Calculating the Costs and Benefits of Advance Preparations for the Next Pandemic

Rachel Glennerster
University of Chicago

Christopher M. Snyder
Dartmouth College

Brandon Joel Tan
International Monetary Fund

June 2022

Abstract
To date, the Covid-19 pandemic is estimated to have caused over 7 million deaths and reduced economic output by over $13 trillion, with losses continuing to mount. While vaccines were developed and deployed at unprecedented pace, pre-pandemic investments could have accelerated their widespread introduction and thus saved millions of lives and trillions of dollars. Combining estimates of the distribution of annual mortality risk from pandemics with estimates of the economic costs associated with pandemics of varying severities, we calculate that the expected cost of pandemics to the world is $329 billion every year. Spending $60 billion up front to expand production capacity for vaccines and supply-chain inputs and $2.2 billion every year thereafter would ensure there was sufficient production capacity to vaccinate 70% of the global population against a new virus within six months. This investment would generate expected net benefits of $28 billion annually. Investing in systems that reduce the time to gain regulatory approval for vaccines by one month (for example making ethical human challenge trials faster) would add net benefits of $900 million annually.

1 Introduction

The Covid-19 pandemic is estimated to have reduced economic output during 2020–24 by $13.8 trillion relative to pre-pandemic forecasts (International Monetary Fund 2022) while excess deaths during the pandemic are estimated at between 7 and 13 million (Economist, 2022). Loss of future productivity and earnings as a result of school closures is estimated at between $10 trillion and $17 trillion (World Bank 2022). While vaccines against Covid-19 were developed, approved, and distributed at record speed, sharply reducing both the economic and social losses, their distribution was highly uneven creating both injustice and more deaths and lost output that a more optimal distribution. One of the key lessons of the pandemic has been the huge value in accelerating the production and distribution of vaccines: if the production capacity to produce 1.5 billion courses of vaccine annually had been accelerated by just three months (from April 2021 to January 2021) this acceleration would have been worth $1.3 trillion (Castillo et al. 2021).

While Covid-19 is still with us, it is not too early to start preparations for the next pandemic which may take the form of a new virus, a drug resistant bacterium, or a deadly, vaccine-escaping, mutation of the existing, widely circulating, Covid-19 virus. While most of the recent pandemics have been viral, until the invention of antibiotics many of the worst pandemics, including the Black Death, were bacterial. A growing number of multidrug resistant bacterial strains underlines the risk that a highly infectious, deadly, multidrug resistant bacterial strain will emerge causing both health and economic damage. Much of the analysis and recommendations in this paper are specific to reducing the cost of future viral pandemics, but the larger message of the benefit of preparation also holds for preparations for potential bacterial threats. For example, development of new antibiotics to be kept in reserve for use only in combination therapy for multidrug resistant strains would reduce the probability of a highly damaging bacterial pandemic. Fortunately, scaling up antibiotics tends to be easier, cheaper and faster than scaling up vaccines, hence our focus on vaccines.

In this paper, we estimate the economic return to investing now in specific pandemic preparedness strategies that would enable the rapid production and distribution of vaccines worldwide in the event of a new pandemic and thus substantially reduce the economic and social costs. To estimate the returns to pandemic preparedness we first need to estimate the probabilities of future epidemics of different magnitudes and the economic and social costs of these epidemics.

2 Probability of Future Pandemics

An analysis of the frequency of pandemics in history and of emerging trends suggest that a pandemic of at least the magnitude of Covid-19 is a one in 138-year event. Mariani et al. (2021) document 476 epidemics since 1600 of which 271 have data on duration and deaths, forming the main basis of their estimations.[1] They show the distribution of annual intensity of epidemics i as indexed by mortality (specifically, deaths per thousand people per year) is well described by a generalized Pareto distribution with cumulative distribution function

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

where $\mu = 10^{-3}$ is the threshold below which epidemics are too small to leave a detectable record, $a = 0.62$ is the probability that the epidemic is below the threshold of detectability, and $\sigma = 0.0113$ and $\xi = 1.41$ are maximum-likelihood estimates of the shape parameters from Mariani et al. (2021). The complement to the cumulative distribution function, $\overline{\Phi}_0(\mathrm{i}) = 1 - \Phi_0(i)$, sometimes called the exceedance probability, has the useful interpretation as the annual probability that an epidemic of intensity i or more occurs.

The support of the distribution has a natural upper bound at 1,000 deaths per thousand, corresponding to the whole population dying out. Equation (1) reflects a further adjustment we make to avoid contaminating computations of expected values by projecting outside the support of the data. Small changes in the mass in the fat tail can have a large influence on the expected value of a Pareto random variable. But projecting the mass in the tail outside of the data support is difficult given the large confidence intervals there and the poor approximation that the Pareto law must provide as the intensity approaches the population size. As a conservative approach to address this issue, we cap the maximum epidemic intensity at the upper bound of the support of their data: $\mu'' = 5.7$ deaths per thousand per year for the Spanish flu.[2] We perform various sensitivity

[1] The authors' exclusion of ongoing epidemics like HIV and Covid-19 biases the estimated probability of a large epidemics downward since large epidemics are more likely to be ongoing that small ones.

[2] Our adjustment adds an atom of mass $\overline{\Phi}_0(\mu'')$ at $i = \mu''$. Mariani et al. (2021) leave the distribution of intensity unspecified for $i < \mu'$. The specification in equation (1) fills this gap in by adding an atom of mass a at $i = 0$ and positing zero mass for $i \in (0, \mu')$.

analyses for alternatives to the baseline distributions of pandemic risk. One of these doubles the upper bound on intensity to $\mu'' = 11.4$.

Starting with the basic distribution in (1), Mariani et al. (2021) transform it in a way that maintains a constant distribution of relative intensities but allows the rate of epidemics to vary over time. The transformation also accounts for the fact that there are few epidemics in any year, invalidating the asymptotic arguments required for equation (1) to hold directly. They transform (1) via the metastatistical extreme value distribution (MEVD), averaging the distribution of the maximum order statistic from n_t draws corresponding to the number of epidemics in year t. The resulting formula is

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

where w is the width of the window of years under consideration. Equation (2) has a particularly simple form if, following Mariani et al. (2021), we take the window to be the most recent 20 years in their dataset, during which, according to their Supplementary Figure S1(a), there were 13 years without a detectable epidemic, six years with one, and three years with two. Substituting those data, (2) becomes

$$\Phi(i) = \frac{1}{20}[13 + 6\Phi_0(i) + 3\Phi_0(i)^2]. \tag{3}$$

The rate of epidemics over the last 20 years, which factors into $\Phi(i)$ as we have just seen, turns out to be historically low. The historically low rate reflects two opposing forces. The invention of antibiotics led to a sharp decline in the probability of bacterial epidemics including the plague. Working in the other direction, models of likely mammal movements due to climate change suggests estimates of the frequency of high intensity pandemics based on past data (even recent past data) may be an underestimate of future frequency. Carlson et al. (2021) simulate likely hotspots for zoonotic spillover: (the transmission of viruses from animals to humans) based on the predicted movement of mammals through 2070 as a result of climate change. The mechanism is as follows. Climate change will force mammals to move from their existing location to locations with a (new) climate which is closer to their ecological niche. These mass movements will generate more mixing of mammals that have hitherto had little contact leading to spread of diseases across mammal species. This increases the chance of a disease spread to those mammals with the ability to pass a disease on to humans (either because of their interaction with humans or related biology).

Climate and mammal modeling suggests much of the increased mammal interaction is likely to take place in areas in Asia with high population density, further increasing the risk of zoonotic spillover.

A final factor to consider when estimating the frequency of epidemics of different severity is our changing ability to mitigate them. Observed deaths reflects both the underlying disease dynamics and our behavioral response. Taking observed deaths, Covid-19 is about a one in 60-year pandemic. However, this death toll reflects unprecedented action including working from home in a way impossible in previous pandemics and the rapid invention and roll out of vaccines. According to one study, mortality would have been eight times higher with no mitigation. A pandemic with that latent level of severity is about a one in 200-year event.

Given the uncertainties inherent in these estimations we conduct sensitivity analysis of our calculations to different distributions of pandemic probabilities.

