics. In a fourth step, we exclude epidemics with unidentified pathogens. Pathogen identities are required to make a definitive determination of novelty used in the final step, and without such information, the conservative choice is to exclude them. In the final step, we include only the first observed epidemic of each pathogen, ensuring the sample reflects the threat from outbreaks for which countermeasures are not yet in place. Influenza is the sole exception to this rule, due to its heightened propensity for antigenic shift.[4] In later tests of the sensitivity of the results to various steps in the data curation, we find that our approach leads to a more conservative forecast of pandemic risk.

We supplement the dataset with two additional observations, COVID-19 and HIV/AIDS, which would have qualified for inclusion in the dataset in Marani et al.[1] but for their decision to exclude ongoing epidemics. The authors had a practical reason for doing so. Their outcome variable of interest is epidemic intensity—deaths per year—which they compute by dividing the epidemic's total death toll by its duration. Duration requires an ending date for the epidemic, which is undefined for ongoing epidemics. For the only two ongoing pandemics that remain in our curated dataset, COVID-19 and HIV/AIDS, there is sufficient data to allow us to use a straightforward alternative to compute annual intensity: dividing the death toll to date by duration up to 2024, the latest date for which we have mortality data. Excluding ongoing pandemics introduces a selection bias against long-running epidemics, which we hope to avoid. Table A1 lists the observations comprising the sample used in our primary pandemic forecast.

Pandemic arrival and severity The generalized Pareto distribution (GPD) has been widely used for modeling the distribution of extreme natural events since Pickands[5] showed it is the limiting form for independent observations whose values exceed some high threshold. Applied to outbreak severity s, total deaths per 10,000 population over the outbreak's duration, the GPD cumulative distribution function (cdf) can be written

$$F_0(s) = \begin{cases} 0 & s < \underline{s} \\ 1 - \left[1 + \dfrac{\xi_s(s - \underline{s})}{\sigma_s}\right]^{-1/\xi_s} & \underline{s} \le s \le \bar{s} \\ 1 & s > \bar{s}, \end{cases} \tag{A1}$$

where ξ_s and σ_s denote the GPD shape and scale parameters, $\underline{s}$ denotes the threshold above which the GPD provides a good approximation of pandemic severity, and $\bar{s}$ denotes the upper bound at which pandemic severities are truncated. Assuming that epidemics above threshold $\underline{s}$ arrive according to a Poisson distribution at rate $\lambda(\underline{s})$ and are distributed according to a GPD conditional on arrival, one can show[6] that the annual maximum has a generalized extreme value (GEV) distribution with cdf

$$F(s) = \exp\big(-\lambda(\underline{s})\big[1 - F_0(s)\big]\big). \tag{A2}$$

The probability $\bar{F}(s) = 1 - F(s)$ that the most severe pandemic in a given year exceeds severity s (also referred to as the annual exceedance probability) thus equals

$$\bar{F}(s) = 1 - \exp\big(-\lambda(\underline{s})\big[1 - F_0(s)\big]\big). \tag{A3}$$

In our baseline specification, we set $\underline{s} = 0.01$, reflecting the recording threshold adopted by our data source, Marani et al.[1] We set $\bar{s} = 200$ in the baseline, conservatively truncating future pandemics at about the maximum severity observed in our dataset, that of the 1918 flu.

We obtain estimates of parameters $\lambda(\underline{s})$, ξ_s, and σ_s and their sampling distributions via maximum likelihood estimation (MLE). To ensure that the estimated sampling distributions respect the parameters' natural bounds, we apply MLE to the parameters transformed by the so-called "inverse softplus" function,

$$\Lambda(\underline{s}) = \ln\left(e^{\lambda(\underline{s})} - 1\right) \tag{A4}$$

$$\Sigma_s = \ln\left(e^{\sigma_s} - 1\right) \tag{A5}$$

$$\Xi_s = \ln\left(e^{1+\xi_s} - 1\right), \tag{A6}$$

ensuring $\lambda(\underline{s}), \sigma_s > 0$ and $\xi_s > -1$. The natural lower bound for ξ_s is -1 since in the limit $\xi_s \downarrow -1$, the GPD converges to a uniform distribution, representing an extreme case of light-tailed behavior over its bounded support. MLE yields an asymptotically unbiased estimate of the parameters and their associated variance–covariance matrix. While our curated dataset is small—making the asymptotic approximation harder to justify—our prior work in Li et al.[7] indicates that pandemic-risk forecasts using this asymptotic frequentist approach lead to slightly more conservative severity forecasts than an objective Bayesian alternative.

Pandemic duration We specify that the duration d of a pandemic in years is a lognormally distributed random variable with horizontal shift parameter $d_0 = 0.5$. The parameters μ_d and σ_d denote the mean and standard deviation, respectively, of $\ln(d - d_0)$. Our preliminary analysis could not reject the independence of severity and duration, so for simplicity we estimate the two MLEs separately. We estimate μ_d and σ_d via MLE from pandemic durations observed in our curated dataset. The estimated exceedance function for duration is plotted in Figure A2.

Pandemic risk by pathogen The benefits of some preparedness investments depend on the specific pathogen families generating the pandemic. In particular, neglected pathogen R&D—modeled as developing prototype vaccines for pathogen families currently lacking them—pays off if one of the target pathogen families happens to cause a future pandemic. If there is only a remote chance that pathogen family causes a future pandemic, the chance that the targeted investment pays off is correspondingly remote, reducing the investment's expected NPV.

This rich feature of the model requires forecasting risk separately by pathogen family, indexed by $j \in \{1, \ldots, J\}$. The families we consider are the nine named viral families on the priority lists maintained by CEPI[8] and the WHO,[9] plus two catch-all bins: one for other known viruses, and one for the emergence of an as-yet-unknown virus, which both CEPI and the WHO refer to as Disease X (see Table A4 for an overview of the pathogens). We then asked our expert respondents, of pandemics caused by any pathogen severe enough to spur world action, the relative proportion caused by each pathogen on the list (see the "Expert Survey" section for more details on the question). We assign pathogen families to future pandemics according to the average of experts' responses. Figure A3 plots these arithmetic average of the relative frequencies of outbreaks for specific viruses after rescaling expert responses to sum to 100%.

Estimating social losses

Pandemics create disruption and harm across multiple dimensions of life. When estimating social losses, we thus consider not only mortality but also economic and educational impacts, presenting a unified monetized valuation of these dimensions. Our approach closely follows the methodology used by our coauthor team in Glennerster et al.[2]

Mortality losses Our forecast of mortality losses from future pandemics depends on how much the population grows by the time the pandemics hit. Since severity s measures deaths per 10,000 population, which is a relative measure, a pandemic of severity s will result in 10% more deaths if it breaks out in a 10% larger population. Long-run global population trends are uncertain.[10] While there is little doubt population will continue to grow over the near term, it could slow, and possibly reach a peak at some uncertain future date before shrinking. To sidestep the need to forecast population trends, and to avoid creating the impression that optimistic population assumptions are driving pessimistic pandemic-harm projections, we hold the global population fixed throughout at its 2024 level of $N_0 = 8.14$ billion.

We also abstract from pandemic waves, supposing the annual death toll from a pandemic is constant each year over its duration. Let d_t be the duration (years) and s_t the severity (deaths per 10,000 over the duration of a pandemic) of the pandemic active in year t. Letting x_t denote deaths from this pandemic in year t, we have

$$x_t = \left(\frac{s_t}{d_t}\right)\left(\frac{N_0}{10{,}000}\right). \tag{A7}$$

To facilitate aggregating mortality with other social losses, we express mortality in monetary terms by multiplying lives lost times an estimate of the average global value of statistical life (VSL) of $V = \$1.3$ million.[11] Using a single average VSL avoids the issue of ascribing more value to lives in some countries than others. For readers uncomfortable with monetizing lives lost, mortality losses calculated using this VSL turn out to be several times greater than all other social losses combined, implying that the relative ranking of various vaccine programs is similar whether we use monetized total social losses or focus only on a count of lives saved.

Since VSL captures a person's estimated willingness to pay to avoid death subject to their budget constraint, it naturally rises with income. We assume that VSL scales proportionally with GDP per capita, which, following Guillemette and Turner,[12] we take to grow at long-run rate $g = 1.6\%$ annually. Mortality losses in year t are thus

$$ML_t = \left[V(1+g)^t\right] x_t. \tag{A8}$$

Output losses We model output losses as increasing with pandemic severity, reflecting that mortality, morbidity, and risk-mitigating behavior all obstruct economic activity. Formally, for the pandemic active in year t, we specify output lost as a percentage of global GDP over the duration of a pandemic, denoted ΔGDP, as increasing in the pandemic's severity s_t according to the power law

$$\ln\big(\Delta GDP(s_t)\big) = \beta_0 + \beta_1 \ln(s_t). \tag{A9}$$

We estimate coefficients β_0 and β_1 using Poisson quasi–maximum likelihood estimation[13] under the conditional mean assumption

$$E[\Delta GDP(s_t)] = \exp\big(\beta_0 + \beta_1 \ln(s_t)\big). \tag{A10}$$

We fit the model on the dataset of mortality and economic losses from past epidemics reported in Table A2, which builds on the analogous dataset compiled previously by members of our coauthor team in Glennerster et al.[2] We reviewed and corrected several entries based on the original sources and added an observation for the 1968 Hong Kong Flu.[14] The results are reported in Table A3 and Figure A4. The regression model fits the data almost perfectly, with a deviance of 0.006.

Defining fitted value $\widehat{\Delta GDP}(s_t) = \exp\big(\hat{\beta}_0 + \hat{\beta}_1 \ln(s_t)\big)$, our estimate of output losses in year t is

$$OL_t = \widehat{\Delta GDP}(s_t) y_0 N_0 (1+g)^t, \tag{A11}$$

where y_0 is the global GDP per capita in the base year.

Learning losses The World Bank conservatively estimated $10 trillion in lifetime learning losses resulting from the COVID-19 pandemic.[15] Following earlier work by our coauthor team in Glennerster et al.,[2] we calibrate learning losses for future pandemics by assuming that their ratio with output losses is the same as it was during COVID-19. The IMF estimated economic losses from COVID-19 at $13.8 trillion,[16] implying that learning losses in a pandemic in a future year t are given by

$$LL_t = \frac{10}{13.8} OL_t. \tag{A12}$$

Total social losses Define total social losses from the pandemic active in year t as the sum of the mortality, output, and learning losses described above:

$$TL_t = ML_t + OL_t + LL_t. \tag{A13}$$

The present discounted value of total social losses from pandemics over the entire T-year simulation run is

$$PV(TL) = \sum_{t=1}^{T} \frac{TL_t}{(1+r)^t}, \tag{A14}$$

where r is the social discount rate. The right-hand side is well defined because we prune epidemic draws to ensure there is at most one in any given year t.

We sometimes express these losses on an annualized basis, computing $AV(TL)$, the constant annual loss that would have to be experienced each of the T years of the simulation run to add up to the $PV(TL)$ total in present discounted value terms:

$$\sum_{t=1}^{T} \frac{AV(TL)}{(1+r)^t}. \tag{A15}$$

Using the standard formula for a geometric series of length T and rearranging yields

$$AV(TL) = \left[\frac{r(1+r)^T}{(1+r)^T - 1}\right] PV(TL). \tag{A16}$$

Monte Carlo simulation The complexity of the pandemic model does not admit closed-form solutions for the expected value of social losses, lives lost, and other outcomes nor other moments of their distributions. We therefore turn to Monte Carlo simulation to derive estimates of these outcomes. We simulate $I = 1{,}000{,}000$ runs, each of length $T = 200$ years using the following steps.

Step 1: Draw parameters for pandemic-severity distribution. In the first step of Monte Carlo run $i \in \{1, \ldots, I\}$, we draw parameter vector $(\Lambda(\underline{s})^{(i)}, \Sigma_s^{(i)}, \Xi_s^{(i)})$ from the asymptotic distribution of pandemic severity obtained via MLE. The asymptotic distribution is multivariate normal with mean given by the MLE parameter estimates $(\hat{\Lambda}(\underline{s}), \hat{\Sigma}_s, \hat{\Xi}_s)$ and variance-covariance matrix given by the inverse Fisher information matrix from the MLE. These estimates are transformed back into their natural analogs $(\lambda(\underline{s})^{(i)}, \sigma_s^{(i)}, \xi_s^{(i)})$ via the softplus function, the inverse of the inverse softplus transformation employed in equations (A4)–(A6), to ensure the estimated sampling distribution respects the parameters' natural bounds.

Step 2: Draw parameters for duration distribution. Next, we draw parameters $(\mu_d^{(i)}, \sigma_d^{(i)})$ from the asymptotic distribution of duration obtained via MLE. The asymptotic distribution is bivariate normal with mean given by the MLE parameter estimates $(\hat{\mu}_d, \hat{\sigma}_d)$ and variance-covariance matrix given by the inverse Fisher information matrix from the MLE.

Step 3: Simulate pandemics each year. For each year $t \in \{1, \ldots, T\}$ in Monte Carlo run i, we draw pandemic severities $s_t^{(i)}$ from the GEV distribution $F(s)$ in equation (A2) resulting from parameters $(\lambda_a^{(i)}, \sigma_s^{(i)}, \xi_s^{(i)})$ drawn in step 1. With probability $\exp(-\lambda_a^{(i)})$, the severity does not exceed the recording threshold $\underline{s}$. The GEV specification automatically returns severity draw $s_t^{(i)} = 0$ in this case. For all t with a positive severity draw, we also draw duration $d_t^{(i)}$ from the shifted lognormal distribution with parameters $(\mu_d^{(i)}, \sigma_d^{(i)})$, and round it to the nearest whole year. Finally, we draw the type of pathogen $j_t^{(i)}$ from the distribution of relative frequencies estimated from the expert survey.

Step 4: Prune overlapping pandemics. Since pandemics exceeding severity $\underline{s}$ are rare, it is unlikely that any given pandemic is overlapped by another. With many simulation runs and many years, however, it is possible that some pandemics overlap. This will happen if two pandemics are drawn in close succession and the second starts within the duration of the first. We resolve overlapping pandemics by cutting short the less intense outbreak in favor of the more intense one and rescaling the severity of the first outbreak in proportion to the reduction in its duration, maintaining the intensity of the first outbreak as it was initially drawn. With the large number of simulations we run (one million), a second outbreak is drawn within the duration of an ongoing one in a tiny fraction of cases (0.08% of draws).

Step 5: Compute outputs. The last step is to compute model outputs. To compute the expected number of lives lost, we sum lives lost over the 200-year horizon for each simulation run i and

then average this over the I simulations. We compute expected total social losses the same way except we average the present discounted value over the 200-year horizon according to equation (A14) rather than the simple sum. Other moments of the outcome distribution can be computed as desired from the saved Monte Carlo results.

A simpler approach for drawing random deviates would be to fix the distribution parameters at their MLE estimates and draw from that fixed distribution. Our approach described above adds a preliminary stage in which parameters from the asymptotic distribution of MLE estimates are drawn; random deviates are then drawn from this randomly drawn distribution. This approach better accounts for estimation error. Since the outcomes we seek to estimate are nonlinear functions of the parameters, by Jensen's Inequality, their expected values (let alone higher moments) can be biased if estimation error is not appropriately accounted for.

To minimize simulation error, we maintain a consistent set of pandemic draws for any vaccine program studied. There is a complication in doing this for the improved early-warning system, which is assumed to generate more false alarms, requiring a different set of pandemic draws. We address this issue while maintaining consistency across vaccine programs to the extent possible by drawing a superset of false alarms reflecting the highest false-alarm rate across any policy considered. We then omit a subset of false alarms to obtain the desired false-alarm rate for the policy under consideration.

Modeling status quo vaccine response

This subsection describes how we model vaccine response in the status quo (i.e., without further enhancement from preparedness investments). Figure 1 provides an overview of the dynamics of the vaccine response model. Our model of vaccine response takes the simulated sequence of pandemics as an input and overlays modules representing outbreak detection, vaccine development and regulatory approval, manufacturing scale up, vaccine administration, and harm mitigation. We calibrate the status quo response based on observations from the COVID-19 pandemic and responses from our survey of experts.

Initiating response We assume that decision-makers only initiate vaccine response for a pandemic if its severity s_t exceeds threshold $\hat{s}$ for which decision-makers are willing to invest in surge vaccine R&D and manufacturing capacity. Our baseline specification sets that threshold at $\hat{s} = 6.14$ deaths per 10,000, which is half our estimate of what the severity of COVID-19 would have been in the counterfactual case of no vaccine response. We obtain this counterfactual estimate by inverting our response model, as explained in a subsection below.

Detecting outbreaks How quickly the world responds to a potential disease outbreak depends on how early it is detected, which in turn depends on global biosurveillance capabilities. We model a stylized process of detecting outbreaks whereby policymakers detect and respond to outbreaks three months after their onset.

The three-month lag is roughly consistent with the early trajectory of several past outbreaks. For COVID-19, modeling studies suggest the index case may have been contracted as early as mid-October 2019,[17,18] roughly two and a half months before the Wuhan Municipal Health Commission

reported the outbreak to the WHO.[19] Of course, formal international recognition lagged considerably behind initial detection, with the WHO declaring COVID-19 a "Public Health Emergency of International Concern" on January 30, 2020. For the 2009 H1N1 flu pandemic, retrospective analyses place the first human cases sometime in February or March of 2009,[20,21] with the U.S. Centers for Disease Control confirming the first cases on April 17, 2009.[22] For the 2013–2016 West Africa Ebola, the detection lag was nearly four months: the outbreak began in December 2013, but the WHO was not notified until March 23, 2014.[23]

As discussed below, an enhanced early-warning system can shorten the detection lag, with the drawback of introducing some false alarms.