3 Economic Costs of Epidemics

The literature on the economic and social cost of epidemics suggests costs are dominated by the few, relatively rare epidemics with high intensity. In this section we seek to estimate more precisely the relationship between the severity of an epidemic measured by mortality and its economic and social costs. Two main methodologies have been used in the literature to estimate the economic costs of epidemics: one calculates the deviation from trend for gross domestic product (GDP) for impacted countries: examples include the International Monetary Fund (IMF) estimates for Covid-19 and World Bank estimates of Ebola. An alternative method is to estimates impacts on specific sectors (such as tourism) and days of work lost to sickness and sum these to generate a total impact (see, e.g., Keogh-Brown and Smith 2008 and Lee and McKibbin 2004). Both have limitations.

Deviations from projected GDP do not take into account other shocks which may coincide with the epidemic. For example, the 2014 West African Ebola epidemic coincided with a sharp fall in commodity prices and deviations from projected GDP almost certainly overestimate the true impact in this case.[3] Deviation from GDP projections do not capture reductions in the stock of human capital which impact GDP over the long term, so these estimates need to be augmented

[3] Iron ore represented nearly 70% of exports in Sierra Leone pre-pandemic; prices fell by 146% during 2014 (Mihalyi 2015).

with the value of lives lost and lower productivity from closed schools as we do below. To be conservative we use the lower end of the range in Azevedo et al.

In contrast, micro-grounded studies often assume no flexibility in an economy: for example, colleagues of sick workers do not take up the slack, workers in tourism do not shift to other sectors temporarily, or the timing of tourism/business trips are changed rather than canceled. A few are based on surveys of reported losses where there is a risk of strategic over reporting of losses (Glennerster, M'cleod and Suri 2015). Keogh-Brown and Smith (2008) show that micro-grounded studies conducted at the time of the epidemic estimated that SARs could cost between $30 to $100 billion. However, for most countries GDP during the impacted quarters was unchanged or only slightly lower than the previous trend and higher than trend growth post pandemic suggests very limited losses. Where there is the option, we therefore give preference to studies based on deviations in GDP.

We provide novel estimates of the economic costs of pandemics by undertaking a meta-analysis of studies listed in Table 1. There were six pandemics for which we could find credible economic-loss estimates, which we could pair with mortality estimates. Along with the disease and an estimate of the total mortality losses, each row reports an estimate of the total economic losses over the duration of the pandemic gleaned from the listed source. The last column converts the economic-loss estimate, which is reported in different ways—sometimes in absolute terms in current dollars, sometimes as a percentage of global GDP—into a consistent absolute loss in billion 2019 dollars.

[Insert Table 1 about here.]

Figure 1 plots annual economic harm against annual mortality with log scales on axes for the six observations. The figure plots the best-fitting regression line, estimated to be

$$\ln G = 1428.8 + 0.5085 \ln i. \quad (4)$$

The fit is quite good, with an $R^2 = 0.53$.

[Insert Figure 1 about here.]

The regression can be paired with the distribution of pandemic intensity from equation (*4*)

to compute expected annual economic losses from pandemics:[4]

$$EL = \int_{\mu}^{1000} G(i)d\Phi(i). \tag{5}$$

Substituting from (4) and integrating yields $EL = \$67.9$ billion annually.

Expected mortality is just the expected value of i. To convert expected mortality into a monetary value, we use the value of a disability adjusted life year (DALY) implicit in the World Health Organization (WHO) standard for cost-effective health interventions. As reported in Marseille et al. (2015), the WHO judges a health intervention in a country to be cost effective if the required spending is less than three times that country's GDP per capita. Multiplying three times global GDP per capita ($17,000 in 2019) times 15 DALYs per life lost yields $765,000. Expected mortality losses are then

$$ML = \$765{,}000 \cdot \int_{\mu}^{1000} \frac{P}{1000} i d\Phi(i), \tag{6}$$

where $P/1000$ is world population in thousands, needed to convert intensity into deaths. Integrating, $ML = \$212.1$ billion annually. This is considerably higher than economic losses. The reason is that mortality losses have about twice the elasticity with respect to i than economic losses (unit elastic compared to an elasticity of 0.5085). The fat Pareto tail interacts with the higher elasticity to generate a high expected value for ML.

A final source of losses we consider is learning losses, denoted LL. We derived LL by assuming that it is proportional to EL, and use estimates for Covid-19 for the proportionality constant. In particular, we take the conservative end of the World Bank's estimated range for learning losses at an aggregate $10 trillion in lifetime earnings in present value (Azevedo et al. 2021)[5] and the IMF's estimated $13.8 trillion reduction in economic output relative to pre-pandemic forecasts (International Monetary Fund 2022), which yields

[4] The use of a Lebesgue integral in (5) accounts for the fact that or the fact that intensity i is a mixed random variable with an atom of mass at μ'' as discussed in a previous footnote.

[5] Azevedo et al 2021 use the correlation between years of schooling and wages to calculate the return to an additional year of schooling and hence the cost of closed schools. If wages reflect marginal product and private return to education reflect social returns then these estimates reflect future losses in GDP, not just individual wage losses. While Mincer (1974) equations do not measure the causal effect of education on earnings, Duflo (2001) conclude causal

$$LL = \frac{10}{13.8} EL, \tag{7}$$

or $LL = \$49.2$ billion annually. Totalling these losses up,

$$TL = EL + ML + LL, \tag{8}$$

For $TL = \$329.2$ billion annually.

The Spanish flu is the most severe pandemic for which we have estimates of social and economic harm. One uncertainty is whether the same relationship between mortality and economic and social harm stays the same for pandemics of greater severity than Covid-19 and the Spanish flu. Once someone is working from home or schools are close, a more deadly pandemic would have relatively little impact on that individual's productivity or child's learning. As severity increases some costs would continue to rise: deaths would rise by construction and maintaining essential services would be more dangerous, more costly to provide, and eventually less possible to provide. Note that more severe pandemics lasting longer is already baked into severity probabilities as these are the probability of deaths in a given year.

Unfortunately, past data is little guide to the economic costs of pandemics larger than Spanish flu. The Black Death in Europe and the Middle East and seventeenth century plague in Italy killed 30% or more of the population (Alfani 2013). But while there has been work on the impact on wages and land prices (and thus distribution) and long run institutions there has been little on GDP losses. More importantly, the evolution of our behavioral response to pandemics means any estimates of GDP loss from 1300s or 1600s may not be informative about future losses from a pandemic of a similar scale.

Given the unclear theoretical predictions of how costs vary with pandemics of every increasing severity and the paucity of data we rely on sensitivity analysis to give a sense of the range of possible costs. As our base case, we take a conservative assumption that pandemic severity is capped at the level of the Spanish flu and economic and social costs are also capped at the level of the Spanish flu. In our most conservative scenario we simply halve the intensity distribution of future pandemics. In our less conservative scenario, we allow for the possibility of

estimates are similar to cross-sectional estimates of the income benefits of education where the two can be compared. In a Spence (1973) model, education can be rewarded with higher wages even if it does not increase productivity suggesting private returns are above social returns to education. However, education also generates positive spillovers suggesting social returns are higher than private returns. We assume that on average social returns equal private returns and that mincer regressions give an adequate estimate of the social returns to education.

pandemics of greater severity than the Spanish flu, doubling the limit on epidemic intensity.

4 Reducing Duration of Pandemic Harm

Pandemic preparedness can sharply reduce the economic and social costs of pandemics by reducing the time taken from a new disease emerging to achieving widespread distribution and uptake of effective vaccines which can control a pandemic at relatively low cost to the economy. Of the roughly 2 years it took between Covid-19 emergence to sufficient vaccine being produced to be able to fully immunize 70% of the world's population, the most time consuming step was that between a safe and effective vaccine being approved and sufficient vaccine being produced to fully vaccinate 70% of the world's population. We therefore focus most of our attention on policies to reduce this step. In Section 7 we discuss policies to reduce the time taken to complete other steps.

4.1 Lessons from Covid-19

In assessing how best to accelerate large scale vaccine production in a pandemic it is important to understand key characteristics of the economics of vaccine production and learn the lessons from alternative approaches used by different countries during the Covid-19 pandemic. Past data suggest vaccines have a 72% failure rate (Ahuja et al. 2021). Production is hard and not simply a matter of copying a formula: getting high yields requires skilled technicians and time to iterate in a given facility. Regulatory approval must be given for each production facility. This not only means very large-scale production needs to be initiated as early as possible, but there are benefits in production being relatively concentrated. It is easier for the limited number of skilled technicians with intimate knowledge of the challenges of a particular vaccine to oversee more lines in an existing (already regulatory approved) facility than to oversee work in a new facility.

Social and political pressure to keep prices of vaccines low during a pandemic mean there is a high divergence between the private and social return to a vaccine. Covid-19 vaccines sold for between $6 and $40 while capacity to produce one course annually had a social value of $5,800 in early 2021 (Castillo et al. 2021).

The combination of high risk and lower private than social returns mean firms will be reluctant to commit to producing at large scale prior to regulatory approval. They will also be unwilling to pay the social value of inputs whose supply is inelastic. One of the strategies that was

effective in securing larger quantities of vaccine early during Covid-19 was reducing the risk to vaccine producers by providing financial support for large scale production even before vaccines had received regulatory approval. This was done through directly subsidizing production and by signing large scale vaccine purchase commitments prior to regulatory approval. Ahuja et al. (2021) show that the most cost-effective way to speed large scale production of Covid-19 vaccines during the pandemic was to pay the cost of putting in place additional production capacity for multiple candidate vaccines prior to them achieving regulatory approval. This allows manufacturing processes to be improved in parallel to clinical trials and regulatory scrutiny of results. It also allows stockpiles of vaccines to be accumulated which are ready to be distributed once approval is achieved.

In Spring 2020 the investment portfolio that would have produced the highest net benefits for the world consisted of investing in eight different vaccine candidates, sufficient capacity to produce 2.3 billion vaccine courses per month, of which in expectation, 500 million would have been successful (Ahuja et al. 2021). The expected benefit for the world was $137 per capita, while the cost was $37 per capita. High-income countries optimizing on their own would have invested more ($143 per capita). However, it was worth even low-income countries investing at risk: an investment of just 26 cents per capita would have generated benefits of 58 cents per capita. Below we discuss potential financing arrangements that would enable low-income countries to make such at-risk investments in future pandemics.

Concerns have been raised about the potential negative spillover to other countries of larger, early, at risk investment of this type. Disagreement about whether these had negative spillovers and should be discouraged slowed agreement on a global vaccine response. The debate arose from the concern that there was a fixed supply of vaccine production capability generated in part from a fixed supply of key inputs. However, with hindsight it is clear additional investments were able to expand total supply and there is reason to think that additional at-risk investment would have expanded it further and help in relaxing bottlenecks which then has positive spillovers to others (Ahuja et al. 2021). As supply is more elastic in the long than short run, concerns of negative spillovers are even lower for prepandemic investments.

One lesson of Covid-19 was that political realities mean countries will not delay protection of their own citizens during a pandemic in order to agree on a set of global rules on pandemic response. Even prior to a pandemic, countries will be reluctant to tie their hands to an arrangement

that might slow their response to a pandemic by forcing them to move at the pace of global consensus and thus the slowest mover. The approach to pandemic preparedness we propose below has the advantage that it is incentive compatible in that it is the optimal investment for individual countries while also generating positive spillovers for other countries.

National strategies designed to secure adequate vaccines can take two approaches: they can seek to secure more of an existing vaccine supply or seek to expand (and/or accelerate) total production. The former has negative spillovers on other countries, the latter has positive spillovers. Even if the additional production capability funded by a given country initially serves only that country, that country's demand is met sooner than otherwise; and the capacity can then be used to meet the demand of other countries. Production capacity in this paper is defined as all the necessary inputs into the production process not just equipment in factories. Ideally then, pre-pandemic preparations should be focused on approaches that increase total vaccine production capability.

4.2 Pre-pandemic Approach to Accelerating Vaccine Production

If the lesson of Covid-19 was the large benefit of accelerating vaccine capacity once a new disease emerges, what are the implications for pre pandemic preparedness? Here we evaluate the benefits of securing vaccine production capability before a new pathogen emerges by paying to expand total vaccine capacity and then paying an annual fee to reserve vaccine capacity that could be quickly switched to the production of a vaccine for an emerging threat. While we use global figures to illustrate the point, as we discuss below, we expect that much of this investment would be done by individual countries to protect their own population and in the next section give examples of what this would look like for a High Income Country (HIC) and Middle Income Country (MIC).

At its peak, the world was producing 580 million doses of mRNA Covid-19 vaccine per month: enough supply in a month to fully vaccinate 290 million people. At this rate it would take 2.25 years to have enough vaccine to fully vaccinate the world. To have sufficient capacity to vaccinate the world against a new pandemic in a timely manner using mRNA technology would thus mean keeping all the existing mRNA capacity in a functioning state and building substantially more. Instead, in response to the falling demand for Covid-19 vaccines, existing mRNA capacity is being shut down. A policy of paying the owners of existing capacity an annual fee to keep their existing capacity in place and paying others to build new mRNA capacity would, as we show below have a high expected return.

Contracts to ensure sufficient vaccine capacity was in place to vaccinate the world or an individual country for the next pandemic would need to guarantee that such capacity was functional and up to date. (During Covid-19 some reserve vaccine capacity failed.) Given the billions of dollars at stake, appropriate monitoring systems could be devised. One way to do this is to allow (and even encourage) contracted reserve capacity to be used for the production of other vaccines. For example, mRNA technology was originally invented to produce vaccines to address neglected tropical disease. Using mRNA reserve capacity to test and produce vaccines for a variety of existing diseases would keep capacity operating and up to date, as well as generate much needed learning about what type of viruses mRNA vaccines are most effective at combatting and how to improve their effectiveness.

We have used the example of mRNA capacity as an illustration, but it is also a useful technology for pandemic preparedness as it is easier to scale rapidly than many other vaccine technologies. However, the world has limited experience with mRNA vaccines and thus there is no guarantee they will be the most effective vaccine technology for all viruses.

There are 5 main vaccine platforms. Live attenuated vaccines grow a modified form of the live virus by, for example, growing it in chicken embryos for many cycles until it mutates into a virus that is less dangerous to humans. Viruses can also be directly modified in the lab. Inactivated viral vaccines treat the virus to prevent it replicating before introducing it to the human body. Subunit or conjugate vaccines take one or more parts of a virus such as a protein, sugar or coating of a virus, and replicate them and use those to generate an immune response. mRNA vaccines use the human body to multiply a part of the virus that will generate an immune response by taking part of the virus RNA code, putting it in a fat, and inserting it into the body. The human body then codes part of the spike protein from the virus which your body reacts to. Viral vectors are an alternative way to deliver mRNA or other genetic code into the body but instead of using a fat as a delivery they use another virus (eg from the common cold) to carry the mRNA into the body.

Which vaccine platform and which vaccine within a platform is most effective and fastest to develop and scale is likely to depend on the characteristics of the virus as well as the quality (and luck) of different vaccine candidates within a platform. Some viruses and parts of viruses are harder to grow at scale than others and different manufacturers will be better and worse (or unlucky) at growing them. This uncertainty and dependence on virus type is a particular challenge with live attenuated and conjugate vaccines. Live attenuated vaccines also require high levels of

biosecurity. Inactivated viruses are a very known and relatively simple technology, but they still require growing viruses with the inevitable unpredictability. Inactivated viral vaccines in addition may provide less lasting protection. Viral vector vaccines also require growing viruses at scale. They do have the advantage of using a known (rather than the new) virus as a base so it is possible to practice producing at scale in advance. However, once a new virus emerges the relevant parts of the active DNA/RNA must be inserted into the known viral vector and the new adapted virus must be grown.

The new mRNA platform has the potential to be less uncertain and virus dependent in its development and time to scale than previous platforms because it is less dependent on growing viruses as the RNA is delivered through a fat. It is important to test this potential, but the platform has already proved to be easier to scale than other vaccine approaches. For this reason and because mRNA vaccines use a more distinct production process than other vaccine platforms, our estimation assumes priority is given to mRNA platforms in reserve capacity investments.

As there will be substantial uncertainty at the beginning of a pandemic about which platform and which candidate within each platform is going to be most effective it will be important to have enough capacity to start work producing multiple vaccine candidates. Candidates typically fail and drop out throughout the process including before stage 1 trials. As candidates drop out it is possible to repurpose capacity to other candidates: experience during Covid-19 suggests this process takes about three months. As production facilities are most different for mRNA vaccines, we assume conversion between mRNA and non-mRNA production facilities is not possible within a relevant time period. In the next section we therefore estimate the cost of maintaining sufficient capacity to vaccinate the world (or individual countries) within 6 months for two platforms: mRNA and non-mRNA (henceforth “traditional” vaccines).