False alarms When a novel disease first breaks out, it is difficult to tell how severe it will ultimately be. Some fade quickly; others escalate into pandemics. Prudent policy may call for response efforts to be initiated before it is certain which path the outbreak will take given the potentially high returns to early response. Pre-emptive response, however, also risks wasting resources and disrupting society. A prominent example is the 1976 swine flu outbreak in the United States, where a localized cluster of influenza infections thought to resemble the 1918 flu prompted the rapid vaccination of over 43 million people, but the anticipated pandemic never materialized.[24]

To capture this dynamic, we incorporate "false alarm" outbreaks into our model. Let $f \in [0, 1)$ denote the false-alarm rate—the probability that an outbreak triggering response ultimately does not have significant health impacts. Treating both true outbreaks and false alarms as Poisson arrivals, the share f pins down the false-alarm arrival rate at

$$\lambda_f(\hat{s}) = \frac{f}{1-f}[-\ln F(\hat{s})]. \tag{A17}$$

The bracketed factor $-\ln F(\hat{s})$ is the annual arrival rate of at least one true outbreak exceeding $\hat{s}$, obtained by rearranging equation (A2).

We simulate false alarms by adding a step to the Monte Carlo procedure following step 4. In each year t in which no outbreak with severity $s > \hat{s}$ is ongoing, we draw false-alarm indicator A_t equal to 1 with probability $1 - \exp(-\lambda_f(\hat{s}))$ and 0 otherwise. False alarms are assumed to have no severity but trigger the sinking of some response investment, described in more detail below. Our baseline specification sets $f = 0$, reflecting (to the best of our knowledge) the lack of any documented false alarms triggering global vaccine response since 2000. By contrast, we allow investments in improved early-warning systems, described later, to entail a positive false-alarm rate.

Response R&D As soon as a pandemic exceeding the response threshold is detected, the world begins to invest in R&D to develop a vaccine against the new pathogen. We label this *response R&D* to distinguish it from *advance R&D* discussed later. Response R&D attempts to translate the scientific frontier about a pathogen family into a successful vaccine against the circulating variant. By contrast, advance R&D attempts to push that frontier outward by developing prototype vaccines for pathogen families that do not yet have them. The existence of a prototype for a pathogen family accelerates response R&D if that family turns out to be the source of the pandemic.

The world invests in response R&D for both traditional and mRNA technology platforms, each of which is manufactured using a distinct stock of manufacturing capacity. Let $\ell \in \{0, 1\}$

index the technology platform, where $\ell = 0$ represents traditional vaccines and $\ell = 1$ represents mRNA vaccines. This R&D takes time to complete before success or failure is realized. We calibrate success probabilities using responses to our expert survey and calibrate the time needed to complete response R&D based on the timeline for vaccine approval during COVID-19 together with estimates from an expert survey conducted by Jeyapragasan et al.[25]

The probability that vaccine R&D succeeds can differ markedly across diseases,[26–28] depending both on the scientific challenge presented by the pathogen and the amount of R&D directed toward relevant vaccines.[29,30] We therefore turned to expert elicitation, asking their subjective probabilities of technical and regulatory success (PTRS) across viruses and technology platforms. To simplify their task, the question asked respondents to report which equal-width decile bin $[\underline{p}_{ij\ell}, \bar{p}_{ij\ell}]$ for their subjective belief rather than a point value. To estimate sample means requires a method that accommodates interval rather than point observations.

Interval regression is a common technique for modeling the latent values underlying interval-censored outcome data.[31] The typical implementation of interval regression assumes that latent outcomes are normally distributed, which would violate the bounds of the unit interval $[0, 1]$ our observations were constrained to be in. We therefore used interval-censored beta regression suited to this purpose, modeling respondents' latent PTRS as draws from a beta distribution parametrized by a generalized linear model. More specifically, for each latent PTRS, denoted $p^*_{ij\ell}$,

$$p^*_{ij\ell} \sim \text{beta}(\mu_{pij\ell}, \sigma^2_{pij\ell}) \tag{A18}$$

where μ_{pije} is the distribution's mean and σ^2_{pije} is a convenient rescaling of the variance

$$\sigma^2_{pije} = \frac{\text{Var}(p^*_{ij\ell})}{\mu_{pij\ell}(1 - \mu_{pij\ell})}. \tag{A19}$$

Mean $\mu_{pij\ell}$ is assumed to be linear in the logit function:

$$\text{logit}(\mu_{pij\ell}) = \alpha_j + \alpha_\ell + \beta_i, \tag{A20}$$

where α_j and α_ℓ are pathogen and platform fixed effects and $\beta_i \sim \text{normal}(0, \sigma^2_\beta)$ is one of a set of normally distributed random effects capturing the general optimism of each respondent i toward the development of vaccines. The logit function is defined $\text{logit}(x) = \ln\big(x/(1-x)\big)$.

Each observation, consisting of a report that $p^*_{ij\ell} \in [\underline{p}_{ij\ell}, \bar{p}_{ij\ell}]$, contributes likelihood

$$B(\bar{p}_{ij\ell}) - B(\underline{p}_{ij\ell}), \tag{A21}$$

where B denotes the beta cdf. We estimate the parameters $\text{Var}(p^*_{ij\ell})$, $\{\alpha_j\}_{j=1}^J$, $\{\alpha_\ell\}_{\ell\in\{0,1\}}$, and σ^2_β via maximum likelihood estimation using the `gamlss` package in R.[32]

We estimate $\hat{p}_{j\ell}$, the conditional mean PTRS for pathogen j and platform ℓ, by integrating over the distribution of respondent random effects:

$$\hat{p}_{j\ell} = E_i(p^*_{ij\ell}) = \int_{-\infty}^{\infty} \text{logit}^{-1}(\hat{\alpha}_j + \hat{\alpha}_\ell + \beta)\frac{1}{\hat{\sigma}_\beta}\varphi\left(\frac{\beta}{\hat{\sigma}_\beta}\right)d\beta, \tag{A22}$$

where φ is the standard normal probability density function (pdf). We evaluate this integral using Gauss-Hermite quadrature, capturing estimation error via bootstrapping, resampling 1,000 times from respondent-level clusters. Panel (a) of Figure A6 reports our central estimates $\hat{p}_{j\ell}$.

The timeline for developing a vaccine for virus v depends on whether it has a prototype vaccine. When a prototype is available, we assume vaccine development takes $m^a = 11$ months, mirroring the time to first vaccine approval in the United States during COVID-19.[33] Although no coronavirus vaccine had received regulatory approval before the emergence of COVID-19, prototype vaccines for SARS and MERS, two closely related coronaviruses, had been developed using the adenovirus and mRNA platforms.[29,30] We set R&D completion without a prototype at $m^a = 18$ months, adding to the 11 month timeline the roughly seven month speedup that Jeyapragasan et al.[25] attribute to prototype availability on average.

Vaccine success is simulated as a draw of a Bernoulli random variable equaling 1 with probability $\hat{p}_{v\ell}$. If a 0 is drawn, the world has to respond to the pandemic without a vaccine in platform ℓ.

Manufacturing capacity The rate at which vaccines can be manufactured and administered depends on available manufacturing capacity.[34] Two sources contribute to the capacity available in the status quo: first, the stock of baseline capacity Q^b_ℓ available throughout each simulation as a result of peacetime vaccine demand and, second, surge capacity q^s_ℓ, which is constructed when a suspected pandemic exceeding the response threshold $\hat{s}$ is detected. (As discussed below, the advance-capacity program contributes an additional source, Q^a_ℓ.) Capacity stocks are indicated with an uppercase variable and flows with a lowercase variable. Capacity is specific to a technology platform ℓ and is measured in units of the vaccine courses that the facilities can produce annually.

We surveyed experts regarding global baseline capacity and calibrated the model using their mean response: $Q^b_0 = 5.9$ billion for traditional vaccines and $Q^b_1 = 2.9$ billion for mRNA vaccines.

When a new outbreak is detected, the world invests *at-risk* in surge capacity, that is, before vaccine success is guaranteed. At-risk capacity investment can waste resources—if the pandemic turns out to be a false alarm or the vaccine ends up failing to achieve technical or regulatory success. The advantage is that the capacity is available sooner than it would be if it only started to be built after regulatory approval[34]. We set the maximum feasible capacity surge following[2]: $\bar{q}^s_\ell = 2.25$ billion for both platforms $\ell \in \{0, 1\}$.

Baseline capacity is a stock, which persists throughout the simulation. Surge capacity accumulates stochastically during simulations since a portion of the surge capacity installed in past outbreaks is retained for use during future outbreaks. Let $Q^s_{\ell t}$ denote the stock of surge capacity for platform ℓ that has accumulated by the start of year t, initialized at $Q^s_{\ell 0} = 0$. As we saw from reports during COVID-19, the world lets some of the capacity spun up during a pandemic lapse after the pandemic wanes[35,36]. We use responses to our expert survey to calibrate the portion of surge capacity retained for the future as $\phi_q = 50\%$. This is added to the stock of surge capacity available for the next pandemic; the rest of surge capacity is allowed to lapse.

We impose an absolute ceiling $\bar{Q}$ on how much capacity the world lets accumulate. We set $\bar{Q} = 24.42$ billion—the amount needed to vaccinate 70% of the current global population within four months under the assumption that 30% of capacity will not be repurposable to manufacture successful vaccines. We further impose an individual ceiling on each technology platform ℓ of $\bar{Q}/2 = 12.21$ billion. The actual amount of surge capacity installed for technology ℓ in the event of a pandemic in year t is constrained by both the maximum feasible flow $\bar{q}^s_\ell$ and the absolute limit

$\bar{Q}/2$ on the global capacity stock for that technology:

$$q^s_{\ell t} \leq \min\left\{\bar{q}^s_\ell, \frac{\bar{Q}}{2} - (Q^b_\ell + Q^s_{\ell t})\right\}. \tag{A23}$$

We further treat condition (A23) as an equality, supposing that in the event of a pandemic the world installs the maximum capacity allowed by the constraints. Let

$$Q_{\ell t} = Q^b_\ell + Q^s_{\ell t} + q^s_{\ell t} \tag{A24}$$

be total vaccine capacity available to use in year t.

Once vaccines are approved, the installed capacity can start to produce vaccines. The model shifts from a yearly to a monthly time scale to allow a richer set of intra-year events to transpire. Following Glennerster et al.,[2] we assume that of the capacity $Q_{\ell t}$ available to use in year t for technology ℓ, 30% happens to be dedicated to successful vaccine candidates and can commence production immediately after vaccine approval. A further 40% is dedicated to unsuccessful candidates but can be repurposed to manufacture the successful vaccines after a delay. The remaining 30% is not able to be repurposed for any successful candidates. Based on responses from our expert survey, we set the delay to repurpose traditional capacity at $m^r_0 = 8$ months and to repurpose mRNA capacity to $m^r_1 = 2$ months. Overall, the model yields the following formula for $Q^*_{\ell t}(m)$, the manufacturing capacity for technology ℓ that can actively produce in month m:

$$Q^*_{\ell t}(m) = \begin{cases} 0 & m \in [0, m^a) \\ 0.3\dfrac{Q_{\ell t}}{12} & m \in [m^a, m^a + m^r_\ell) \\ (0.3 + 0.4)\dfrac{Q_{\ell t}}{12} & m \in [m^a + m^r_\ell, 12d_t]. \end{cases} \tag{A25}$$

The world fully employs all capacity available for active production in month m, totaling $Q^*_t(m) = Q^*_{0t}(m) + Q^*_{1t}(m)$.

Loss mitigation Vaccines mitigate more of the social loss from pandemics the more people that are vaccinated. Let $Z_t(m)$ denote the share of the population covered by the vaccine m months into a pandemic arising in year t. Active manufacturing capacity $Q^*_t(m)$ determines the rate at which vaccines can be administered during month m before reaching ceiling $\bar{Z}$:

$$Z_t(m) = \min\left\{Z_t(m-1) + \frac{Q^*_t(m)}{N_0}, \bar{Z}\right\}. \tag{A26}$$

We set the maximum possible coverage to $\bar{Z} = 70\%$, justified on several grounds. The maximum coverage could reflect the threshold for herd immunity beyond which vaccination provides no increase in protection.[37] Prior work adopted a 70% threshold for herd immunity.[38] The maximum coverage could also reflect vaccine hesitancy among the the remaining individuals. A cross-country survey found a vaccine acceptance rate of 75% in 2021, up from 72% the previous year.[39]

To model the social losses mitigated by coverage $Z_t(m)$, we build on an approach taken in previous evaluations of programs accelerating vaccine availability.[2,40] In those articles, the fraction

of social losses mitigated is specified as an nondecreasing, concave function H of coverage $Z_t(m)$ satisfying $H(0) = 0$ and $H(\bar{Z}) = 1$, the latter equality implying that full herd immunity is achieved at the coverage ceiling. The main modification we make here to their construction is to select the relevant (mortality) component for our H rather than averaging over multiple components, as those studies do.

The function H used in this study (Figure A5) is piecewise linear in the share of the population vaccinated, with three segments of successively declining slope. This concavity reflects diminishing marginal returns to vaccination: initial doses are allocated to the highest-priority groups, generating disproportionate reductions in mortality; later doses reach populations of lesser priority. Specifically, H is parameterized by two interior kink points, Z' and Z'', and a threshold for herd-immunity, $\bar{Z}$. In the first segment, from 0 to Z', vaccinations target high-priority populations and each dose reduces harm at a rate θ times that of a dose administered to a non-priority individual. In the second segment, from Z' to Z'', the slope falls to the level corresponding to vaccinations of non-priority individuals. In the third segment, from Z'' to $\bar{Z}$, the slope falls by half again, reflecting further diminishing returns as coverage advances toward herd immunity. At $\bar{Z}$, herd immunity is reached and no further pandemic harm can be averted: $H(\bar{Z}) = 1$. We later include the "haircut" factor $\gamma = 0.5$ to ensure that reaching herd immunity does not imply that vaccines eliminate all pandemic harm.

The first kink Z' corresponds to the share of the global population in high-priority groups. Following Castillo et al.[40] and Ahuja et al.,[38] we define these as the elderly and healthcare workers—those who face elevated morbidity risk or who are likely vectors of onward transmission. Data from the World Bank indicate that approximately 10% of the global population is aged 65 or older in 2024.[41] Healthcare workers account for a further roughly 0.8%, based on an estimated 65 million health workers globally in 2020.[42] Together, these groups sum to approximately $Z' = 11\%$.

A scale parameter ζ governs the steepness of the initial segment relative to the second segment, capturing the elevated benefits of vaccinating priority populations. Ahuja et al. calibrate ζ to range from 5 in the lowest GDP per capita country (Burundi) to 10 in the highest (Monaco), scaling linearly with per-capita GDP.[38] This calibration is inspired by epidemiological models of vaccine prioritization strategies,[43] which find that prioritizing older adults yields disproportionate mortality reductions relative to their population share. The gradient is smaller in lower-income countries, where age-mortality profiles are less steep. Because our analysis is conducted at the global level, we adopt a single value of ζ corresponding to the global average per-capita GDP. Applying the calibration using a global average GDP per capita of \$13,631, a minimum of \$219 (Burundi), and a maximum of \$288,001 (Monaco)—all from World Bank data in 2024[44]—yields

$$\zeta = 5 + 5 \times \frac{13{,}631 - 219}{288{,}001 - 219} \approx 5.23.$$

The second kink is set at $Z'' = 40\%$, after which, following Castillo et al.[40] and Ahuja et al.,[38] we assume the slope falls by half. Harm mitigation reaches its ceiling at $\bar{Z} = 70\%$. One of the justifications offered in the text for imposing a coverage ceiling $\bar{Z} = 70\%$ relates to its interpretation as the threshold for herd immunity. According to a standard susceptible-infected-recovered (SIR) epidemiological model, a perfectly effective vaccine achieves herd immunity above threshold

$$1 - \frac{1}{\mathcal{R}_0}, \tag{A27}$$

where $\mathcal{R}_0$ denotes the outbreak's basic reproductive number.[37] The specific 70% value for the threshold was arrived at by Castillo et al.[40] and Ahuja et al.[38] in the context of their evaluation of COVID-19 vaccination programs. As a starting point, the meta-analysis in Davies et al. puts the central estimate at $\mathcal{R}_0 = 2.7$ for COVID-19,[45] which implies a 63% threshold for herd immunity by equation (A27). The articles adjusted the threshold slightly upward to 70% to account for imperfect vaccine efficacy.

For future pandemics, the relevant $\mathcal{R}_0$ will of course vary and, along with it, the implied threshold for herd immunity. Still, the 70% threshold is broadly consistent with epidemiological evidence on pathogens with pandemic potential. A meta-analysis of studies of such pathogens finds that the class of highly transmissible coronaviruses has a median $\mathcal{R}_0$ well above 2.7.[46] While the median $\mathcal{R}_0$ for other classes is below 2, a value of 2.7 is well inside the 95% confidence intervals for virtually all of them. The assumption of $\bar{Z} = 70\%$ therefore represents a plausible central value for a large set of potential outbreaks.

The function H captures potential harm reduction from vaccine coverage, but several other factors can drive a wedge between potential and realized harm reduction. The vaccine may be imperfectly effective, herd immunity may not in fact be achieved at 70% coverage, or a treatment may emerge obviating the need for vaccination. We capture these contingencies by multiplying H by dampening factor $\gamma \in [0, 1]$. To be conservative in projecting the surplus from vaccine programs, we set $\gamma = 0.5$, in effect taking a 50% "haircut" off the top of the harm that the model predicts vaccines to mitigate.

According to the model, therefore, social losses mitigated by the status quo vaccine program in month m in year t a pandemic is active equal $\gamma H(Z_t(m))TL_t/12$. Dividing by 12 converts annual total social loss TL_t into the monthly time scale used to model pandemic response. Over the year t a pandemic is active, mitigated social losses amount to

$$\frac{\gamma TL_t}{12} \sum_{m=1}^{12} H(Z_t(m)). \tag{A28}$$

The analogous expression for lives saved is similar, replacing TL_t with x_t.

Program costs The status quo vaccine response program involves an array of costs: spending on response vaccine R&D, investment in surge capacity, vaccine production costs, and spending to maintain the fraction of surge capacity the world decides to hold on to. This subsection explains how those various components of cost are modeled. Throughout the discussion, monetary values are expressed in 2024 U.S. dollars.

For response R&D, we assume the world expends about what it did to develop vaccines during COVID-19, approximately $K_t^r = \$6.9$ billion.[47] In the event of a false alarm, the world invests only a fraction $\phi_f = 1/3$ of the intended response R&D before discovering the alarm was false. Hence, conditional on a pandemic emerging in year t, response R&D cost equals

$$[1 - (1 - \phi_f)A_t]K_t^r. \tag{A29}$$

A second category of cost is spending on surge capacity. A small amount of surge capacity for technology ℓ—up to a threshold $\hat{q}_\ell = 100$ million annual courses—can be installed at linear cost k_ℓ per unit. Above this threshold, the marginal cost of installing surge capacity increases, exhibiting

an elasticity of capacity supply ϵ, which we set to the round number 1 following Glennerster et al.[2] In the event that the surge was triggered by a false alarm, only fraction ϕ_f of surge capacity costs are incurred. Conditional on a pandemic emerging in year t, the cost of installing surge capacity for technology ℓ as such equals

$$\begin{cases} [1-(1-\phi_f)A_t]k_\ell q^s_{\ell t} & q^s_{\ell t} \le \hat{q} \\ [1-(1-\phi_f)A_t]k_\ell \left(\dfrac{q^s_{\ell t}}{1+\epsilon}\right)\left[\dfrac{\epsilon\,\hat{q}}{q^s_{\ell t}} + \left(\dfrac{q^s_{\ell t}}{\hat{q}}\right)^{\epsilon}\right] & q^s_{\ell t} > \hat{q}. \end{cases} \tag{A30}$$

Apart from incorporating false alarms, this functional form for surge capacity costs is identical to that in Castillo et al.[40] and Glennerster et al.[2]

Vaccine production incurs marginal cost c_ℓ per course for technology ℓ, resulting in production costs in month m for technology ℓ of $c_\ell Q^*_{\ell t}(m)$. Over year t during which a pandemic is active, production costs for vaccines in technology platform ℓ amount to

$$\sum_{m=1}^{12} c_\ell Q^*_{\ell t}(m). \tag{A31}$$

Surge capacity that is retained after a pandemic incurs maintenance costs. Based on responses from our expert survey, we take manufacturing capacity to depreciate at annual rate $\delta_q = 9\%$. We assume that the maintenance expense exactly offsets the depreciation for each technology ℓ, thus equaling

$$\delta_q k_\ell Q^s_{\ell t}. \tag{A32}$$

In the first year t a pandemic is active, total costs for the part of the status quo vaccine program using technology ℓ equal the fixed setup costs in equations (A29)–(A30) and the production costs in (A31). In subsequent years the pandemic is active, total costs equal the production costs in equation (A31) and the maintenance costs in equation (A32). In years with no active pandemic, total costs equal the maintenance costs in (A32).

Net social surplus Net social surplus from the status quo vaccine program in a year t experiencing a pandemic equals the mitigated social losses in equation (A28) net of the costs from Equations (A29)–(A32) that are relevant that year, added across the technologies $\ell \in \{0, 1\}$. The time series of net social surplus across years $t \in \{1, \ldots, T\}$ in a simulation run can be discounted in a similar way as in equation (A14).

We evaluate program performance by running Monte Carlo simulations as described previously, inserting the vaccine program after step 4 and the generation of false alarms to determine the losses mitigated by the program before computing outcomes in step 5. Most of the program features proceed deterministically, but three are stochastic. False alarms and success in vaccine response R&D are treated as draws from Bernoulli random variables with the probabilities specified above. Time to R&D completion is a deterministic function of the prototype availability for each outbreak's causal pathogen, a random variable we draw in step 3. The remainder of program features are fully deterministic, including the proportion of at-risk capacity that is repurposable for successful vaccines, the proportion of surge capacity retained post pandemic, and the timelines for repurposing manufacturing capacity.

Modeling preparedness enhancements

This section details how the status quo program is enhanced by the four preparedness investments we analyze—advance capacity, advance R&D, universal flu vaccine R&D, and early warning systems—discussed in turn.

Advance capacity We consider an advance-capacity program under which the world gradually adds sufficient production capacity to be assured of achieving the goal of $Q^g = 16.28$ billion annual courses in a future pandemic, split evenly across the two at $Q^g_\ell = 8.14$ billion annual courses each. The combined goal is the capacity needed to vaccinate 70% of the current global population within six months under the assumption that 30% of capacity will not be repurposable for successful vaccines. The advance capacity installed for platform ℓ equals the gap between the goal and baseline capacity:

$$Q^a_{\ell t} = Q^g_\ell - Q^b_\ell. \tag{A33}$$

The world takes 30 years to complete the program, installing $1/30$ of the target each year. Capacity built at this more relaxed pace in advance avoids the rising marginal costs exhibited in equation (A30), only costing k_ℓ per unit. With the benefit of advance capacity, we assume the world reduces its per-pandemic surge capacity investment by up to $\theta_q = 25\%$, such that the surge capacity constraint in equation (A23) becomes

$$q^s_{\ell t} \leq \min\left\{\bar{q}^s_\ell - \min\left(\theta_q \bar{q}^s_\ell, Q^a_{\ell t}\right), \frac{\bar{Q}}{2} - \left(Q^b_\ell + Q^s_{\ell t} + Q^a_{\ell t}\right)\right\}. \tag{A34}$$

Maintenance of advance capacity is modeled identically to maintenance of retained surge capacity: spending to reverse annual depreciation, amounting to $\delta_q k_\ell Q^a_{\ell t}$ by analogy to equation (A32).

Advance vaccine R&D We model an advance R&D program as the development of prototype vaccines for each of the identified pathogen families with pandemic potential that currently lack them: Rift Valley Fever, Lassa, and Crimean-Congo Hemorrhagic Fever (see Table A4). Prototype vaccines increase the speed and the probability of technological and regulatory success (PTRS) of response vaccine R&D.

As discussed above, the availability of a prototype vaccine for a pathogen family reduces the vaccine development timeline m^a from 18 to 11 months. We also discussed that we derived an estimate $\hat{p}_{j\ell}$ of the PTRS for all identified pathogens j from the expert survey, including the three noted here without a prototype vaccine in the status quo. We re-use the expert survey responses to estimate $\widehat{\Delta p}_\ell$, the percentage-point improvement in PTRS for platform ℓ attributable to prototype availability. Ideally, we would estimate the effect of an available prototype on PTRS and response R&D lags by including an indicator $\alpha_{p\ell}$ for having a prototype in equation (A20). To streamline the survey, however, experts were not asked how the existence of a prototype would change probabilities of success, so the survey lacks necessary data to directly estimate $\alpha_{p\ell}$. Instead, we estimate $\Delta\hat{p}_\ell$, the effect of a prototype on the PTRS of response vaccines developed for technology platform ℓ, by taking the average difference in estimated PTRS across pathogens in set J_1 that have a prototype

vaccine in the status quo and pathogens in J_0 that do not:

$$\Delta\hat{p}_\ell = \frac{1}{|J_1|}\sum_{j\in J_1}\hat{p}_{j\ell} - \frac{1}{|J_0|}\sum_{j\in J_0}\hat{p}_{j\ell}. \tag{A35}$$

Panel (b) of Figure A6 reports the estimates.

Unlike response R&D, advance R&D to develop a prototype vaccine has the advantage of occurring during peacetime at a deliberate pace. We thus take the average costs, probabilities of success, and timelines for past preventive vaccines as representative for prototype vaccine R&D. To estimate costs, we start by averaging the costs and success probabilities of preclinical R&D, clinical trials from phase one through three, and regulatory approval reported by Sertkaya, Beleche, and Jessup,[48] excluding seasonal influenza vaccines. We then calculate the expected cost of a vaccine candidate as the sum of the R&D stage costs multiplied by the probability that a candidate reaches the given stage. Using this approach, vaccine candidates have a 9.4% probability of progressing from preclinical development through regulatory approval at an expected cost of $46 million. We assume that advance R&D investment will sponsor enough candidates to achieve a 90% probability that at least one candidate succeeds. Under the simplifying assumption that vaccine candidates for the same virus have independent success probabilities, this requires sponsoring roughly 23 candidates, implying a total advance vaccine R&D cost of $1.06 billion per disease.

By construction, then, the probability that the advance vaccine R&D program successfully develops a prototype is 90%. We set the timeline for program completion to 11 years, following average timelines outlined in Sertkaya, Beleche, and Jessup.[48] We assume the $1.06 billion R&D spending is spread evenly over that period.

Universal flu vaccine R&D A longstanding goal in vaccine innovation is to develop "universal" vaccines that protect not just against the currently circulating strain, but against a broader class of related pathogens.[49] For flu, this idea has largely focused on eliciting immunity to relatively conserved viral regions rather than highly mutable strain-specific sites.[50] Analogous efforts for coronaviruses aim for protection across variants of SARS-CoV-2 (the virus causing COVID-19) and other sarbecoviruses, making broadly protective vaccine design a plausible long-run objective for pandemic preparedness.[51]

Since flu was judged by our expert survey to be the most likely cause of future pandemics (see Figure A3), we focus on analyzing a program engaging in advance R&D to develop a universal flu vaccine. If the program successfully develops such a vaccine, it provides several advantages over ordinary prototypes or response vaccines. First, the universal vaccine creates a base level of preemptive global immunity against flu since individuals will substitute it for seasonal flu shots. We calibrate this base level of immunity, which starts providing protection immediately in month $m = 0$, at $U_t(0) = 10\%$. This value is based on the WHO's forecast of 850 million flu vaccine doses administered in 2024,[52] divided by the contemporaneous world population.[53]

Another modeled benefit of a universal flu vaccine is that, since it does not need to be tailored to the particular flu variant causing the future pandemic, it is available to be produced and administered as soon as a flu pandemic is detected. A modeled drawback is that the universal vaccine is assumed to be less effective than a vaccine resulting from response R&D that is tailored to the specific pandemic variant.

Suppose that a flu pandemic arrives in year t and a universal flu vaccine was already successfully

developed. While the world waits for response R&D to succeed in a tailored vaccine, all available production $Q_t^*(m)$ is devoted to producing the universal flu vaccine. Hence, universal vaccine coverage equals

$$U_t(m) = \min\Big\{U_t(m-1) + \frac{Q_t^*(m)}{N_0}, \bar{Z}\Big\}. \tag{A36}$$

for $m > 0$ and $U_t(0) = 0.1$.

If and when a response vaccine is developed, production shifts to this more effective vaccine from that month on. Unvaccinated individuals are prioritized first. Each month's output $Q_t^*(m)$ goes first to immunizing the unvaccinated, proceeding at a monthly rate denoted $\Delta Z_t^0(m)$. Once the target $\bar{Z} = 70\%$ of the population has received some sort of vaccine, whether universal or response version, the program returns to individuals who received a universal vaccine and upgrades them to the more effective response vaccine, proceeding at a monthly rate denoted $\Delta Z_t^u(m)$. The upgrading occurs at the maximum rate allowed by available capacity until all $\bar{Z} = 70\%$ of the population has received the tailored response vaccine. Letting $U_t(m)$ denote coverage by the universal vaccine, distinct from coverage by the response vaccine, denoted $Z_t(m)$ above, we have

$$\Delta Z_t^0(m) = \min\big\{Q_t^*(m)/N_0, \bar{Z} - Z_t(m-1) - U_t(m-1)\big\} \tag{A37}$$

$$\Delta Z_t^u(m) = \min\big\{Q_t^*(m)/N_0 - \Delta Z_t^0(m), U_t(m-1)\big\} \tag{A38}$$

$$Z_t(m) = Z_t(m-1) + \Delta Z_t^0(m) + \Delta Z_t^u(m) \tag{A39}$$

$$U_t(m) = U_t(m-1) - \Delta Z_t^u(m). \tag{A40}$$

By construction, total coverage by one or the other type of flu vaccine, $Z_t(m) + U_t(m)$, never exceeds the maximum possible, $\bar{Z}$.

To calculate harm mitigated by a universal flu vaccine, we need to extend the function H appearing in equation (A28). The earlier specification applied to a single vaccine type with a single level of efficacy, but addition of a universal vaccine leads to two types of vaccines with different efficacies. We extend the function so that harm is mitigated in proportion to the relative efficacies of the two vaccine types:

$$H\big(Z_t(m) + e_u U_t(m)\big), \tag{A41}$$

where $e_u \in [0, 1]$ is the relative efficacy of the universal vaccine compared to the response vaccine, set to $e_u = 0.5$.

We use the same cost structure for universal flu vaccine R&D as we did for advance R&D for prototype vaccines. As before, the program pursues 23 candidates each having a 9.4% chance of success to achieve a 90% chance of at least one success. The only difference is that the costs are scaled up so that costs at each stage are higher by factor $\kappa_u = 2$ to reflect the consideration that a universal vaccine is a more technologically ambitious target.

Although a key benefit of universal flu vaccine R&D is that it provides pre-pandemic immunization against flu outbreaks, we do not account for the cost of maintaining this level of immunization, as we assume the universal vaccine simply replaces existing seasonal flu vaccination at no added cost. If anything, by focusing only on the benefits of a universal vaccine during a pandemic—and omitting its potential advantage over seasonal influenza vaccines in preventing seasonal flu—we likely provide a conservative estimate of the net surplus generated by a universal flu program.

Early warning Investments enhancing early warning systems are assumed to implement a global biosurveillance infrastructure enabling an outbreak to be detected one month earlier than in the status quo. A one-month speedup is slightly better than that found for wastewater surveillance across a range of diseases studied in liu et al.,[54] who used a simulation-based model to predict outbreak detection time as a function of the detection system and a disease's epidemiological parameters.

We assume that the enhancement to the early-warning system is imperfect in two respects. First, there is a 30% chance that the improved early-warning system fails to speed detection, which then remains the same three months as in the status quo. Second, the effort to speed detection leads to more false alarms. Specifically, we assume that there are no false alarms in the status quo: an outbreak is not regarded as detected and response initiated until it is certain that it is a significant threat. The early-warning enhancement introduces false alarms. We set the false-alarm rate—the probability that an apparent outbreak ultimately does not result in significant harm—to $f = 55\%$. We estimated the false-negative rate and false-positive rate from a study evaluating the capability of the Kraken2 tool to detect pathogens in human and environmental samples.[55] These error rates are based on their estimates of the system's estimated specificity (one minus the false-negative rate) and precision (one minus the false-positive rate) at the confidence threshold that maximizes specificity (see their Figure 3). We round to the nearest 5% to avoid conveying undue precision. Error rates may be lower if the biosurveillance system is able to monitor continuously but higher if the system is required to detect unknown as well as known pathogens. Of course, the performance of future enhancements to biosurveillance systems might be quite different from Kraken2. In light of the considerable uncertainty concerning these error rates, we conduct a sensitivity analysis in Table A7, allowing for substantially different error rates.

We discussed above how false alarms at rate f are simulated in our Monte Carlo exercises. In simulations of the enhanced early-warning program, we also need to simulate the detection lag. Conditional on drawing an apparent outbreak, the detection lag is always reduced to two months if a false alarm is drawn. It is also reduced to two months in the case of a true outbreak if a draw of a Bernoulli random variable indicates earlier detection is achieved, which equals one with probability $100 - 30 = 70\%$, reflecting the 30% chance of no speedup discussed previously.

Calibrating the costs of enhancements to the early-warning system is challenging as evidence is scarce and estimates range widely.[56] One estimate puts the needed additional investment in global disease surveillance at several hundred million U.S. dollars;[57] another puts it at tens of billions.[58] Estimates vary not only because of limited data but also because authors differ on the scope of "disease surveillance," whether limited to pandemic preparedness or incorporating the functionality of the health system more generally.

We base our cost estimate on an advanced pathogen-agnostic biosurveillance system outlined for the United States by the Nucleic Acid Observatory (NAO).[59] Their proposed system would establish wastewater sample collection at five treatment plants and fifteen airports, in addition to nasal-swab sampling at thirteen U.S. airports, with samples transported to two private-sector laboratories for metagenomic sequencing. They project total system costs at $50.8 million annually.

To scale this proposed U.S. system to a global system, we first identify countries in which more than 30% of wastewater is treated.[60] We compute each country's treatment-weighted population by multiplying its population by its wastewater treatment rate and aggregate countries into the seven World Bank regions. We then scale the U.S. deployment figures by the ratio of each region's weighted population to the U.S. weighted population, round up to the nearest whole number, and sum across the seven regions to obtain the deployments needed for a global system. This yields 47

wastewater treatment plants, 140 airports for wastewater surveillance, 117 airports for nasal swab surveillance, and 23 sequencing machines. Substituting these values into the NAO's cost model produces an estimated annual implementation cost of $490 million.