The existing volume of non-mRNA vaccine capacity is larger than for mRNA vaccines. However, some of the vaccines being produced are of sufficiently high value that health authorities might not want these facilities to switch to producing pandemic vaccines even during a pandemic (during Covid-19 very few facilities producing childhood vaccines switched to producing Covid-19 vaccines). Thus, despite the larger level of production capacity for traditional vaccines, much of the reserve capacity may also need to be new build.

There is a risk that none of the different vaccine platforms or candidates are effective against a virus. There are at least three reasons why it might not be possible, or very hard, to

generate a vaccine for a new virus. Like HIV, a virus that attacks the immune system is hard, though not impossible, to vaccinate against. Like Dengue there may be multiple strains and having antibodies against one strain can lead to worse outcomes if infected by another strain. Like Scarlet fever, the most dangerous reaction may be the body's reaction to the antibodies rather than the virus. In our model we adjust for this risk by including a probability that no vaccine works. To cover these risks, however, we need to accompany any pre-pandemic preparedness in vaccines with investment in research and development on antiviral drugs and securing supply lines for likely raw materials. Securing raw materials has been a constraint in scaling Paxlovid, an effective treatment for Covid-19. Fortunately, scaling production of drugs tends to be easier than scaling vaccines so keeping spare capacity for drug manufacturing pre-pandemic is less necessary.

In the next section we estimate the costs and benefits of such an approach. In the subsequent section we discuss other policy options to reduce the time between the emergence of a new disease with the characteristics to become a dangerous pandemic and widespread vaccine uptake.

5 Estimating Costs and Benefits of Contracts for Reserve Vaccine Capacity

The model in this section adapts the framework from Castillo et al. (2021) which was used to estimate the benefits of investing in Covid-19 vaccines early in the pandemic to estimate the benefits of making pre-pandemic investments. It builds on the work of the Accelerating Health Technologies team, academics and policy makers who came together during the pandemic to support policy makers seeking to accelerate access to Covid-19 vaccines.

5.1 Setup

There is a worldwide investment in pandemic preparedness that results in available capacity to produce z_R courses per year of mRNA vaccines and z_T courses per year of traditional vaccines. Each year, an epidemic of intensity i takes place following density function $\phi(i)$. Whenever a pandemic takes place, defined as an epidemic with sufficient intensity to be part of the historical record ($i > \mu'$), countries increase the production capacity up to x_R and x_T courses per year of mRNA and traditional vaccines, respectively. Note that the probability of a recordable epidemic is $a = 0.38$.

5.2 Benefits

5.2.1 Effective Vaccination Capacity

Suppose the world has capacity to produce x_R courses per year of mRNA vaccines and x_T courses per year of traditional vaccines. Let p_S denote the probability of success of some vaccine, approved as safe and effective for use against the newly emergent virus. Let p_R denote the probability that only an mRNA vaccine is successful, p_T denote the probability that only a traditional vaccine is successful, and p_B denote the probability that both mRNA and traditional vaccines are successful. Conditional on some vaccine succeeding, these are exhaustive and mutually exclusive events, implying $p_B + p_R + p_T = p_S$.

If mRNA vaccines succeed, a fraction f_R of the factories to produce them are successful and are available immediately after approval, and a fraction g_R are not and must be repurposed to use the technology of the successful factories, which means they are available with a delay Δ_R. The remaining fraction $1 - f_R - g_R$ cannot be repurposed and is therefore not useful during the pandemic. Traditional vaccines are modeled similarly, with a fraction f_T being immediately available and the remaining fraction being available with a delay Δ_T. Letting Δ_A denote the lag in months in approval time, production capacity at time t is then given by

$$X(t) = \begin{cases} 0 & t \leq \Delta_A \\ y_R x_R f_R + y_T x_T f_T & t \in (\Delta_A, \Delta_R + \Delta_A] \\ y_R x_R (f_R + g_R) + y_T x_T (f_T + g_T) & t > \Delta_R + \Delta_A, \end{cases}$$

where y_R and y_T are dummies for whether mRNA and traditional vaccines are successful, respectively.

The total number of people vaccinated at time t is then the integral of $X(\cdot)$ from time zero to time $\bar{T}$.

5.2.2 Vaccination Benefits during a Pandemic

We take a form for benefits of vaccination similar to that in Castillo et al. (2021) model for Covid-19 with some simplifications. At any given point in time, vaccination reduces the harm from the pandemic by the proportion $g(\lambda(t))$, a function of the fraction $\lambda(t) = X(t)/P$ of the world population P that has been vaccinated. Assume g is a continuous, concave, piecewise linear function such that $g(0) = 0$ and $g(\lambda) = 1$ for $\lambda > 0.7$, the threshold for herd immunity. It has

two additional kinks at 0.13 (the fraction of high-risk population) and 0.5, and the fraction of benefits at these kinks are 0.395 and 0.816, respectively.

The supplemental material for Castillo et al. (2021) explains in detail why this functional form is a good approximation for the benefits of vaccination during the Covid-19 pandemic. The function is concave because initial vaccines give larger benefits, since they are given to more vulnerable populations (such as frontline workers and the elderly, in the case of Covid-19). Other diseases might have different characteristics, but benefits are still likely to be concave given that different demographics might be affected differently by the disease. For that reason, although the benefits from vaccination might change, they are likely to follow a function similar to the one we describe here.

The total benefits from vaccination conditional on pandemic of intensity i are computed by integrating benefits over a horizon:

$$B(x_R, x_T, i) = \left\{G(i) + 765{,}000\frac{P}{1000}i + \frac{10}{13.8}G(i)\right\}\int_0^{\bar{T}} g\big(\lambda(t)\big)dt,$$

where the term in braces includes the components of economic losses, mortality losses, and learning losses conditional on i.

5.2.3 Expected Yearly Benefits of Vaccination

With the previous elements, we can compute the expected benefits $B(x_R, x_T, i)$ of having capacities x_R and x_T during a pandemic of intensity i. We discount total benefits by γ (set to 50% in the baseline) to account for the fact that some measure other than vaccines might address a future pandemic such as the possibility of improved treatments, mitigation strategies such as effective contact tracing, etc. that would preclude an important share of the benefits from vaccination. With that, the expected yearly benefits of vaccination are equal to

$$\gamma\int_{\mu'}^{1000} B(x_R, x_T, i)d\Phi(i) = \gamma\, TL\int_0^{\bar{T}} g\big(\lambda(t)\big)dt.$$

5.3 Costs

5.3.1 Preparedness Costs

For vaccine technology $V \in \{R, T\}$, there is a long-term investment cost c_V per dose per year. The yearly cost of maintaining a production capacity z_V is thus $(r + d)c_V z_V$. We assume that such capacity can be rented out to pharmaceutical firms for routine vaccine production for a fraction ϕ of the yearly cost, which means that the effective yearly cost is $(1 - \phi)(r + d)c_V z_V$. This also means that the up-front investment must be

$$c_R z_R + c_T z_T.$$

5.3.2 Pandemic-time Costs

The counterfactual to pre-pandemic investment is paying for vaccine capacity during a pandemic. During a pandemic, countries need to invest in additional capacity $x_V - z_V$ at a cost $C_V(x_V - z_V)$, where C_V is a cost function exhibiting increasing marginal costs. We assume that $C_V'(0) = c_V$.

In addition to this, we assume that countries need to incur a marginal cost of production m_V per course. Thus, every year there is a pandemic, the world incurs a cost

$$\left(\frac{m_R x_R + m_T x_T}{x_R + x_T}\right) P.$$

The total expenditure that will be necessary during pandemic years is then

$$C_R(x_R - z_R) + C_T(x_T - z_T) + \left(\frac{m_R x_R + m_T x_T}{x_R + x_T}\right) P.$$

5.3.3 Total Costs

The annual social cost is then given by

$$(1 - \phi)(r + d)(c_R z_R + c_T z_T) + (1 - a)\left[C_R(x_R - z_R) + C_T(x_T - z_T) + \left(\frac{m_R x_R + m_T x_T}{x_R + x_T}\right) P\right].$$

5.4 Scenarios

In the baseline scenario, suppose that the world doesn't make any investments in pandemic preparedness (i.e., $z_R = z_T = 0$). We assume that in that scenario countries invest some quantities x_R^0 and x_T^0 in years in which a pandemic takes place.