We emphasize that this cost estimate is only a rough approximation of implementing improved early warning capabilities at a global scale. However, given the large program benefits we estimate, even large discrepancies in estimated program costs have little effect on the program's estimated NPV.

Inverting the response model to determine counterfactual pandemic harm

To estimate the benefits of vaccine preparedness and response programs, we first forecast pandemic severity in the counterfactual absence of vaccines, before overlaying each program onto that unmitigated counterfactual. Our forecast must therefore be based on outbreak mortality estimates that reflect the lack of a vaccine. Among the epidemics in our primary sample for pandemic risk estimation, all but one had no vaccine available over the course of their duration (Table A1). The exception is COVID-19, for which record-speed vaccine development and deployment played a large role in addressing the societal costs of the pandemic. Using the observed COVID-19 severity as our data point would therefore bias our estimate of unmitigated pandemic risk downward and understate the benefits of the studied programs.

To avoid this bias, we invert our vaccine-response model to back out how severe COVID-19 would have been in the counterfactual absence of vaccines. In particular, we solve for the values of two latent variables—the severity s^* of COVID-19 in the absence of vaccines and the proportion of harm γ^* mitigated by COVID-19 vaccines—that satisfy two moment conditions based on observable information. Our notational convention is to add stars to indicate latent variables and omit stars for their observable analogs.

The first moment condition is that applying the vaccine-response model to unmitigated severity s^* yields observed severity s:

$$s = \sum_{m=1}^{12d} \frac{s^*}{12d} \left[1 - \gamma^* H(Z_t(m))\right]. \tag{A42}$$

We set $s = 9.17$ to match the observed severity of COVID-19, $d = 5$ to match its observed duration in years, and the vector $Z_t(m)$ to match the cumulative vaccination coverage observed in each month m of the COVID-19 outbreak as reported by Our World in Data.[61] The second moment condition is that the model's output of total lives saved by vaccines over the duration of the COVID-19 pandemic matches the estimate of $L = 2.53$ million lives saved from Ioannidis et al.:[62]

$$L = \sum_{m=1}^{12d} \left(\frac{s^*}{12d}\right) \left(\frac{N_0}{10{,}000}\right) \gamma^* H(Z_t(m)). \tag{A43}$$

We set $m^a = 11$ months to match the observed lag between the start of the COVID-19 pandemic and the approval of the first successful vaccines in the U.S.. We use MATLAB's nonlinear equation solver to find the tuple(s^*, γ^*) satisfying equations (A42)–(A43).

Expert survey

Several parameters, such as pathogen-specific pandemic arrival rates, pathogen specific vaccine success probabilities, and inputs into vaccine capacity modeling, proved difficult to estimate or ground in evidence using publicly available data. We therefore surveyed experts with relevant domain knowledge to provide their subjective estimates of these parameters. Survey respondents were drawn from CEPI's network of experts, which includes academics, industry leaders, and policymakers. Their expertise extends across vaccine manufacturing, vaccine R&D, biosurveillance, and epidemiology. The survey instrument was developed in collaboration with CEPI staff and with Linksbridge, a public-health consultancy; the study protocol was approved by the University of Chicago Institutional Review Board (IRB24–0012). Potential survey participants were sent an email describing the survey and asking if they were willing to complete it. Informed consent was provided by anyone who started the survey. We reproduce the specific survey questions used to calibrate various model parameters in the relevant appendix sections below.

Pandemic arrival rates varying across pathogens

To estimate pathogen-specific pandemic arrival rates, we began by assembling a list of pathogens with pandemic potential by combining the list of priority pathogens from CEPI[8] with the list of priority diseases from the WHO.[9] To the list we added influenza and five additional, more diffuse categories: "drug-resistant bacterial infection"; "non-viral, non-bacterial (e.g., parasite, or fungus, please specify)'; "unknown virus"; and "other known virus."

Survey respondents were presented with this list and asked to provide answers to the following query.

> Assume there is a pandemic in the next 50 years. Below is a list of diseases. Please indicate a rough estimate of the proportion of pandemics (defined by the U.S. CDC as a disease outbreak that has spread across several countries or continents with a large number of cases) to spur world action that will come from *[name of pathogen]* over the next 50 years (values should sum to 100%).

Six experts provided responses. Where an expert's responses did not sum to 100%, we rescaled them so the rescaled values sum to 100%, preserving their relative proportions. Given that experts assigned zero or low probabilities to non-viral categories and our model focuses on viral outbreaks only, we ultimately omitted non-viral categories from the analysis, rescaling frequencies a second time so that they would sum to 100% across viruses alone.

Probability of vaccine success

While efforts to develop COVID-19 vaccines were an unprecedented success, beating the previous record for quickest time to approval by nine years,[63] similar successes in future pandemics are not guaranteed. To help calibrate the probability of success for future vaccines in our model, we asked members of our expert panel (indexed by i) to report their subjective views on a plausible range for the probability of technical and regulatory success (PTRS) for vaccines of a pandemic-time vaccine for each pathogen j and vaccine platform ℓ:

> Assume there is an outbreak of a pandemic-causing *[name of pathogen]*. Given current R&D on vaccines for pathogens within their respective virus families, please list the plausible range of PTRS for developing *[mRNA / traditional / either]* vaccines against this outbreak. Please enter one of the following in the boxes below: 0-9%, 10-19%, 20-29%, 30-39%, 40-49%, 50-59%, 60-69%, 70-79%, 80-89%, 90%+.

The survey had a large numbers of entries, not all of which were filled in by every respondent, generating an unbalanced panel, including between two and five responses for each pathogen and vaccine platform pair.

Baseline manufacturing capacity

To calibrate Q_ℓ^b, the status quo production capacity for each vaccine technology platform ℓ, we included the following question in our expert survey:

> Approximately how much global *[mRNA / egg-based / non–mRNA or egg-based]* vaccine production capacity currently exists that would require minimal repurposing to produce new vaccines (e.g., for emerging pathogens such as MERS or Lassa)?
>
> 1. 2.5 billion doses per year,
> 2. 5 billion doses per year,
> 3. 7.5 billion doses per year,
> 4. 10 billion doses per year,
> 5. Other: _____.

We received seven or eight responses for each vaccine type. Grouping the non-mRNA platforms together under the heading of "traditional" capacity, the mean of the respondents' estimates yield 8.3 billion and 4.1 billion doses of repurposing-ready traditional and mRNA capacity, respectively.

We convert the responses to annual courses by dividing by two (assuming a two-dose course). A final conversion is to scale up the respondents' estimates to a higher level of installed capacity required to yield the repurposable capacity envisioned by the respondents given some capacity will fail to be repurposable. Consistent with the assumption in the model that a total of 70% of at-risk surge capacity can eventually be repurposed for approved vaccines and 30% not, we scale up the expert responses by the reciprocal of 70% to yield the correct amount of repurposing-ready capacity. After the two rescaling steps, the status quo capacity is sufficient to manufacture $Q_0^b = 5.9$ billion courses of traditional vaccine and $Q_1^b = 2.9$ billion courses of mRNA vaccine annually.

Capacity costs

We calibrate the unit cost of building vaccine capacity in advance, k_ℓ, for each vaccine platform $\ell \in \{0, 1\}$, using the following question from our expert survey.

> For each of the below parameters, please list your best estimate. [...]
>
> - Unit cost of *[mRNA / traditional]* capacity in advance.

For traditional capacity, three experts responded, all estimating a unit cost of \$3 per annual course. For mRNA capacity, four experts responded. Their mean response was \$1.88 per annual course and median was \$1.50. Based on these responses , we set $k_0 = \$3$ for traditional capacity and $k_1 = \$1.5$ for mRNA.

Advance capacity and surge capacity retained after a pandemic must be maintained at a cost. As the discussion around equation (A32) indicates, we model the annual maintenance cost for a unit of capacity as the expenditure required to reverse the annual depreciation per unit, $\delta_q k_\ell$.

We obtained estimates of depreciation rate δ_q by posing the following question in our expert survey.

> For each of the below parameters, please list your best estimate. Our current estimates are listed in parenthesis. [...]
>
> - Depreciation rate of vaccine manufacturing facilities (8%).

Three experts responded to this question, with an average response of 8.7%. We set the depreciation rate by rounding to the nearest percent: $\delta_q = 9\%$.

Time to repurpose capacity

To calibrate the time required to repurpose available capacity for pandemic vaccines m^r_ℓ, we used the responses to the following question in the expert survey:

> For each of the parameters below, please record your best estimate in column 3 ("Your Estimate"). Baseline values and corresponding sensitivity ranges are shown in column 2.
>
Parameter	Pilot value (range for sensitivity)	Your estimate
Time to repurpose mRNA capacity	2 months (1–3 months)	
Time to repurpose traditional capacity	6 months (3–9 months)	

The question mentions values used in some of our pilot analysis to induce the experts to adjust that up or down.

On average, respondents suggested that traditional manufacturing capacity could be repurposed in 8.3 months, while repurposing mRNA capacity should be easier, requiring only required 2.3 months. For modeling purposes, we round these to the nearest whole number: $m^r_0 = 8$ and $m^r_1 = 2$.

Capacity retained after pandemic

To calibrate ϕ_q, the share of surge capacity maintained after a pandemic ends, we included the following question in our expert survey.

Total vaccine production capacity at the end of 2022 was roughly 15 billion doses per year across modalities. Roughly, how much of this capacity is still in use today and will be maintained going forward?

1. 50%,
2. 75%,
3. 90%,
4. 100%,
5. Other: _____.

We received ten responses with an average answer of 55%. This motivates the round number $\phi_q = 50\%$ as our baseline parameter choice.

Complementarity among preparedness investments

Figure A7 reports the results from our analysis of the complementarity among the four preparedness investments. This analysis simulates the outcome from each of the four run as standalone programs and then them running in all paired combinations of them.

Sensitivity analysis

The baseline value of model parameters are provided in Table A5. We proceed with the sensitivity analyses described below by varying the parameters one at a time from these baseline values.

Pandemic losses

Figure A8 presents estimates of expected pandemic losses as we vary underlying parameters one-by-one. Several parameters cause large changes in results. The severity ceiling on our risk forecast is the most consequential. Raising the ceiling to 1,000 deaths per 10,000 increases expected annual losses to $7.4 trillion. Raising it further to 10,000 deaths per 10,000—allowing outbreaks to threaten human extinction—raises expected annual losses to $21.2 trillion. Varying the social discount rate r—in effect putting more or less weight on harms farther out in the 200-year simulation horizon—causes substantial swings in estimated losses, from $3.1 trillion at a 6% rate to $7.3 trillion at a 2% rate. Remaining model parameters are less consequential, changing expected annual losses by 15% or less from the baseline when they are varied.

The last group of bars in Figure A8 presents expected losses as we vary the dataset used for estimating our risk forecast. Including only airborne pathogens removes the devastating HIV epidemic from the sample, reducing our estimate of expected losses by 35%. Restricting the sample to pandemics since 1950, despite removing the 1918 flu (by far the most severe outbreak from our sample), has surprisingly little effect on the estimates. The removal of the 1918 flu is offset by the window for estimating arrival rate λ narrowing to a period during which novel viral outbreaks were more frequent. The remaining sample changes serve to increase the social loss estimates. Including the full sample of outbreaks since 1900 raises the social loss estimate by 83%, corresponding to

the substantial effect of this sample change on estimates of pandemic risk seen in Figure 3 in the main text.

Status quo vaccine response

Table A6 documents how the expected NPV of status quo vaccine response changes as we individually vary model parameters. The largest changes come from varying the "haircut" parameter taken off the top of program benefits and raising the severity ceiling of the underlying pandemic forecast. A 40% reduction in the haircut—from 50% to 30%—increases the program's NPV by roughly the same 40%, from $17.6 trillion to $24.7 trillion. Raising the severity ceiling from 200 to 1,000 deaths per 10,000 raises NPV by more than $14 trillion, an 81% increase. Expanding the pathogen data used to calibrate the model from our default set to the full universe of epidemics since 1900 above the recording threshold raises NPV by $5.7 trillion, or 32%. Among response R&D parameters, shortening the R&D timeline for candidates with an existing prototype from 11 to 6 months raises NPV by 15%, while lengthening it to 16 months lowers NPV by 12%. The remaining parameter changes reported in the table move NPV by 11% or less in either direction.

Preparedness investments

Parameter sensitivity analysis Table A7 reports estimates of the expected NPV of the modeled preparedness investments as we individually vary key parameters. We discuss the changes in the results for each program in turn.

The NPV for universal flu vaccine R&D is so large to begin with that it remains large for all parameter changes. The program's NPV is roughly proportional to the assumed probability of R&D success. For example, cutting the probability that universal flu vaccine R&D succeeds by 45%—from 90% to 50%—cuts the program's NPV by about the same 45%. The program's NPV is also sensitive to the relative efficacy of the universal flu vaccine. Increasing its relative efficacy to parity with the standard response vaccine raises the NPV of the universal flu vaccine program by more than $600 billion.

The NPV of the improved early warning program depends on how much it speeds detection. If the speedup is assumed to be two months rather than one as in the baseline, the program's NPV almost doubles, from $436 to $841 billion. The program's NPV also depends on the error rates in the biosurveilance system or equivalently on the complements to these error rates: sensitivity (one minus the system's false-negative rate) and precision (one minus its false-positive rate). Raising the system's sensitivity from 70% to 90% adds just under $100 billion in NPV. The system's sensitivity can be reduced all the way to 0%—only generating false alarms—yet the program still maintains some positive NPV from the extra capacity retained after "mistaken" response to the false alarms. Remarkably, lowering the program's precision to 20%, such that the system on average generates four false positive for every one accelerated detection, raises the program's NPV by $106 billion, underscoring the immense value of additional warm capacity for pandemic response.

The NPV from prototype vaccine R&D is most sensitive to the changes in response vaccine R&D timelines, rising over 60% when the completion time for response vaccine development in the presence of a prototype falls from 11 to 6 months and falling by more than 40% when it is extended to 16 months. Reducing the prototype R&D success probability from 90% to 50% produces a similar downside effect.

The advance-capacity program's NPV is relatively insensitive to unit capacity costs since program costs are only a small fraction of program benefits in any event. It is more sensitive to the other parameters since they directly or indirectly affect the incremental capacity it delivers. The biggest change in the program's NPV reported in Table A7 comes from reducing the amount of surge capacity retained after a pandemic. If instead of retaining half of surge capacity, none is retained, the program's NPV increases by 53%, from $269 billion to $412 billion. Advance capacity is all the more important if it is the only source of capacity available at the start of pandemics after the first.

Additional sensitivity analyses in Figure A10 go beyond changes to program dimensions, investigating how the programs' outcomes are affected when alternative specifications of the pandemic forecast are used. The programs' NPVs are higher using specifications that entail more pandemic harm for them to mitigate. Increasing the severity ceiling from 200 to 1,000 deaths per 10,000 increases the NPV for all four programs by at least 88%. Using the less conservative sample of all outbreaks since 1900 also substantially increases the expected NPV of all our programs.

Program returns to scale The advance capacity and prototype vaccine R&D investments are both of a nature where they can be flexibly scaled: The advance-capacity program can vary the amount of added capacity, and the advance R&D program can research a smaller set of neglected pathogens. However, the programs exhibit different patterns of returns to scale (Table A8). Prototype vaccine R&D displays positive but decreasing returns to scale: increasing the number of neglected pathogens the program researches from one to three increases the program's NPV while reducing its benefit–cost ratio. The pattern of returns to scale is more complicated for the advance-capacity program. Doubling the program's scale from 50% to 100% of the capacity in the standard program increases both its NPV and benefit–cost ratio. A further doubling of the program's scale— from 100% to 200% of the capacity in the standard program—increases the program's NPV but reduces its benefit–cost ratio. The advance-capacity program thus exhibits increasing returns to scale at low capacity levels and decreasing returns to scale at high capacities. The explanation is that the displacement of surge capacity modeled in equation (A34) subtracts a fixed amount from the program's effective capacity gain, so doubling the program's scale more than doubles the effective capacity it adds. At larger program scales, that source of increasing returns to scale is dominated by other factors contributing to decreasing returns to scale, such as the concavity of the damage-mitigation function H.

Appendix exhibits

Conventions

All monetary figures in appendix exhibits are measured in 2024 U.S. dollars. The abbreviation "yrs." stands for years, "mo." for months, "bil." for billions, "mil." for millions, "prob." for probabilities, "pct." for percentile, "NPV" for net present value, and "R&D" for research and development.