In alternative scenarios in which investments in pandemic preparedness are nonzero ($z_R > 0$ and $z_T > 0$), we assume that such investments reduce to some extent countries' investments during the year of a pandemic, so that $x_V = x_V^0 + (1 - \theta)z_V$.

Table 2 summarizes the parameters used in the model. Our results are robust to sensitivity analysis using the plausible ranges in the last column.

[Insert Table 2 around here.]

5.4 Results

We consider the benefits of investing in sufficient up front production capacity for vaccines in order to vaccinate 70% of the world's population, the threshold for herd immunity under certain assumptions, within six months. This would cost $60 billion up front to expand production capacity for vaccines and supply chain inputs and $1.8 billion annually thereafter to maintain capacity and would be a very high return investment.

Under our baseline risk distribution scenario, the expected net benefits would be $27.6 billion per year. This would fund 24 billion courses of capacity, allowing the repurposing of sufficient capacity to vaccinate the world for several vaccine candidates, so that capacity is ready for whichever of the candidates is successful at achieving health and safety goals. At this level of investment, decreasing the time to approve vaccines by one month would yield net benefits of $600 million per year. The expected net benefits would be $12.8 billion per year if we were to cut probability of observed pandemic in half ("half baseline risk" scenario) and $32.1 billion per year if we double the limit on epidemic intensity.

[Insert Table 3 around here.]

Even larger investments would be worthwhile. An up-front investment of $112 billion to install 45 billion courses of annual capacity would have an expected net benefit of $31.8 billion

per year. Smaller up-front investments would still be worthwhile but would have lower expected net benefits. For example, an upfront investment of $30 billion to install 12 billion courses of annual capacity would have expected net benefits of $20.7 billion per year.

These estimates for expected net benefits take into account the cost of producing vaccines for the world's population during a pandemic. They have been adjusted down by 50% to account for the fact that some measure other than vaccines might address a future pandemic. The cost could be lower depending on supply chain efficiencies and the ability to productively use the capacity outside of a pandemic.

Next, we consider the benefits of investment for a high-income country such as the United States versus that of a middle-income country such as Brazil. We compute the economic, health and learning loss harm for each country based on their world GDP share, population share, and World Bank estimates of their school-hours-lost during Covid-19 respectively. For the United States, under our baseline risk scenario, the expected net benefits from investing in sufficient upfront capacity to vaccinate 70% of its population in six months would be $2.3 billion per year, or $7 per capita. On the other hand, for Brazil, the expected net benefits would be $600 million per year, or $3 per capita. The difference in benefits per capita is driven by the greater economic losses faced by high income countries during a pandemic, while only partially offset by longer school closures in LMICs.

6 Who Pays for Pre-pandemic Investments and Should it be Coordinated?

Most of our analysis thus far has been at the global level but this does not mean that action should be taken at the global level. As the previous section indicates, it is in the interests of individual countries to invest in securing adequate vaccine capacity to be able to scale vaccine production, and thus have timely access to sufficient vaccine to rapidly immunize their population in the event of a new pandemic. Should this effort be coordinated among different countries?

In the early stages of the Covid-19 pandemic many (including the authors) assumed that investment in developing and scaling a vaccine would best be done in a coordinated way. However, it rapidly became apparent that there were relatively low benefits of coordinating. Diversification of a country's vaccine portfolio was possible even without coordination: countries could simply purchase some of many different vaccine candidates (whether or not they were produced in their country). Early-stage research and development on vaccines and clinical trials are global public

goods and thus countries could benefit from sharing costs across countries. However, production capacity greatly outweighs R&D in total costs. While there are positive externalities from one country investing in this capacity, production capacity is rival and excludable, and thus not a global public good. An important benefit of coordination is insurance against uncertainty about which country will be hit hardest and at what time.

Coordination with an insurance objective would involve countries making a pooled investment but then receiving vaccines based on how bad their case rate or death rate was. Even after the start of the pandemic there was considerable uncertainty about both the relative severity and timing of waves in different parts of the world suggesting insurance even during a pandemic could have a benefit. For example, India initially thought it would not have high mortality rates, only to be hit badly by the delta wave. As with any insurance mechanism, there is a risk of moral hazard, i.e., the risk that countries take less stringent control measures than they would because higher cases will increase their vaccine supply. Given the high costs of the pandemic even with vaccine access it is unlikely that moral hazard would be significant in this case. However, the perception that an insurance type of allocation would "reward" bad performers might still undermine efforts to agree to insurance type coordination. COVAX, a large, coordinated vaccine purchase mechanism for Covid-19 did not include any insurance element but instead initially allocated vaccines without regard to cases, mortality rates or even demand (as proxied by whether previous shipments had been utilized). The behavior of COVAX may reflect the political challenges of insurance-based coordination.

In contrast to the relatively small benefits, coordination takes time and during a pandemic time is extremely valuable. Another challenge for coordinating vaccine investment is allocation of cost. Our initial assumption, as well as the initial reaction of many others', was that an effective coordination mechanism would involve countries paying into a common pool on the basis of their income and receiving vaccines on the basis of their population (or number of high-risk individuals). However, such a mechanism would involve substantial redistribution from richer to poorer countries. Coordination with redistribution is only incentive compatible for HIC if the benefits of coordination are larger than any redistributive tax or HIC will be better off going it

alone. As coordination benefits were small in this case, HIC quickly realized it was not in their interests to join a coordination mechanism that imposed redistribution.[6]

Investments in pandemic preparedness must learn these lessons. There is a case for coordination on R&D, for example testing mRNA vaccines against multiple different existing viruses so that we learn when and where they work, development of a universal coronavirus vaccine, and development of new antiviral drugs. However, as with investment during a pandemic, the main cost is of pandemic preparedness is investment in building and maintaining capacity which is not a global public good. The insurance benefits of coordination prepandemic are larger than during a pandemic as there is greater uncertainty about which countries will be hit: the next viral pandemic may be more subject to temperature or humidity for example. Pooled investment would allow countries to pay in less than under individual investment strategies, in exchange for vaccine being prioritized on the basis of susceptibility to the virus. These rules for allocation would need to be agreed ahead of time and countries would need to trust the allocation process would be fair. There would also have to be agreement about how much of what types of capacity to maintain and an agreed pooled monitoring system. Countries at different income levels might well have very different preferences and there would be strong pressure for redistribution to be built into the system. The latter may well make pooled funding infeasible outside a grouping of likeminded, similar income countries (e.g., the European Union).

An intermediate approach would be for individual countries to make investments on their own but commit to lend their facilities to others in certain situations or to trigger a switch to producing a pandemic virus even if they themselves have not yet been badly hit. Then if the virus never took off in the country that made the investment, others could purchase the output which would be available much more rapidly than would otherwise have been the case. In other words, the country that made the investment would have priority, but others would reap some benefit.

Would individual countries going it alone with pandemic preparedness lead to an inefficient concentration or dispersal of production facilities? There is some benefit to dispersion of vaccine production globally. The heavy reliance on India as a source of Covid-19 vaccines was

[6] Redistribution does not have to simply be monetary. At one point, it was muted that countries wanting to join COVAX should commit that they would not vaccinate more than 10% of their population until every other country in COVAX had vaccinated at least 10% of their population. This effectively meant redistribution of supply that HIC had bought early and at risk. Needless to say, HIC did not agree to this condition but debate about redistribution conditions delayed the launch of COVAX.

a challenge for COVAX and MICs and LICs more generally when India required that supply to India be prioritized during the delta wave. However, it is also unlikely to be efficient to have production facilities dispersed to every country in the world. There are important economies of scale in vaccine production: each production facility needs to receive regulatory approval for each vaccine produced so the same capacity dispersed over 8 vs 4 facilities requires twice as much regulatory approval which can take months and a lot of scarce human capital. There are also fixed costs per facility in terms of the technical team running production.

This tradeoff between the benefits and costs of dispersal vs concentration of vaccine capacity should be fully internalized by individual countries suggesting that leaving it to individual countries should not lead to an inefficient allocation of vaccine capacity around the world. The caveat to this is political economy considerations. To the extent that countries making big investments in vaccine capacity will feel domestic pressure to locate the capacity in their home country or region for political economy reasons, there could be inefficient fragmentation of capacity. Overall, location of production capacity is second order to quantity of capacity.

One reason countries state for wanting regional production capacity is to hedge against the risk of export bans. However, a much better hedge against export bans during a pandemic is placing manufacturing capacity in experienced but small (in population) countries. Any country will, in the midst of a pandemic, want to secure their own vaccine supply before others and any pledges not to do so pre-pandemic are time inconsistent and thus may well be breached when the time comes. But once the home population is fully vaccinated, the pressure for export bans will be relieved. For example, at its peak the world was producing 580 million doses of mRNA a month, or over 19 million doses a day. The population of Singapore (a major vaccine producer) is 5.8 million. In other words, if all the mRNA production was in Singapore, even if Singapore decided they wanted to vaccinate their entire population with at least one dose before exporting any to the rest of the world, this would delay export by a 1/3 of a day. If all the production capacity was in India it would take over 70 days to vaccinate the entire Indian population with at least one dose. This is a more effective hedge against export bans than distributing vaccine production capacity in different continents as trade barriers within continents are often higher than between continents (the trade barriers within Africa for example are higher than between African countries and the European Union).

It is important to distinguish between the need for some dispersal of vaccine capacity and the dispersal of vaccine access. While export bans did occur during Covid-19, a much stronger determinant of unequal access to vaccines was inequality in the speed of procuring vaccines. Agarwal and Reed (2022) show that LICs and MICs received vaccines late mainly because they (and those who purchased for them) made purchases later than HICs. In particular, neither COVAX nor purchases through the World Bank involved purchasing vaccines at risk. Pre-pandemic investments by MICs are precisely designed to ensure early access through early investment. In the effort to ensure that MICs and LICs have rapid access to vaccines in future pandemics, where the capacity is based, is less important than how much capacity there is and how quickly MICs and LICs (or those purchasing for them) secure capacity it. Pre-pandemic purchasing as suggested here achieves both more total capacity and secures early access. If lower middle-income and low-income countries do not borrow to invest in purchase of vaccine capacity pre-pandemic we discuss ways to ensure they can speed up purchasing and thus accessing vaccines during a pandemic in the next section.

7 Additional Policies to Accelerate Widespread Vaccination

Above we show how reserve contracts that guaranteed the rapid availability of vaccine production capacity to produce vaccines for an emergent disease would have high economic and social returns. In this section we discuss a number of other policy interventions that could reduce the time from disease emergence to large scale vaccine production.

7.1 Accelerating Disease Detection

Better surveillance would help reduce the time between a disease emerging and detection. Disease emergence could either be as a result of a mutation of an existing disease already present in humans or the result of a virus jumping from animals to humans. Surveillance, unlike some of the other pandemic preparedness approaches discussed in this paper, would need to be broadly distributed across the world and a lot of the burden would fall on health systems of LICs and MICs (LMICs). As reducing the time from emergence to detection of a new disease is a global public good and there is a high return to health activities in LMICs, any investment in surveillance should be funded globally and not displace existing health operations in LMICs. Community health workers, for example, are often tasked with surveillance but they have a long, arguably unrealistically long, list

of responsibilities so there is a real risk that surveillance will displace other high return activities.

7.2 Accelerating Vaccine Development

mRNA technology has the potential to reduce the time between the identification of a new virus to development of a vaccine. For Covid-19, the first mRNA vaccine was developed 2 months after the genomic sequence of Covid-19 was released by Chinese scientists, even though mRNA technology was an extremely new technology.

While mRNA technology has proved highly adaptive, we cannot know for certain it will be successful against all future viruses: even a small chance of mRNA failure implies a high return to investing in more than one vaccine platform for any given disease. For this reason, above we estimate the benefits of making pre-pandemic investments linked to two different vaccine platforms.

Rather than rely on the rapid development of a vaccine after the emergence of a new disease, one approach would be to invest now in the development of a vaccine that would be effective against a range of possible future viruses. This would effectively reduce the time between disease identification and vaccine development from two months to zero. Rather than target a specific spike protein (as most Covid-19 vaccines do) a universal Coronavirus vaccine would target multiple targets on the assumption that any future mutation would not change all of these targets.

7.3 Accelerating Approval

One of the longest steps between disease identification and widespread uptake of vaccines is conducting clinical trials and achieving regulatory approval. The minimum duration of this step during Covid-19 was 8 months. The experience of Covid-19 did demonstrate the potential to speed up the trial and regulatory process. It helped that there were existing procedures for seeking emergency approval which helped streamline the approval process and had been agreed pre-pandemic, although a number of countries also reprioritized regulatory resources to focus on Covid-19 vaccines in ways that were not preplanned. The use of human challenge trials (HCTs) has the potential to dramatically reduce the time taken to conduct clinical trials on new vaccines. In an HCT, participants are deliberately exposed to a virus. In a standard clinical trial, many participants will never be exposed to the virus and there is uncertainty about who was exposed,

increasing the sample size needed to detect an impact of a vaccine with reasonable certainty. Berry et al. (2020) estimate that compared to the roughly 30,000 participants needed for regular clinical trials of Covid-19 vaccines, a HCT would require only 250 participants. Recruiting large numbers of people into a trial can take time: during Covid-19, trials were slowed as teams competed for trial subjects in the best locations (i.e., with high Covid rates but good trial infrastructure).

Based on their assumption that a study could enroll 250 people a day, Berry et al. (2020) conclude that the smaller sample for a HCT would take just 1 day for enrollment to be completed vs 120 days for a standard trial. The gap between the two doses of vaccine and time needed for antibodies to respond to the vaccine are the same under both types of trial. The surveillance period however is just 14 days with a HCT for Covid-19. This is because the exposure date is known for a HCT and directly follows the wait for antibodies to develop (28 days in the case of Covid-19).

Several factors can reduce the speed advantage of HCTs. In some cases, regulators require that the virus is modified or that a less dangerous strain is selected before participants can be "challenged" with the virus. Selecting, developing, and identifying an appropriate challenge virus could take between 30 and 120 days (Berry et al. 2020). Even if a virus is not modified it has to be isolated and grown. An alternative, and faster approach is to deliberately expose someone to an infected person, although this does not guarantee infection and thus requires a slightly larger sample size. Whichever approach is used to challenge participants with a virus, a large safety trial would also have to be run where thousands of participants receive the vaccine to check that it does not generate negative side effects but this could be run concurrently with other steps and some argue would not add to the total time needed for an HCT. Overall, and including time for regulators to process the results, Berry et al. (2020) conclude that a HCT for Covid-19 would have reduced the time needed to conduct the necessary clinical trials by between a third and a half (from 476 days to between 221 to 311 days). Note that these estimates were made in 2020 and actual clinical trials were conducted and approved faster than Berry et al. (2020) assumed. A final issue with challenge trials is that they are typically run with relatively healthy population because of ethical concerns of deliberately exposing higher risk individuals. As a result, efficacy is only known for this subset of the population.

Both the need to modify the virus before challenging participants with it, and restricting entry into the trial to very healthy individuals come from ethical concerns. Eyal, Lipsitch, and Smith (2020) argue that individuals can be made fully aware of the risks involved in an HCT. They

will benefit from close monitoring and good care if they get sick. They argue therefore that HCT with unmodified vaccine is in line with research ethics principle of respect for persons that fully informed individuals should be allowed to take risks for the benefit of others, as firefighters, medical staff and members of the armed forces do every day.

Rather than have these debates about when it is appropriate to use HCTs in the midst of a pandemic it is more effective to agree on appropriate conditions ahead of time. The UK for example approved the use of HCT during Covid-19 but by the time the conditions for these trials were established, regular trials had been underway for some time and there was limited time advantage of an HCT.

7.4 Accelerating Finance

Agreeing financing mechanisms by which low-income countries (LICs) and middle-income countries (MICs) could purchase vaccines was a challenge in Covid-19 even though early access to vaccines would have been extremely valuable. In June, GAVI raised $505 million of grant funding for Covid-19 vaccines for LICs and MICs, and in October 2020, the World Bank announced $12 billion in loan financing for countries to purchase Covid-19 vaccines.