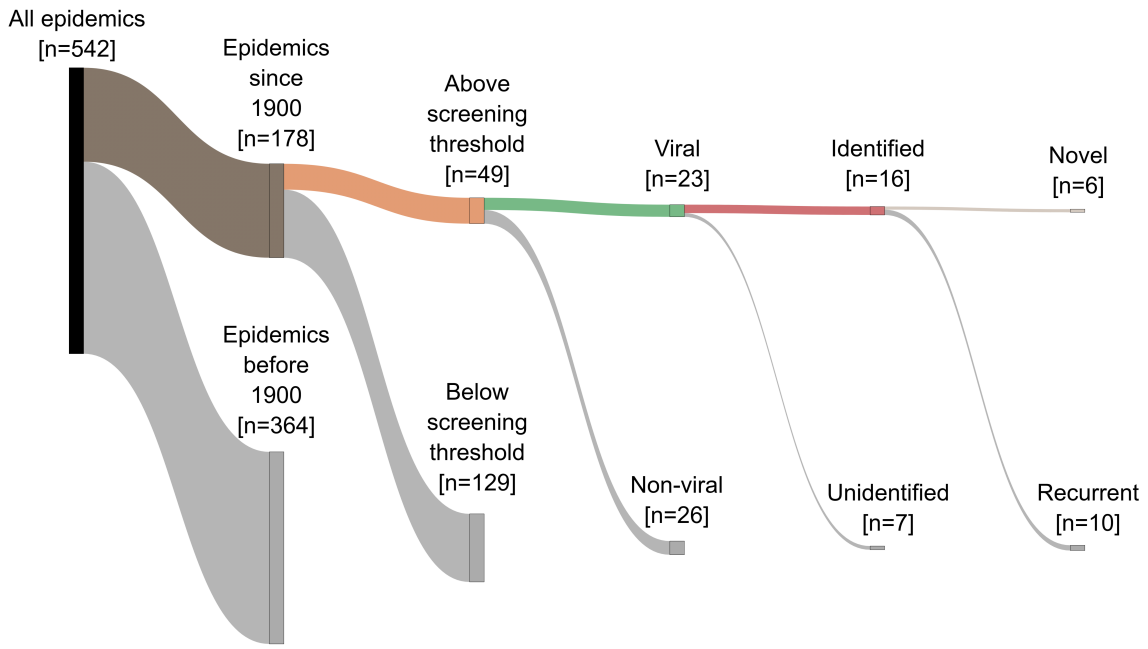


Figure A1: Curating data on historical epidemics to construct primary dataset used to estimate pandemic risk. Each colored node represents a filtration step used to extract a well-defined sample of past epidemic outbreaks for pandemic risk forecasting from the initial dataset assembled by Marani et al.,[1] adding observations for COVID-19 and HIV/AIDS.

Disease	Transmission mode	Start year	End year	Duration (years)	Severity (deaths per 10,000)
Influenza	Droplet	1918	1920	3	171
Influenza	Droplet	1957	1958	2	7.0
Influenza	Droplet	1968	1969	2	2.8
HIV/AIDS	Fluid	1981	2024	44	93
Influenza	Droplet	2009	2009	1	0.41
COVID-19	Airborne / droplet	2020	2024	5	12.3

Table A1: Major outbreaks in the curated dataset used to estimate severity distribution. Computing severity requires dividing by population, which we take to be the global population in the outbreak's start year. For COVID-19 and HIV/AIDS, we define duration as the years from the outbreak's emergence (or the earliest available mortality data) through 2024. Severity for COVID-19 is the figure wee adjusted to reflect the counterfactual severity in the absence of vaccine response, computed by inverting our pandemic-response model.

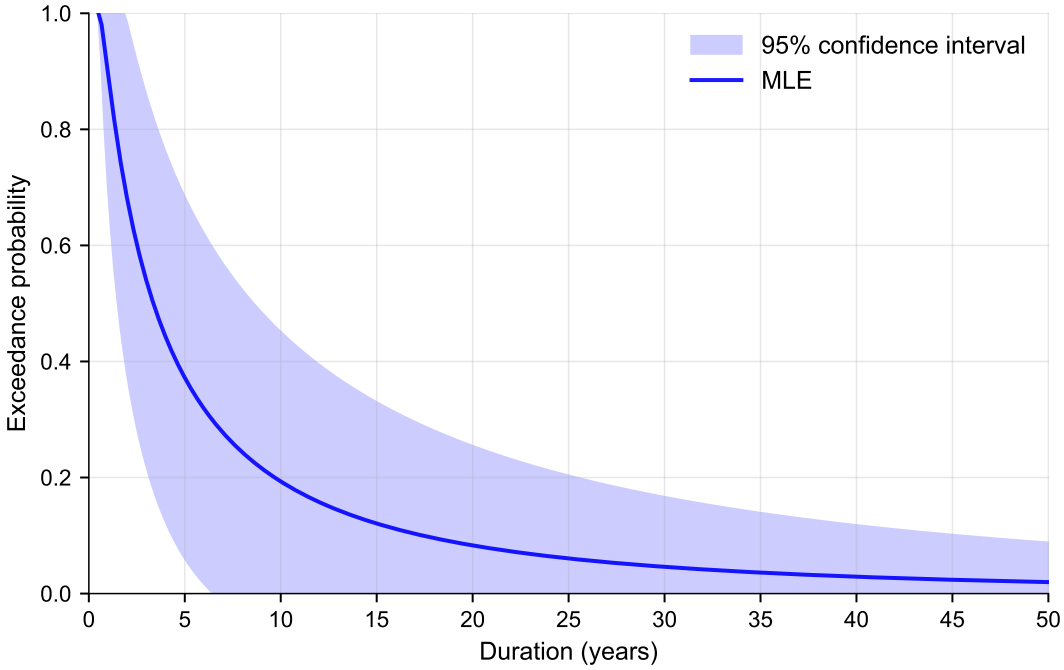


Figure A2: Estimated exceedance function for pandemic duration. Graphs the estimated probability that, conditional on arriving, a pandemic lasts at least as long as indicated on the horizontal axis. Pandemic duration is modeled as following a lognormal distribution. The curve is the exceedance function evaluated at the MLE estimates of the distribution parameters. The shaded area represents a 95% confidence interval for the exceedance function computed from the estimated asymptotic distribution of the MLE using the delta method. Durations are allowed to span the full 200-year simulation horizon, but the horizontal axis is truncated at 50 years for legibility.

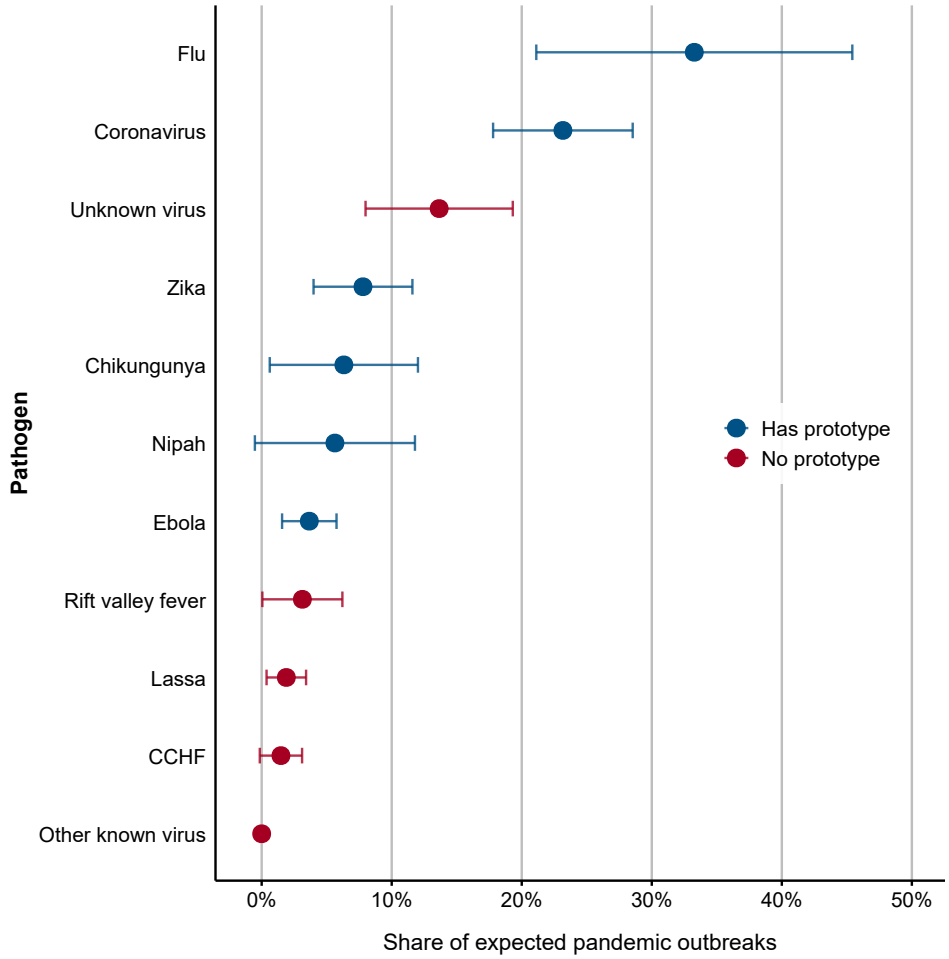


Figure A3: Shares of future viral pandemics generated by various pathogens according to expert survey. Points represent the average across $n = 6$ experts of the forecasted share of future pandemics caused by the indicated pathogen. Bars represent 95% confidence intervals assuming that survey responses are independently and identically distributed normal random variables.

Disease	Start year	End year	Severity (deaths per 10,000)	Total GDP loss	% GDP loss over total duration	Sources
1918 flu	1918	1920	171	3.0×10^{11}	5.53%	[64], [65]
COVID-19	2019	2024	9.17	1.7×10^{13}	2.49%	[66], [16]
Hong Kong flu	1968	1970	2.82	1.89×10^{12}	1.20%	[1], [14]
SARS	2002	2003	0.0012	1.8×10^{10}	0.01%	[1], [67]
Ebola	2013	2016	0.0157	1.9×10^{10}	0.003%	[1], [68]
Zika	2015	2017	5.6×10^{-5}	9.29×10^{9}	0.002%	[69], [70]

Table A2: Selected pandemics with published sources for economic losses. Outbreaks in the last century typically have scores of studies on mortality losses but more rarely on economics losses. The pandemics in this table have published estimates of economic losses, which provide data points that we can use to estimate the association between severity and economic losses.

Variable	Estimate
ln(Severity)	0.401
	(0.040)
Intercept	-4.873
	(0.203)
Deviance	0.006

Table A3: Estimates from Poisson regression relating economic losses to pandemic severity. Estimates of Poisson regression specified in equation (A10) using the six observations from Table A2. Huber–White heteroskedasticity-robust standard errors reported in parentheses below coefficient estimates.

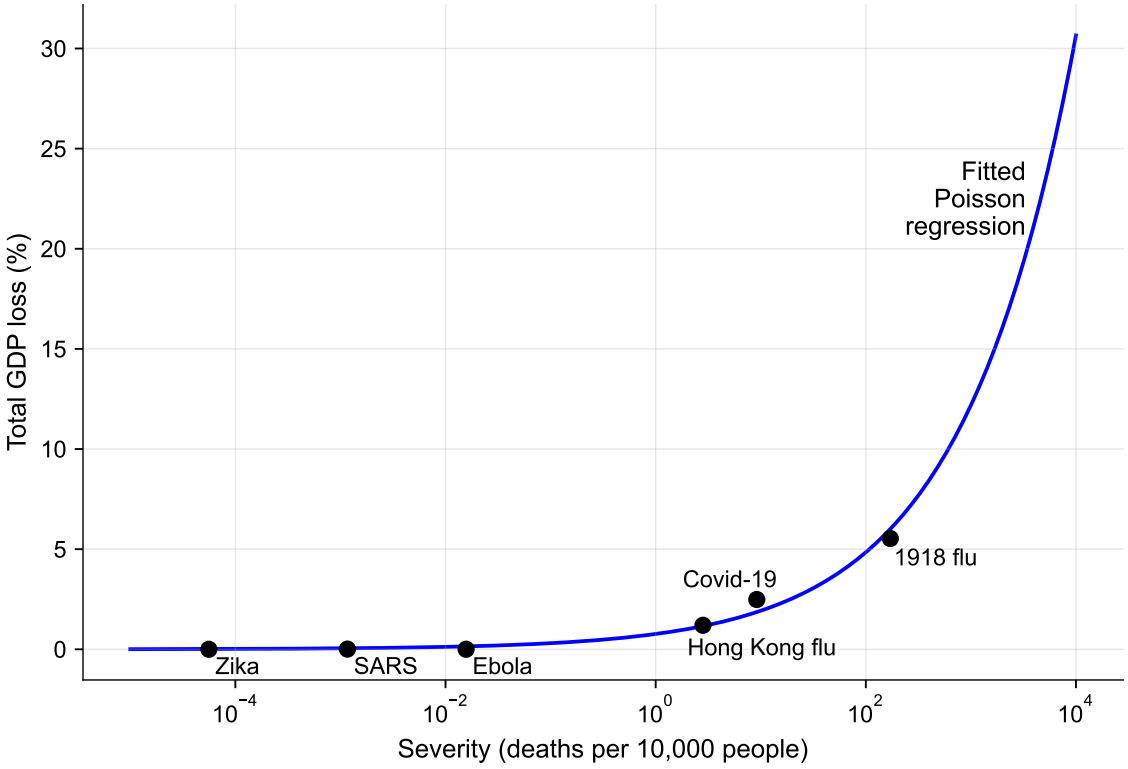


Figure A4: Estimated relationship between pandemic intensity and annual GDP loss. Each point represents one of the historical pandemics from Table A2. Curve is the fitted line from the Poisson regression reported in Table A3.

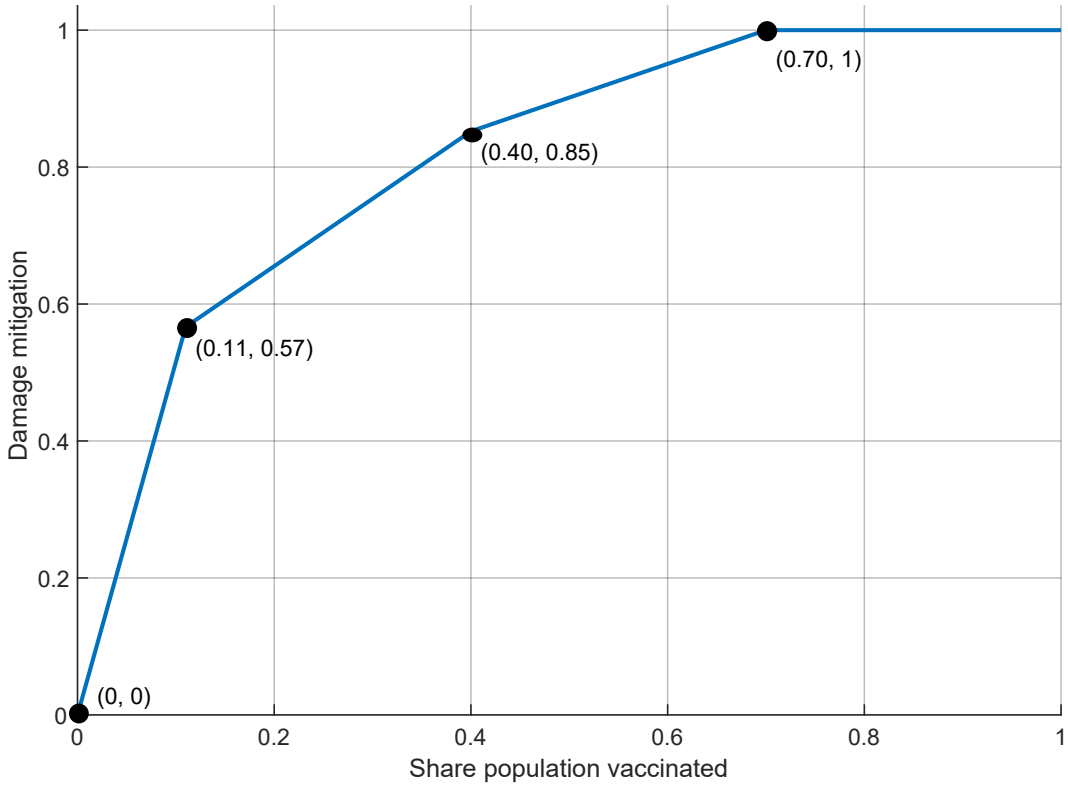


Figure A5: Pandemic harm mitigation function. Graph of the function H mapping vaccine coverage to the fraction of social losses averted. Its concavity reflects diminishing marginal returns to the vaccination of populations that face progressively lower morbidity risk or likelihood of onward disease transmission. It achieves total herd immunity at coverage $\bar{Z} = 70\%$. The function is taken from prior evaluations of programs accelerating vaccine availability during pandemics.[38,40]

Disease	Viral family	Airborne transmission?	Prototype vaccine in status quo	Source
Chikungunya	Togaviridae	No	Yes	[71]
Coronaviruses	Coronaviridae	Yes	Yes	[72]
CCHF	Nairoviridae	No	No	[73]
Ebola	Filoviridae	No	Yes	[74]
Influenza	Orthomyxoviridae	Yes	Yes	[75]
Lassa	Arenaviridae	No	No	[76]
Nipah	Paramyxoviridae	No	Yes	[77]
Rift Valley Fever	Phenuiviridae	No	No	[78]
Zika	Flaviviridae	No	Yes	[79]
Other known virus	N/a	N/a	No	N/a
Unknown virus	N/a	N/a	No	N/a

Table A4: Viruses with pandemic potential in model. Entries are a combination of list of priority pathogens from CEPI[8] and list of priority diseases from the WHO.[9] To this combined list, we added influenza and the residual category of "other known virus." Disease X from those lists has been relabeled "unknown virus." The first nine rows are specific viral families in alphabetical order. The last two rows list broad residual categories included as possible pandemic causes. CCHF is an abbreviation for Crimean–Congo hemorrhagic fever.