During a pandemic there are many calls on a dwindling pool of public finances. Despite the huge returns of vaccine investment, it is hard for policy makers to divert resources from immediate relief towards investments in vaccines that may take months to pay off. Even among those middle-income countries who invested early in vaccines, many more resources were devoted to supporting those who had lost income than were devoted to reducing the duration of the pandemic through purchases of vaccines. The midst of a pandemic is also a challenging time for HICs to prioritize supporting investments that will help other countries when their domestic population is suffering a large economic and social shock.

Financing for vaccines through loans rather than grants would ensure that substantially more funding is available faster than was the case during Covid-19. Given the high returns to investing in vaccines, borrowing to finance vaccine purchases would be economically wise investments. Above we quote estimates for the returns to countries of different income levels from investing at risk in vaccines demonstrating that loans to purchase vaccines would have more than paid for themselves. The limited grant funding made available by donors meant that Gavi could not enter into sufficient purchase agreements with vaccine manufacturers at a time when HICs

were placing such orders. When loan funding became available through the World Bank the quantity of resources made available was 24 times the size of early grant funding.

However, even when loans became available few loans were used for at risk purchases. This was partly because of timing: the best time for at-risk financing was the summer of 2020 while the World Bank package only became available in October. But it was also because it was difficult for actors in LICs and MICs to take on large, at-risk contracts. Individuals in government bureaucracies who considered entering into these contracts on behalf of their governments have shared concerns that they would have been sacked or jailed for corruption if they spent hundreds of millions of dollars on vaccines that ended up not working. Establishing a funding mechanism that would allow LICs and MICs to borrow to finance purchases of vaccines early in a pandemic, even before regulatory approval has been achieved, would be a high return pre-pandemic opportunity.

A mechanism to allow LICs and MICs to finance early purchases of vaccines at risk faces three challenges: reducing the risk of expending resources on a vaccine that might fail which while representing a good investment ex ante could prove challenging politically; contracting challenges in the face of asymmetric information; and the fact that multilateral development banks (MDBs) are designed around country-by-country loans which make coordinated contracts hard.

A highly leveraged, high impact policy for donors would be to commit to pay back the loans taken out by LICs and MICs to finance at-risk vaccine purchases. Their guarantee would allow LICs and MICs to avail themselves of borrowing from MDBs and release substantially more finance for vaccine purchases than the equivalent grant financing. Their guarantee would also help boost the total available vaccine production capability with positive spillovers to the rest of the world. For LICs and MICs, it would enable them to make high-return investments without the political cost that would otherwise be associated with vaccine failure. Setting up this type of mechanism during a pandemic when officials are inevitably distracted is challenging. The mechanism should be established urgently.

A second challenge is that of contracting for at-risk vaccine purchases on a country-by-country basis. HICs had large, highly experienced contracting teams working on vaccine purchasing during 2020 and 2021 which would be hard for LICs and smaller MICs to replicate, especially at short notice. Vaccine developers also struggled to cope with the large number of countries reaching out to negotiate deals. While a level of coordination is one answer to this

challenge, the experience of COVAX also demonstrated how waiting for all countries to agree terms could slow purchasing for early movers. The World Bank struggled from a different problem: set up to do bilateral loans, they were not able to do coordinated contracts across multiple countries. A better approach would be for the World Bank or other MDBs to agree a common template contract that they would support through a loan which countries would be free to change but which they would in most cases end up adopting as is and which would meet the requirements of a donor guarantee against risk of failure. Countries could then pick how much they wanted to purchase and when, but the main loan agreement would be standard and thus not take months to negotiate. Again, many elements of the standard agreement could be agreed prior to the next pandemic. Such a standard agreement is not easy. Discussion about indemnity clauses for example delayed vaccine purchases in many countries. Agreeing a standard contract that all borrowers of the MBDs could (but would not have to) use would provide more bargaining power to LICs and MICs in their dealing with vaccine manufacturers over indemnity and other similar issues.

MDBs could also help LICs and MICs finance pre-pandemic vaccine capacity reservation contracts of the type discussed in this paper. Again, having a standard template for a contract of this kind would be a useful role for MDBs. This would help ensure LICs and MICs got access to vaccines on a timely basis in subsequent pandemics and that the contracts they undertook to achieve this access had positive rather than negative spillovers on other countries.

8 Conclusion

While the world would like to move on from worrying about pandemics, the risk of another pandemic of the magnitude of Covid-19 is relatively high. Combining data on the probability distribution of epidemics of different severity and new estimates of the relationship between epidemic severity (measured in deaths) and economic and social costs of epidemics we estimated that the world should expect the cost of pandemics to be an average of nearly $330 billion every year going forward. Investments, even very large investments, that reduce the cost of the next pandemic can therefore generate very high expected returns. In this paper we discuss a number of such investments that focus on accelerating the time between the emergence of a new disease and sufficient vaccine being available to vaccinate a large proportion of the population. While the estimates in this paper are for global investments, there is no requirement that such investments should be made on a coordinated basis and individual countries would achieve a high return to

making these investments on their own.

Because the longest lag between disease emergence and widespread availability of vaccines was the time taken to scale up vaccine production capacity, we focus particularly on investment to speed the production of vaccines against a new emergent viral threat. We show that for $60 billion upfront investment and $2.2 billion in annual expenditure would be sufficient to fund capacity to produce 20 billion vaccine courses a year and thus vaccinate 70% of the world's population in 6 months. Indeed, under reasonable assumptions even larger capacity would have a high economic and social return.

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

References

Agarwal, Ruchir and Tristan Reed. 2022. Financing Vaccine Equity: Funding for a Day-zero of the Next Pandemic. World Bank working paper no. WPS10067.

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

Alfani, Guido. 2013. Plague in Seventeenth-century Europe and the Decline of Italy: An Epidemiological Hypothesis. *European Review of Economic History* 17(4): 408–430.

Andrabi, Tahir, Benjamin Daniels, and Jishnu Das. 2021. Human Capital Accumulation and Disasters: Evidence from the Pakistan Earthquake of 2005. *Journal of Human Resources* published ahead of print doi:10.3368/jhr.59.2.0520-10887R1.

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

Barrionuevo, Jose M. 1993. How Accurate Are the *World Economic Outlook* Projections? International Monetary Fund Staff Studies for the World Economic Outlook, December.

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

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

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

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

Chappell, Bill. 2022. Governments Have Undercounted the COVID-19 Death Toll by Millions, the WHO Says. National Public Radio Goats and Soda. May 5. [Internet.] Downloaded June 8, 2022 from https://www.npr.org/sections/goatsandsoda/2022/05/05/1096842429/

governments-have-undercounted-the-covid-19-death-toll-by-millions-the-who-says.

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

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

Nir Eyal, Marc Lipsitch, Peter G. Smith, Human Challenge Studies to Accelerate Coronavirus Vaccine Licensure. *Journal of Infectious Diseases* 221(11): 1752–1756.

Fouquet, Roger and Stephen Broadberry. 2015. Seven Centuries of European Economic Growth and Decline. *Journal of Economic Perspectives* 29(4): 227–244.

Glennerster, Rachel, Herbert M'Cleod, and Tavneet Suri. 2015. How Bad Data Fed the Ebola Epidemic. *New York Times* op ed, January 30.

Global Education Evidence Advisory Panel. 2022. Prioritizing Learning during COVID-19: The Most Effective Ways to Keep Children Learning During and Post-Pandemic. Washington D.C.: World Bank.

Gopinath, Gita. 2020. Reopening from the Great Lockdown: Uneven and Uncertain Recovery. *IMF Blog*, June 24. [Internet.] Downloaded March 12, 2021 from https://blogs.imf.org/2020/06/24/reopening-fromthe-great-lockdown-uneven-and-uncertain-recovery/.

Ho, Giang and Paolo Mauro. 2014. Growth: Now and Forever? International Monetary Fund working paper no. 14/117.

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

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

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

Ismail, Kareem, Roberto Perrelli, Jessie Yang, and Rupa Duttagupta. 2020. Optimism Bias in Growth Forecasts: The Role of Planned Policy Adjustments. IMF eLibrary working paper 2020(229): 1–25. [Internet.] Downloaded June 8, 2022 from

https://www.imf.org/en/Publications/WP/Issues/2020/11/08/Optimism-Bias-in-Growth-Forecasts-The-Role-of-Planned-Policy-Adjustments-49804.

Jaume, David and Alexander Willén. 2019. The Long-Run Effects of Teacher Strikes: Evidence from Argentina. *Journal of Labor Economics* 37(4): 1097–1139.