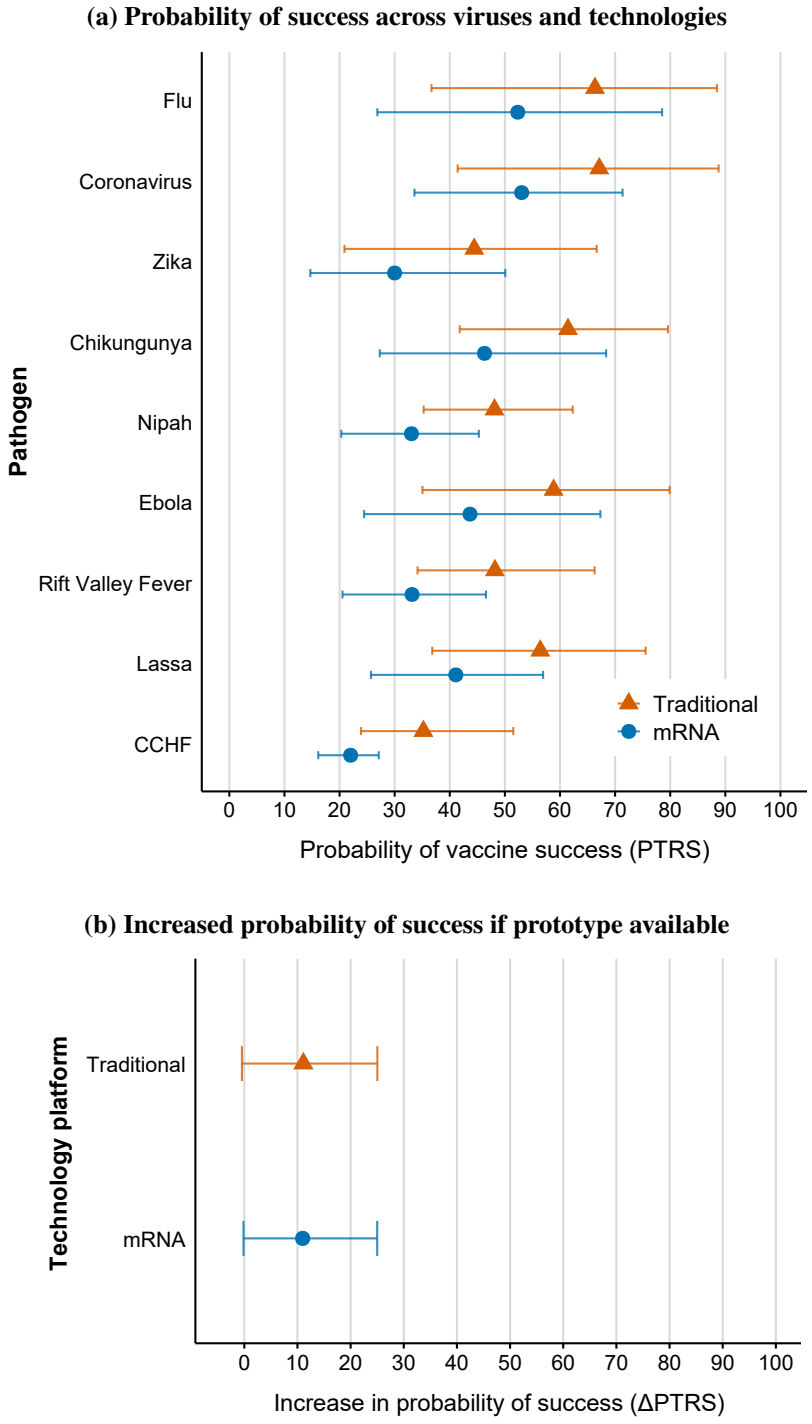


Figure A6: Estimates from expert survey of probability of vaccine success and increase if prototype available. Top panel plots probability of technical and regulatory success (PTRS) for vaccine candidates across viruses and technological platforms. Estimates represent the expected subjective probability of success, based on modeling expert responses using a random-effects beta regression, as detailed in equations (A18)–(A22). Bottom panel reports increase in PTRS if prototype vaccine available, estimated based on responses to the expert survey using the formula in equation (A35). Whiskers denote 95% confidence intervals estimated via a 1000-sample bootstrap, clustered at the respondent level.

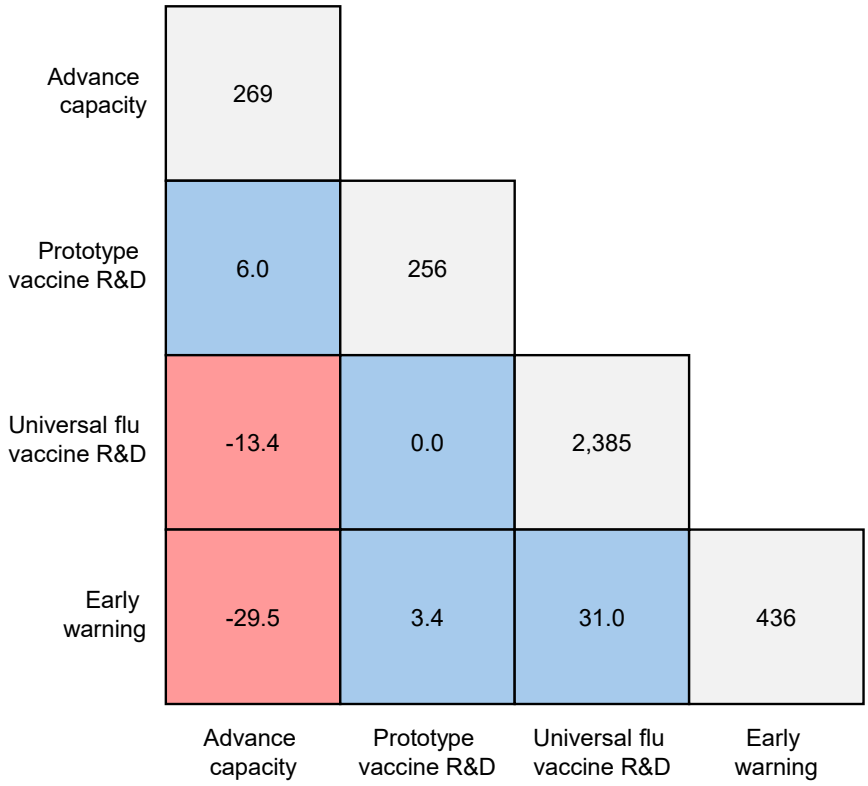


Figure A7: Complementarity between preparedness investments. Gray boxes along diagonal repeat expected NPV estimates from standalone investments provided in Table 2. Colored boxes below the diagonal report the positive or negative synergies between the investments in the row and column headings. The number in the box equals the difference between the two investments' combined NPV and the sum of their standalone NPVs. Box is colored blue if the combined NPV is weakly greater than the sum of their standalone NPVs, indicating the investments are (weak) complements. Box is colored red if the combined NPV is less than the sum of their standalone NPVs, indicating the investments are partial substitutes. Monetary values are discounted.

Parameter	Symbol	Value	Source
Simulation parameters			
• Number of simulations	I	1 mil.	Author choice
• Number of periods simulated	T	200 yrs.	Author choice
Pandemic-risk parameters			
• Severity ceiling (deaths / 10,000)	$\bar{s}$	200	Author choice
• Severity floor (deaths / 10,000)	$\underline{s}$	0.01	Authors, using [1]
• Response threshold (deaths / 10,000)	$\hat{s}$	6.14	Author estimation
Social-value parameters			
• Social discount rate	r	4%	[80]
• GDP growth rate	g	1.6%	[12]
• Base year population (2024)	N_0	8.14 bil.	[53]
• Base year GDP per capita (2024)	Y_0	$13,673	[81]
• Value of statistical life	V	$1.3 mil.	[11]
• Harm mitigated by vaccine	γ	50%	[40]
Capacity parameters			
• Max capacity (annual courses)	$\bar{Q}$	24.4 bil.	Author choice
• Annual capacity maintenance cost	δ_q	9%	Expert survey
• Surge capacity kept after pandemic	ϕ_q	50%	Expert survey
• Advance capacity tailoring fraction		33%	Author choice
• Max surge capacity displacement	θ_q	25%	Author choice
• Elasticity of surge capacity supply	ϵ	1	[2]
• Surge capacity cost kink	$\hat{q}$	100 mil.	[2]
• Capacity immediately usable		30%	[2]
• Repurposable capacity		40%	[2]
• Advance capacity build period		30 yrs.	Author choice
• Capacity unit cost (per annual course)	k_ℓ		
◦ Traditional platform		$3.00	Expert survey
◦ mRNA platform		$1.50	Expert survey
• Production unit cost (per course)	c_ℓ		
◦ Traditional platform		$17	Linksbridge
◦ mRNA platform		$34	Linksbridge
• Status quo capacity (annual courses)	Q_ℓ^b		
◦ Traditional platform		5.9 bil.	Expert survey
◦ mRNA platform		2.9 bil.	Expert survey
• Surge capacity (annual courses)	Q_ℓ^s		
◦ Traditional platform		2.25 bil.	[2]
◦ mRNA platform		2.25 bil.	[2]
Vaccine development parameters			
• Response R&D spending	K^r	$6.86 bil.	[47]
• Max vaccination rate	$\bar{Z}$	70%	Author choice
• Surge spending sunk on false positive	ϕ_f	33%	Author choice
• Advance R&D cost per pathogen		$1.06 bil.	Authors, based on [48]
• Time until advance R&D complete		11 yrs.	[48]
• Universal flu vaccine relative effectiveness	e_u	0.5	Authors
• Universal flu vaccine cost multiplier	κ_u	2	Authors
• Probability prototype successful		90%	[48]
• Probability universal flu vaccine successful		90%	[48]
• Universal flu vaccine peacetime uptake	$U_t(0)$	10%	[52], [53]
• Response R&D timeline (months)	m^a		
◦ Without prototype		18	[25]
◦ With prototype		11	[25]
• Time to repurpose capacity	m_ℓ^r		
◦ Traditional platform		8 mo.	Expert survey
◦ mRNA platform		2 mo.	Expert survey
Biosurveillance parameters			
• Improved biosurveillance annual cost		$490 million	Authors, based on [59]
• Months to pandemic detection			
◦ Status quo		3	Author choice
◦ Early warning		2	Authors, based on [54]

Table A5: Baseline values of model parameters. Where possible, parameter values are based on established literature, expert survey responses, and author estimation. Remaining parameters were based on authors' judgment of reasonable values, meant to keep estimates of risk or program benefits conservative.

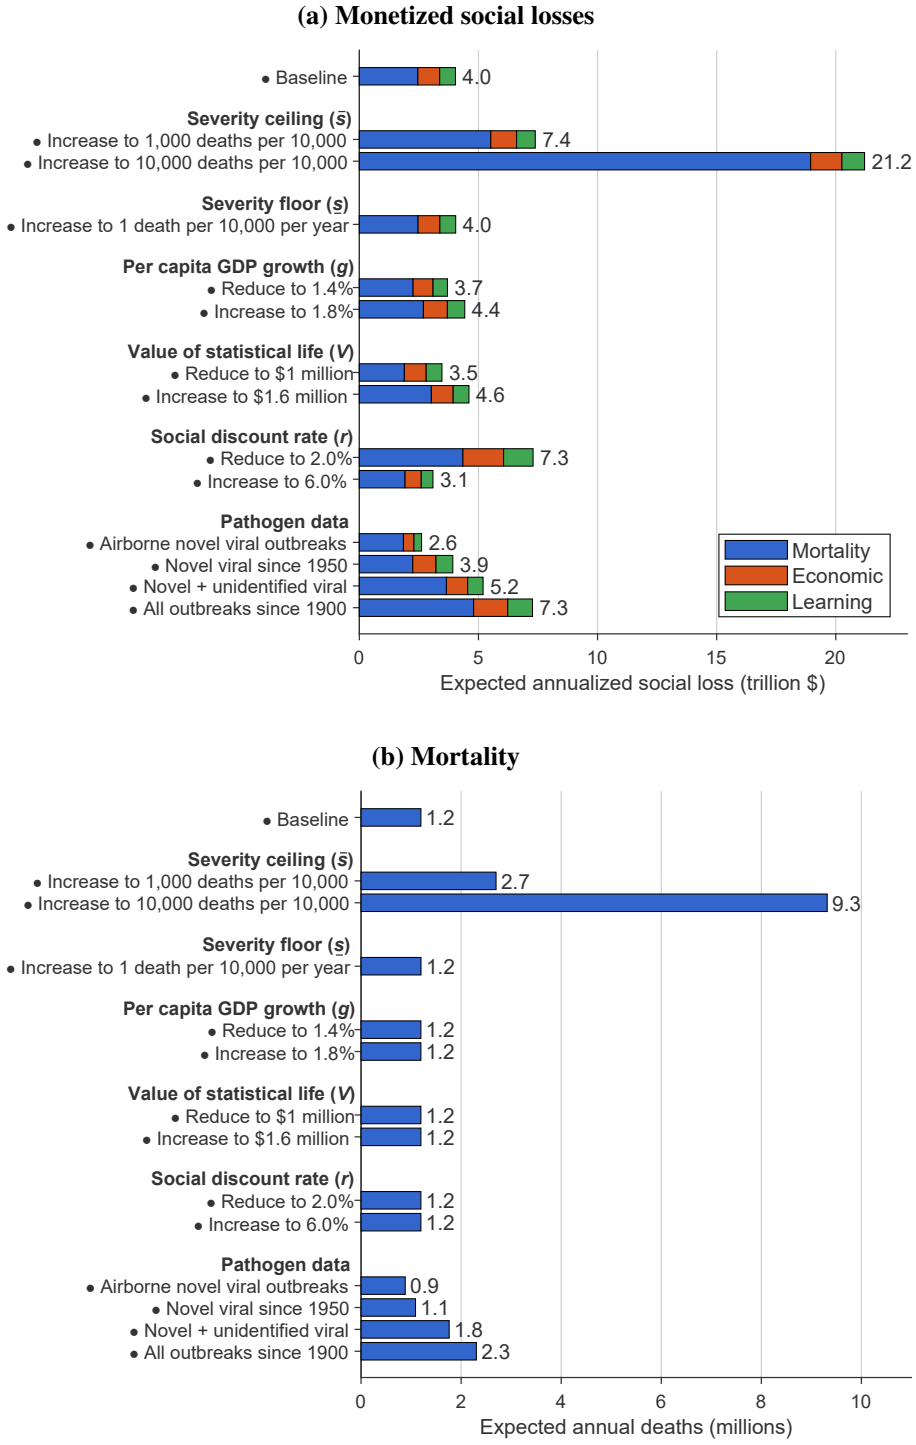


Figure A8: Expected annual pandemic harm in counterfactual absence of vaccine mitigation. Results are means of a million simulation runs for each scenario. The first bar in each panel provides results for scenario with baseline parameters: severity ceiling $\bar{s}$ = 200 deaths per 10,000, severity floor $\underline{s}$ = 0.01 deaths per 10,000, per capita GDP growth rate g = 1.6%, value of statistical life V = \$1.3 million, and social discount rate r = 4%. Bars below the first assess sensitivity to changes in individual parameters from the baseline. Number at end of bar is the total for stacked bar. Panel (a) shows results for monetized social losses, aggregating mortality, output, and learning losses, discounted using social discount rate r, and annualized. Panel (b) shows results for deaths, which have been annualized but are undiscounted.

Parameter	Baseline parameter		Parameter reduction		Parameter increase	
	Baseline parameter value	NPV estimate (trillion $)	Reduced parameter value	NPV estimate (trillion $)	Increased parameter value	NPV estimate (trillion $)
General parameters						
• Harm mitigated by vaccine	50%	17.6	30%	10.5	70%	24.7
• Severity ceiling (deaths / 10k)	200	17.6	100	13.8	1000	31.9
• Pathogen data	Default	17.6	—	—	All 1900+	23.3
Response parameters						
• Months to pandemic detection	3 mo.	17.6	1 mo.	18.7	5 mo.	16.6
• Response threshold (deaths / 10k)	6.1	17.6	3.1	18.3	12.3	16.4
Capacity parameters						
• Max capacity (annual courses)	24.4 bil.	17.6	16.3 bil.	17.5	32.6 bil.	17.6
• Surge capacity (annual courses)						
◦ mRNA platform	2.25 bil.	17.6	1.00 bil.	17.4	3.00 bil.	17.7
◦ Traditional platform	2.25 bil.	17.6	1.00 bil.	17.4	3.00 bil.	17.6
• Surge capacity kept after pandemic	50%	17.6	0%	17.3	75%	17.6
• Surge capacity cost kink	$100 mil.	17.6	$50 mil.	17.5	$200 mil.	17.6
• Elasticity of surge capacity supply	1	17.6	0.75	17.6	1.25	17.5
• Annual capacity maintenance cost	9%	17.6	4%	17.6	12%	17.6
• Capacity cost (per annual course)						
◦ mRNA platform	$1.5	17.6	$0.75	17.6	$5	17.5
◦ Traditional platform	$3	17.6	$1.50	17.6	$6	17.5
• Production unit cost (per course)						
◦ mRNA platform	$34	17.6	$20	17.6	$40	17.6
◦ Traditional platform	$17	17.6	$12	17.6	$30	17.6
• Successful at-risk capacity						
◦ mRNA platform	30%	17.6	15%	17.4	50%	17.8
◦ Traditional platform	30%	17.6	15%	17.0	50%	18.0
• Repurposable at-risk capacity						
◦ mRNA platform	40%	17.6	20%	17.4	60%	17.7
◦ Traditional platform	40%	17.6	20%	17.5	60%	17.6
• Repurposing delay						
◦ mRNA platform	2 mo.	17.6	1 mo.	17.7	4 mo.	17.5
◦ Traditional platform	8 mo.	17.6	4 mo.	17.8	12 mo.	17.5
Response R&D parameters						
• Response R&D spending	$6.86 bil.	17.6	$3.43 bil.	17.6	$9.00 bil.	17.6
• Response R&D success probs.	Mean	17.6	25th pct.	15.8	75th pct.	19.2
• Response R&D timeline						
◦ With prototype	11 mo.	17.6	6 mo.	20.3	16 mo.	15.5
◦ Without prototype	18 mo.	17.6	13 mo.	17.9	28 mo.	17.1

Table A6: Sensitivity of NPV estimates for status quo vaccine response to parameter changes. "Successful at-risk capacity" is the proportion dedicated to successful candidates that can be produced immediately after approval. "Repurposable at-risk capacity" is the proportion dedicated to unsuccessful candidates that can be repurposed with a lag of m_ℓ^r months. For response R&D success probabilities, the baseline parameter values are the mean PTRS estimates for each pathogen (Panel (a), Figure A6), while the reduced and increased parameter values are the 25th and 75th percentile values generated by bootstrapping from the fitted random-effects beta regression.

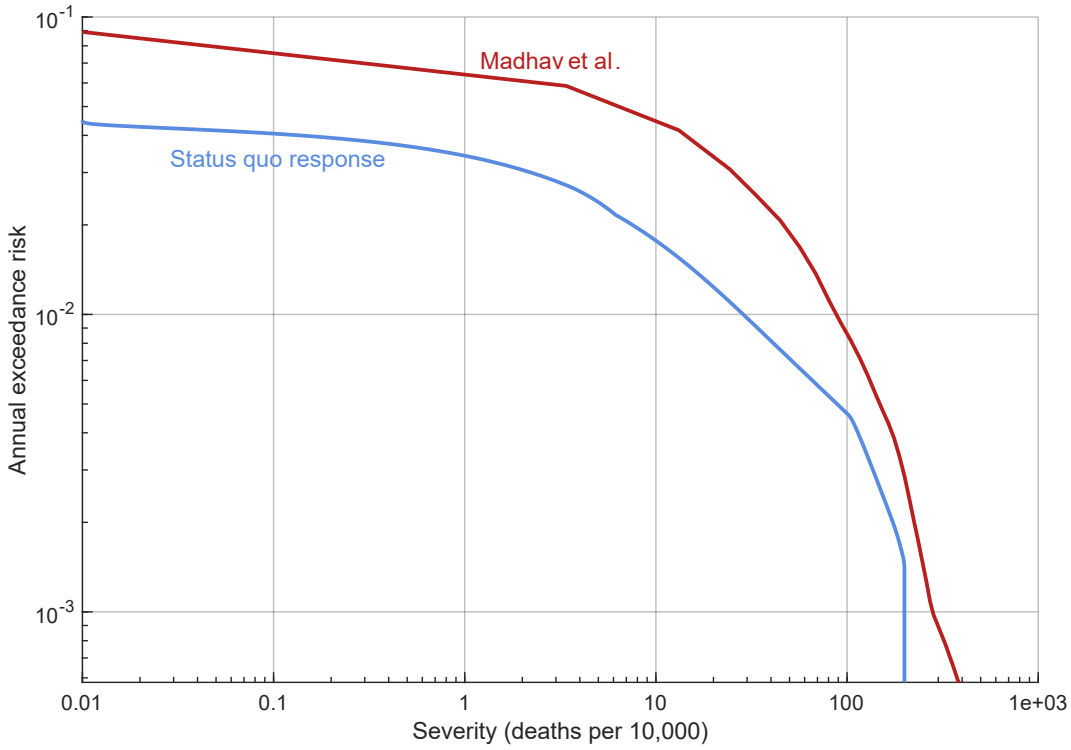


Figure A9: Comparison of exceedance risk estimates with an epidemiological model forecast. The red curve was obtained by digitizing Figure 2.2 from Madhav et al.[82] and replotting with a log scale on the horizontal axis. The exceedance function in that study was derived from a structural epidemiological model of pandemics caused by respiratory pathogens reflecting vaccine response and other countermeasures. The blue curve reproduces our estimate of the exceedance curve from Figure 6 that reflects the status quo vaccine response.

Parameter name	Baseline parameter		Parameter reduction		Parameter increase	
	Baseline parameter value	NPV estimate (billion $)	Reduced parameter value	NPV estimate (billion $)	Increased parameter value	NPV estimate (billion $)
Advance capacity						
• Build period	30 yrs.	269	20 yrs.	312	40 yrs.	233
• Max. surge capacity displacement	25%	269	10%	304	50%	201
• Tailoring fraction	33%	269	10%	264	50%	273
• Max. capacity (annual courses)	24.4 bil.	269	16.3 bil.	139	32.6 bil.	277
• Surge capacity kept after pandemic	50%	269	0%	412	75%	232
• Elasticity of surge capacity supply	1	269	0.75	255	1.25	297
• Capacity unit cost (per annual course)						
◦ mRNA platform	$1.5	269	$0.75	269	$5	267
◦ Traditional platform	$3	269	$1.5	265	$6	277
• Surge capacity (annual courses)						
◦ mRNA platform	2.25 bil.	269	1.0 bil.	398	3.0 bil.	223
◦ Traditional platform	2.25 bil.	269	1.0 bil.	335	3.0 bil.	250
Prototype vaccine R&D						
• R&D cost per pathogen	$1.06 bil.	256	$530 mil.	257	$2.12 bil.	253
• R&D timeline	11 yrs.	256	6 yrs.	287	16 yrs.	226
• R&D success prob.	90%	256	50%	141	100%	285
• Response R&D success prob. increase	Mean	256	25th pct.	205	75th pct.	297
• Response R&D timeline						
◦ With prototype	11 mo.	256	6 mo.	405	16 mo.	144
◦ Without prototype	18 mo.	256	13 mo.	172	28 mo.	377
Universal flu vaccine						
• Efficacy multiplier	50%	2,385	25%	1,773	100%	3,041
• R&D cost multiplier	2	2,385	1	2,386	3	2,384
• Peacetime uptake	10%	2,385	0%	1,811	20%	2,789
• R&D success prob.	90%	2,385	50%	1,321	100%	2,649
• R&D lag	11 yrs.	2,385	6 yrs.	2,687	16 yrs.	2,122
Improved early warning system						
• Annual cost	$490 mil.	436	$245 mil.	442	$980 mil.	424
• System performance						
◦ Precision	45%	436	20%	542	70%	394
◦ Sensitivity	70%	436	40%	272	90%	545
◦ Detection speedup	1 mo.	436	0 mo.	53	2 mo.	841
• Response cost on false alarm	33%	436	10%	398	50%	454
• Surge capacity kept after pandemic	50%	436	0%	357	75%	448

Table A7: Sensitivity of advance investment program net present value estimates. NPV estimates represent the estimated mean difference in discounted net surplus generated by the advance investment and the status quo programs under the same parameter configuration across one million simulations. For "Response R&D success prob. increase," the baseline parameter value is the mean estimated increase in the probability of technical and regulatory success (PTRS) attributable to prototype availability as described by equation (A35). The reduced and increased values for this parameter are the 25th and 75th percentiles of the confidence-interval distribution. Parameters labeled "R&D" are specific to the program indicated in the panel heading.

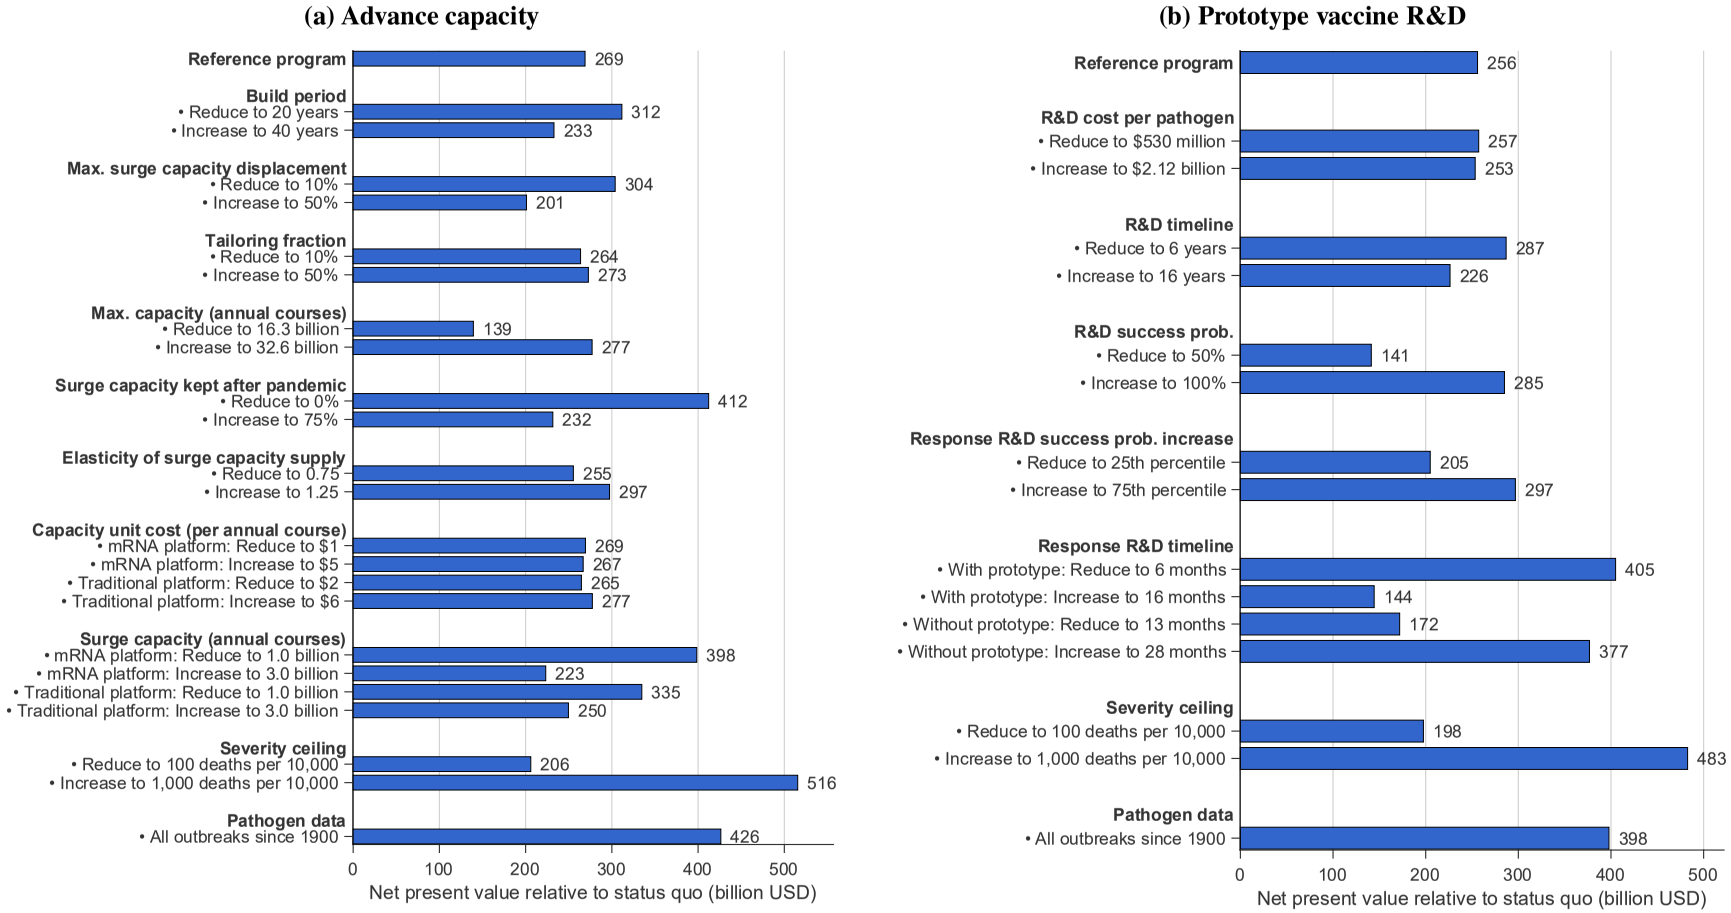


Figure A10: Continues on next page.

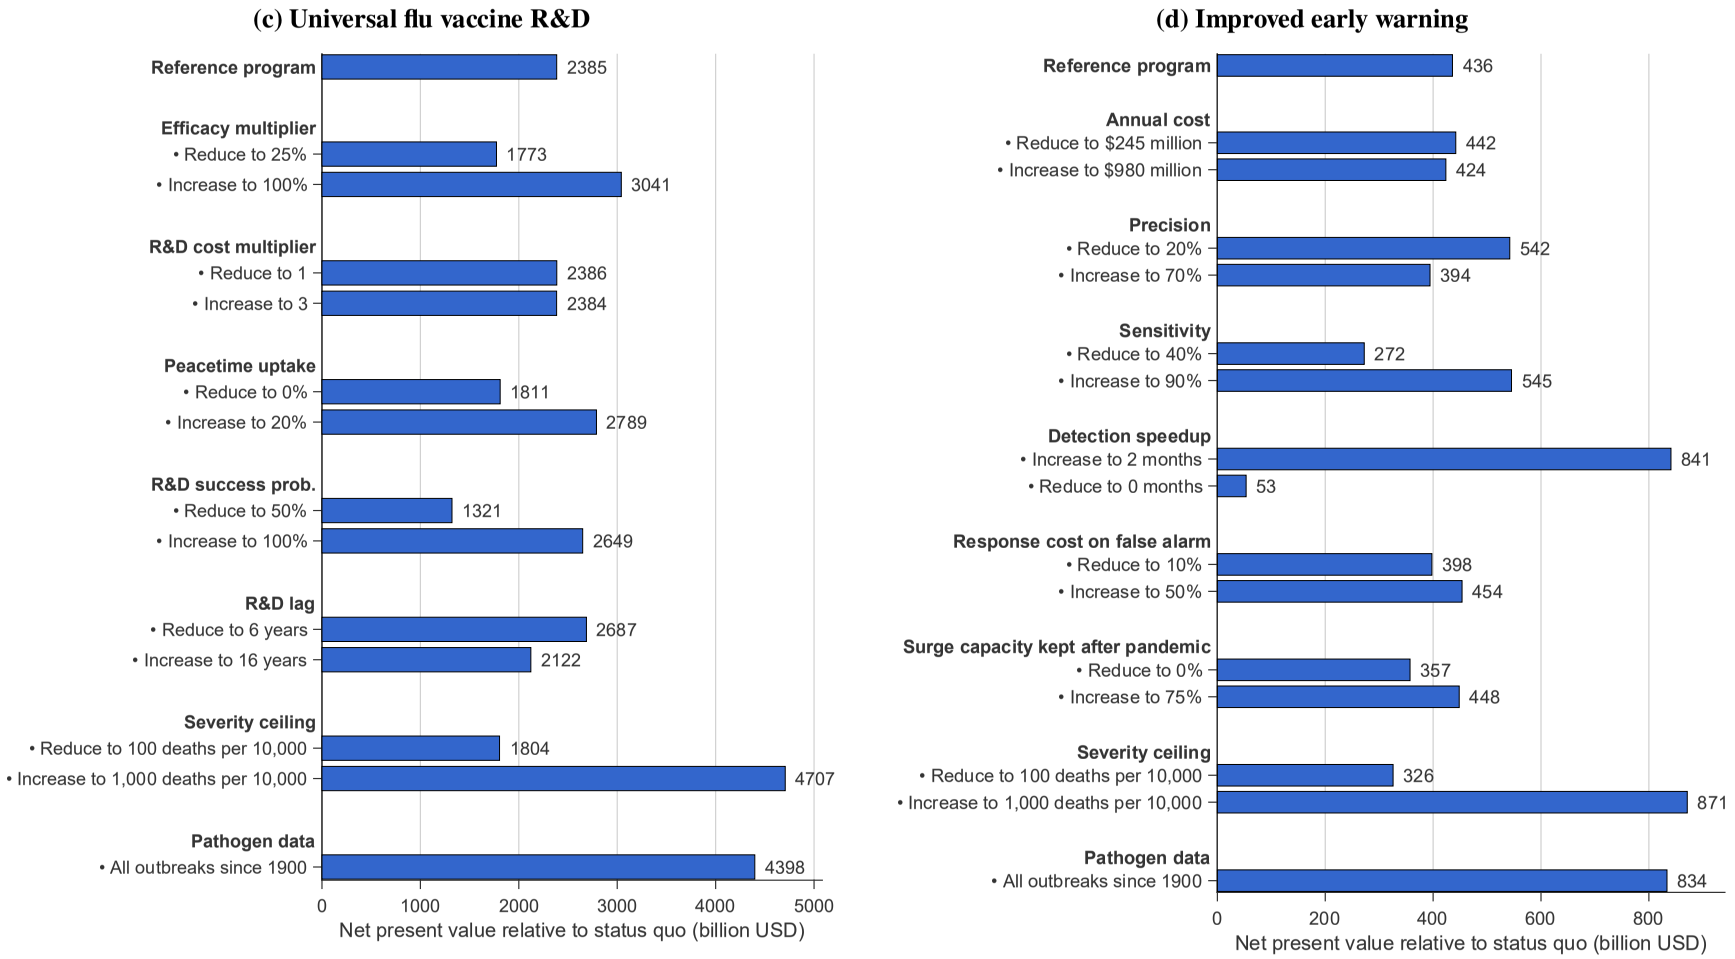


Figure A10: Sensitivity of advance investment program net present value estimates. NPV estimates represent the estimated mean difference in discounted net surplus generated by the advance investment and the status quo programs under the same parameter configuration across one million simulations. Each bar represents one parameter change. Besides reprising results from Table A7, here in graphical form, this figure reports additional sensitivity analyses for the severity ceiling and pathogen data underlying the pandemic forecast.

Program	Present discounted value ($ billion)			Benefit-cost ratio	Lives saved (million)
	Costs	Benefits	NPV		
Advance capacity					
• 50% of standard program	9.4	126	116	13	0.2
• Standard program	19	288	269	15	0.6
• 200% of standard program	38	452	415	12	0.8
Prototype vaccine R&D					
• R&D for top priority pathogen	0.8	123	122	146	0.4
• R&D for top 2 pathogens	1.7	200	198	118	0.6
• R&D for top 3 pathogens	2.5	258	256	102	0.7

Table A8: Estimates of costs and benefits of advance-investment programs of varying scales. Entries are means over one million simulation runs for per program. The top rows report results for advance-capacity programs, varying capacity as a percentage of 7.5 billion annual courses supplied under the standard program design. The bottom rows report results for prototype vaccine R&D programs investing in the top one, two, or three priority pathogens currently lacking a prototype vaccine, ranked by outbreak likelihood from Figure A3. Present discounted values are cumulative values over the 200-year simulation run, discounted at the social discount rate r, and averaged over the million simulations.

Appendix references

[1] Marani M, Katul GG, Pan WK, Parolari AJ. Intensity and frequency of extreme novel epidemics. Proceedings of the National Academy of Sciences. 2021 Aug;118(35):e2105482118.

[2] Glennerster R, Snyder CM, Tan BJ. Calculating the costs and benefits of advance preparations for future pandemics. IMF Economic Review. 2023 Sep;71(3):611-48.

[3] Rubin DB. Inference and missing data. Biometrika. 1976 Dec;63:581-92.

[4] Kim H, Webster RG, Webby RJ. Influenza virus: Dealing with a drifting and shifting pathogen. Viral Immunology. 2018 Mar;31(2):174-83.

[5] Pickands J. Statistical inference using extreme order statistics. Annals of Statistics. 1975;3:119-31.

[6] Davison AC, Smith RL. Models for exceedances over high thresholds. Journal of the Royal Statistical Society B. 1990;52:393-442.

[7] Li G, Quaade S, Snyder CM. Forecasting Pandemic Tail Risk: Model Choice, Specification, and Inference. Dartmouth College; 2025.

[8] Coalition for Epidemic Preparedness Innovation. Priority pathogens; no date. Available from: https://cepi.net/priority-pathogens.

[9] World Health Organization. Prioritizing diseases for research and development in emergency contexts; no date. Available from: https://www.who.int/activities/prioritizing-diseases-for-research-and-development-in-emergency-contexts.

[10] United Nations. World population prospects 2024. New York; 2024. Available from: https://www.un.org/development/desa/pd/world-population-prospects-2024.

[11] Sweis NJ. Revisiting the value of a statistical life: An international approach during COVID-19. Risk Management. 2022 Sep;24(3):259-72.

[12] Guillemette Y, Turner D. The long game: Fiscal outlooks to 2060 underline need for structural reform. OECD Economic Policy Paper number 29. OECD; 2021. Available from: https://doi.org/10.1787/a112307e-en.

[13] Wooldridge JM. Econometric Analysis of Cross Section and Panel Data. 2nd ed. Cambridge, MA: MIT Press; 2010.

[14] Jinjarak Y, Noy I, Ta Q. Pandemics and economic growth: Evidence from the 1968 H3N2 influenza. Economics of Disasters and Climate Change. 2022;6(1):73-93.

[15] Azevedo JP, Hasan A, Goldemberg D, Geven K, Iqbal SA. Simulating the potential impacts of COVID-19 school closures on schooling and learning outcomes: A set of global estimates. World Bank Research Observer. 2021 Mar;36(1):1-40.

[16] International Monetary Fund. World Economic Outlook update: Rising caseloads, a disrupted recovery, and higher inflation. International Monetary Fund; 2022. Available from: https://www.imf.org/en/publications/weo/issues/2022/01/25/world-economic-outlook-update-january-2022.

[17] Pekar J, Worobey M, Moshiri N, Scheffler K, Wertheim JO. Timing the SARS-CoV-2 index case in Hubei province. Science. 2021 Apr;372(6540):412-7.

[18] Roberts DL, Rossman JS, Jarić I. Dating first cases of COVID-19. PLOS Pathogens. 2021 Jun;17(6):e1009620.

[19] World Health Organization. Pneumonia of unknown cause—China. Disease Outbreak News. Geneva: World Health Organization; 2020. Available from: https://www.who.int/emergencies/disease-outbreak-news/item/2020-DON229.

[20] Hsieh YH, Ma S, Velasco Hernandez JX, Lee VJ, Lim WY. Early outbreak of 2009 influenza A (H1N1) in Mexico prior to identification of pH1N1 virus. PLOS ONE. 2011 Aug;6(8):e23853.

[21] Cheng VCC, To KKW, Tse H, Hung IFN, Yuen KY. Two years after pandemic influenza A/2009/H1N1: What have we learned? Clinical Microbiology Reviews. 2012 Apr;25(2):223-63.

[22] Centers for Disease Control and Prevention. Swine influenza A (H1N1) infection in two children—Southern California, March–April 2009. Morbidity and Mortality Weekly Report. 2009;58(15):400-2.

[23] WHO Ebola Response Team. Ebola virus disease in West Africa—The first 9 months of the epidemic and forward projections. New England Journal of Medicine. 2014 Oct;371(16):1481-95.

[24] Neustadt RE, Fineberg HV. The Swine Flu Affair: Decision-Making on a Slippery Disease. Washington, DC: National Academies Press; 1978.

[25] Jeyapragasan G, Graabak J, Luby S, Esvelt KM. Modeling the benefits of virus discovery and pandemic virus identification. bioRxiv. 2024.

[26] Pronker ES, Weenen TC, Commandeur H, Claassen EHJHM, Osterhaus ADME. Risk in vaccine research and development quantified. PLOS ONE. 2013 Mar;8(3):e57755.

[27] MacPherson A, Hutchinson N, Schneider O, Oliviero E, Feldhake E, Ouimet C, et al. Probability of success and timelines for the development of vaccines for emerging and reemerged viral infectious diseases. Annals of Internal Medicine. 2021 Mar;174(3):326-34.

[28] Dattani S, Haria S, Glennerster R. Why we didn't get a malaria vaccine sooner. Works In Progress. 2023 Sep;(issue 12). Available from: https://worksinprogress.co/issue/why-we-didnt-get-a-malaria-vaccine-sooner.

[29] Corbett KS, Edwards D, Leist SR, Abiona OM, Boyoglu-Barnum S, Gillespie RA, et al. SARS-CoV-2 mRNA vaccine development enabled by prototype pathogen preparedness. bioRxiv. 2020 Jun.

[30] Graham BS. Rapid COVID-19 vaccine development. Science. 2020 May;368(6494):945-6.

[31] Stewart MB. On least squares estimation when the dependent variable is grouped. The review of economic studies. 1983;50(4):737-53.

[32] Rigby RA, Stasinopoulos DM. Generalized additive models for location, scale and shape. Journal of the Royal Statistical Society Series C: Applied Statistics. 2005;54(3):507-54.

[33] Fortner A, Schumacher D. First COVID-19 vaccines receiving the US FDA and EMA Emergency Use Authorization. Discoveries;9(1):e122.

[34] Athey S, Castillo JC, Chaudhuri E, Kremer M, Simoes Gomes A, Snyder CM. Expanding capacity for vaccines against Covid-19 and future pandemics: a review of economic issues. Oxford Review of Economic Policy. 2022 Dec;38(4):742-70.

[35] Buntz B. Pfizer to close Durham and Morrisville, North Carolina facilities amid major restructuring effort. Pharmaceutical Processing World. 2023. Available from: https://www.pharmaceuticalprocessingworld.com/pfizer-to-close-durham-and-morrisville-north-carolina-facilities-amid-major-restructuring-effort.

[36] Taylor NP. Moderna halts plans for Kenyan mRNA vaccine plant after taking $1b financial hit. BioSpace. 2024. Available from: https://www.biospace.com/moderna-halts-plans-for-kenyan-mrna-vaccine-plant-after-taking-1b-financial-hit.

[37] Fine PEM. Herd immunity: History, theory, practice. Epidemiologic Reviews. 1993;15(2):265-302.

[38] Ahuja A, Athey S, Baker A, Budish E, Castillo JC, Glennerster R, et al. Preparing for a pandemic: Accelerating vaccine availability. American Economic Association Papers and Proceedings. 2021;111:331-5.

[39] Lazarus JV, Wyka K, White TM, Picchio CA, Rabin K, Ratzan SC, et al. Revisiting COVID-19 vaccine hesitancy around the world using data from 23 countries in 2021. Nature Communications. 2022;13(1):3801.

[40] Castillo JC, Ahuja A, Athey S, Baker A, Budish E, Chipty T, et al. Market design to accelerate COVID-19 vaccine supply. Science. 2021 Mar;371(6534):1107-9.

[41] World Bank. Population ages 65 and above (Available from: https://data.worldbank.org/indicator/SP.POP.65UP.TO.ZS.

[42] Boniol M, Kunjumen T, Nair TS, Siyam A, Campbell J, Diallo K. The global health workforce stock and distribution in 2020 and 2030: A threat to equity and "universal" health coverage? BMJ global health. 2022;7(6).

[43] Bubar KM, Reinholt K, Kissler SM, Lipsitch M, Cobey S, Grad YH, et al. Model-informed COVID-19 vaccine prioritization strategies by age and serostatus. Science. 2021;371(6352):916-21.

[44] World Bank. GDP per capita (current US$); 2026. Available from: https://data.worldbank.org/indicator/NY.GDP.PCAP.CD.

[45] Davies NG, Kucharski AJ, Eggo RM, Gimma A, Edmunds WJ, Centre for the Mathematical Modelling of Infectious Diseases COVID-19 working group. Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: A modelling study. Lancet Public Health. 2020;5(7):E375-85.

[46] Ward J, Gressani O, Kim S, Hens N, Edmunds WJ. The epidemiology of pathogens with pandemic potential: A review of key parameters and clustering analysis. Epidemics. 2025:100882.

[47] Global Health Centre. COVID-19 vaccine R&D investments. Graduate Institute of International and Development Studies; 2021. Available from: https://www.knowledgeportalia.org/covid-19-vaccine-r-d-funding.

[48] Sertkaya A, Beleche T, Jessup A. New estimates of the cost of preventive vaccine development and potential implications from the COVID-19 pandemic. Office of the Assistant Secretary for Planning and Evaluation (ASPE) brief. Office of the Assistant Secretary for Planning and Evaluation; 2024. Available from: https://www.ncbi.nlm.nih.gov/books/NBK611605.

[49] Erbelding EJ, Post DJ, Stemmy EJ, Roberts PC, Deckhut Augustine A, Ferguson S, et al. A universal influenza vaccine: The strategic plan for the National Institute of Allergy and Infectious Diseases. Journal of Infectious Diseases. 2018;218:347-54.

[50] Krammer F. The human antibody response to influenza A virus infection and vaccination. Nature Reviews Immunology. 2019;19:383-97.

[51] Moore KA, Leighton T, Ostrowsky L, Anderson CJ, Danila RN, Ulrich AK, et al. A research and development (R&D) roadmap for broadly protective coronavirus vaccines: A pandemic preparedness strategy. Vaccine. 2023;41:2101-12.

[52] World Health Organization. WHO seasonal influenza vaccine global market study, January 2024. WHO technical document. World Health Organization; 2024. Available from: https://www.who.int/publications/m/item/who-seasonal-influenza-vaccine-global-market-study-january-2024.

[53] World Bank. Population, total;. Available from: https://data.worldbank.org/indicator/SP.POP.TOTL.

[54] Liu AB, Lee D, Jalihal AP, Hanage WP, Springer M. Quantitatively assessing early detection strategies for mitigating COVID-19 and future pandemics. Nature Communications. 2023 Dec;14(1):8479.

[55] Wright RJ, Comeau AM, Langille MGI. From defaults to databases: Parameter and database choice dramatically impact the performance of metagenomic taxonomic classification tools. Microbial Genomics. 2023;9(3):000949.

[56] Fan V, Smitham E, Regan L, Gautam P, Norheim O, Guzman J, et al. Strategic investment in surveillance for pandemic preparedness: Rapid review and roundtable discussion. Center for Global Development policy paper no. 298. Center for Global Development; 2023. Available from: https://www.cgdev.org/publication/strategic-investment-surveillance-pandemic-preparedness-rapid-review-and-roundtable.

[57] Dobson AP, Pimm SL, Hannah L, Kaufman L, Ahumada JA, Ando AW, et al. Ecology and economics for pandemic prevention. Science. 2020 Jul;369(6502):379-81.

[58] Craven M, Sabow A, Van der Veken L, Wilson M. Not the last pandemic: Investing now to reimagine public-health systems. McKinsey & Company; 2021. Available from: https://www.mckinsey.com/industries/public-sector/our-insights/not-the-last-pandemic-investing-now-to-reimagine-public-health-systems.

[59] Grimm S, Justen L, Kaufman J. Scaling US pathogen detection [Text]. Nucleic Acid Observatory. 2025 Sep. Available from: https://naobservatory.org/blog/biothreat_radar/.

[60] Jones ER, van Vliet MTH, Qadir M, Bierkens MFP. Country-level and gridded wastewater production, collection, treatment and re-use. PANGAEA. 2020.

[61] Mathieu E, Ritchie H, Rodés-Guirao L, Appel C, Giattino C, Hasell J, et al. COVID-19 pandemic. Our World in Data. 2020. Available from: https://ourworldindata.org/coronavirus.

[62] Ioannidis JPA, Pezzulo AM, Cristiano A. Global Estimates of Lives and Life-Years Saved by COVID-19 Vaccination During 2020-2024. JAMA Health Forum. 2025;6(7).

[63] Vanderslott S, Dattani S, Spooner F, Roser M. Vaccination. Our World in Data. 2022 Oct. Available from: https://ourworldindata.org/vaccination.

[64] Barro RJ, Ursúa JF, Weng J. 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. National Bureau of Economic Research; 2020. Available from: https://www.nber.org/papers/w26866.

[65] Patterson KD, Pyle GF. The geography and mortality of the 1918 influenza pandemic. Bulletin of the History of Medicine. 1991;65(1):4-21.

[66] COVID-19 vaccine doses administered by manufacturer. Our World in Data. Available from: https://ourworldindata.org/grapher/covid-vaccine-doses-by-manufacturer.

[67] Keogh-Brown MR, Smith RD. The economic impact of SARS: How does the reality match the predictions? Health Policy. 2008 Oct;88(1):110-20.

[68] Huber C, Finelli L, Stevens W. The economic and social burden of the 2014 Ebola outbreak in West Africa. Journal of Infectious Diseases. 2018;218(Supplement 5):S698-704.

[69] Institute for Health Metrics and Evaluation. GBD results;. Available from: https://vizhub.healthdata.org/gbd-results.

[70] United Nations Development Programme. A socio-economic impact assessment of the Zika virus in Latin America and the Caribbean: With a focus on Brazil, Colombia and Suriname. United Nations Development Programme; 2017. Available from: https://www.undp.org/latin-america/publications/socio-economic-impact-assessment-zika-virus-latin-america-and-caribbean.

[71] U S Centers for Disease Control and Prevention. Chikungunya; 2024. Available from: https://www.cdc.gov/chikungunya/about/index.html.

[72] U S Centers for Disease Control and Prevention. Middle East Respiratory Syndrome (MERS); 2025. Available from: https://www.cdc.gov/mers/about/index.html.

[73] U S Centers for Disease Control and Prevention. Crimean-Congo hemorrhagic fever; 2025. Available from: https://www.cdc.gov/crimean-congo-hemorrhagic/about/index.html.

[74] U S Centers for Disease Control and Prevention. Ebola disease basics; 2025. Available from: https://www.cdc.gov/ebola/about/index.html.

[75] U S Centers for Disease Control and Prevention. Influenza vaccination coverage, adults 65 years and older, United States; 2025. Available from: https://www.cdc.gov/fluvaxview/dashboard/adults-65-years-and-older.html.

[76] U S Centers for Disease Control and Prevention. Lassa Fever; 2025. Available from: https://www.cdc.gov/lassa-fever/about/index.html.

[77] U S Centers for Disease Control and Prevention. Nipah virus; 2025. Available from: https://www.cdc.gov/nipah-virus/about/index.html.

[78] U S Centers for Disease Control and Prevention. About Rift Valley Fever (RVF); 2024. Available from: https://www.cdc.gov/rift-valley-fever/about/index.html.

[79] U S Centers for Disease Control and Prevention. Zika virus; 2025. Available from: https://www.cdc.gov/zika/index.html.

[80] Gollier C, Hammit JK. The long-run discount rate controversy. Annual Review of Resource Economics. 2014;6:273-95.

[81] World Bank. GDP (current US$)—United States; 2025. Available from: https://data.worldbank.org.

[82] Madhav NK, Oppenheim B, Stephenson N, Badker R, Jamison DT, Lam C, et al. Estimated future mortality from pathogens of epidemic and pandemic potential. In: Sharma S, Bertozi SM, Fan VY, Jamison DT, Norheim OF, Oshitani H, et al., editors. Disease Control Priorities Fourth Edition (DCP4), Volume 2 Investing in Pandemic Prevention, Preparedness, and Response. 3rd ed. Washington DC: International Bank for Reconstruction and Development, World Bank; 2026. .