Jong-Wha Lee and Warwick J. McKibbin. 2004 Estimating the Global Economic Impact of SARS. In Stacey Knobler, Adel Mahmoud, Stanley Lemon, Alison Mack, Laura Sivitz, and Katherine Oberholtzer, eds., *Learning from SARS: Preparing for the Next Disease Outbreak: Workshop Summary*. Washington DC: National Academies Press.

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

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

Mihalyi, David. 2015. The Miracle That Became a Debacle: Iron Ore in Sierra Leone. National Resource Governance Institute Blog, April 2. [Internet.] Downloaded June 8, 2022 from https://resourcegovernance.org/blog/miracle-became-debacle-iron-ore-sierra-leone.

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

Saunders-Hastings, Patrick R. and Daniel Krewski. 2016. Reviewing the History of Pandemic Influenza: Understanding Patterns of Emergence and Transmission. Pathogens 5(4): 66.

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

Timmerman, Allan. 2007. An Evaluation of *World Economic Outlook* Forecasts. *IMF Staff Papers* 54(1): 1–200.

United Nations. 2020. Covid-19 to Slash Global Economic Output by $8.5 Trillion Over Next Two Years. Department of Economic and Social Affairs, May 13. [Internet.] Downloaded June 8, 2022 from https://www.un.org/en/desa/covid-19-slash-global-economic-output-85-trillion-over-next-two-years#:~:text=The%20global%20economy%20is%20expected,of%20the%20previous%20four%20years.

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

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

	Deaths			Economic losses		
Epidemic (Date)	Estimated total deaths over pandemic	Source	Intensity (% of global population annually)	Estimate of economic loss over pandemic	Source	Annual economic loss (bil. 2019 $)
Spanish flu (1918–20)	32.0 mil.	Mariani et al. (2021)	5.6940515	6% global GDP	Barro, Ursúa, and Weng (2020)	1,751
SARS (2002–03)	744	Mariani et al. (2021) lower threshold		0.1% global GDP	Lee and McKibbbin (2004)	11
H1N1 (2009–01)	284,500	Mariani et al. (2021)		0.5% global GDP (lower bound)	Saunders-Hastings and Krewski (2016)	219
Ebola (2014)	11,325	Mariani et al. (2021)		0.06% global GDP	Huber, Finelli, and Stevens (2018)	64
Zika (2015–17)	1,000	Mariani et al. (2021) lower threshold		0.085% Latin American and Caribbean GDP yearly	United Nations Development Programme (2017)	5
Covid-19 (2020-22)	21.3 mil.	*Economist* (2022) excess deaths		$500 bil. GDP losses monthly	Gopinath (2020)	6,000

Notes: For SARS and H1N1, estimated deaths set to 1,000, the lower bound on observation threshold from Mariani et al. (2021) power-law distribution. For H1N1, set total economic impact to lower estimate in Saunders-Hastings and Krewski (2016) range since estimates are for affected countries only.

Table 2. Model Parameters and Ranges

Notation	Definition	Value	Range for sensitivity analysis
$r + d$	Depreciation plus interest rate	0.12	
ϕ	Fraction of capacity can be rented out	0.7	0.5–0.9
θ	Reduction of pandemic-time investments	0.25	0–0.5
c_R	Long-run cost of mRNA facilities	\$1.50 per annual course	\$0.50–\$2.00
c_T	Long-run cost of traditional facilities	\$3 per annual course	\$1–\$4
$C_R(\cdot)$,$C_T(\cdot)$	Short-run costs of facilities	†	
m_R	Marginal cost of mRNA vaccines	\$6 per course	\$4–\$12
m_T	Marginal cost of traditional vaccines	\$3 per course	\$2–\$6
p_B	Probability both platforms successful	0.5	0.3–0.7
p_R, p_T	Probability platform alone successful	0.15	0.1–0.2
f_R, f_T	Fraction platform capacity succeeds	0.3	0.2–0.4
g_R, g_T	Fraction platform failing capacity repurposable	0.4	0.2–0.6
Δ_R	Time to repurpose mRNA	2 months	1–3 months
Δ_T	Time to repurpose traditional	6 months	3–9 months
Δ_T	Time to approval	12 months	
γ	Fraction of harm avoided due treatments	0.5	0.3–0.7

† Short-run cost of facilities equals long-run cost for respective platform (c_R or c_T) up to 100 million annual courses, increasing with elasticity 1 thereafter.

Table 3. Results for Expected Pandemic Costs and Program Net Benefits

	Pandemic scenario		
	Half baseline risk	Baseline risk distribution	Double upper threshold on intensity
Pandemic Losses (bil. $ annually)			
• Mortality losses (ML)	106.0	212.1	266.1
• Economic output losses (EL)	33.9	67.9	70.3
• Learning losses (LL)	24.6	49.2	50.9
• Total losses (TL)	164.5	329.2	387.3
Annual program benefits, costs (bil. $ annually)			
• Expected annual net benefits	12.8	27.6	32.1
• Expected annual gross benefits	15.0	29.8	35.1
• Expected annual program costs	2.2	2.2	2.2
Up-front program costs (bil. $)	60	60	60

All entries are in billions of 2019 dollars.

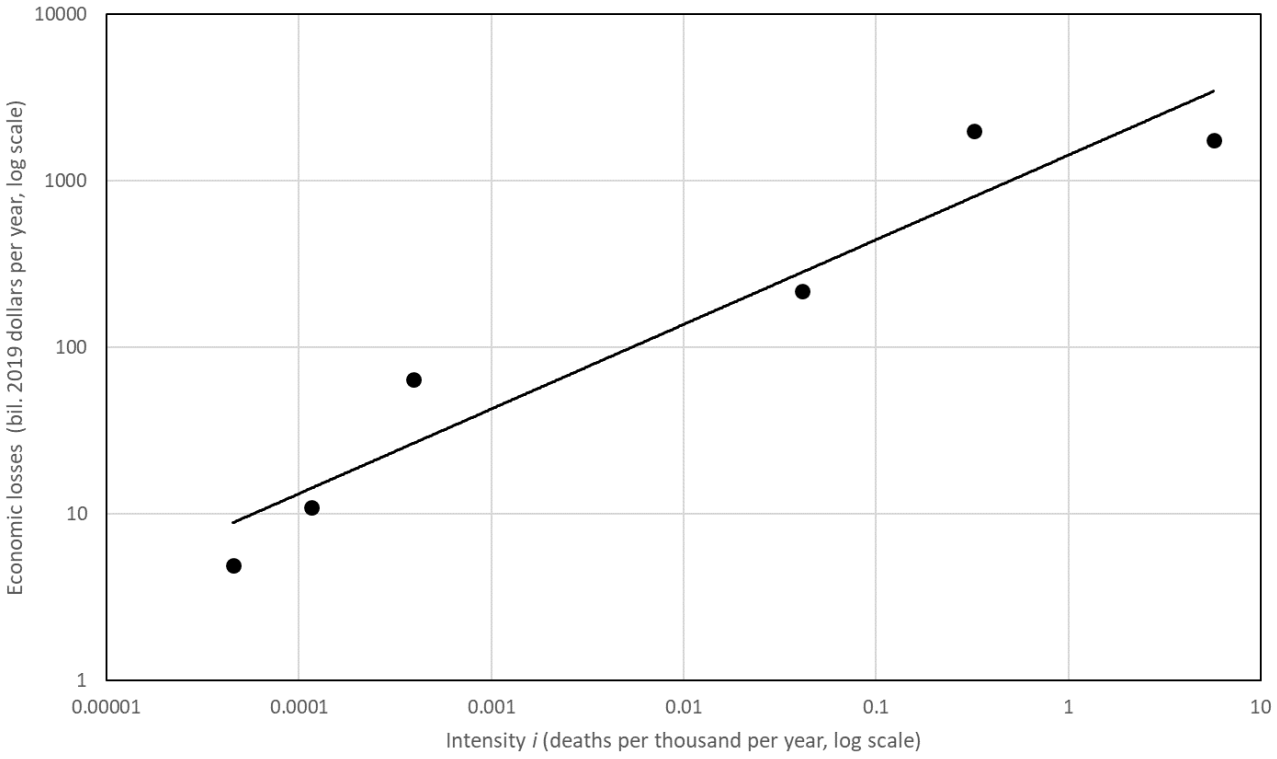

Fig.1 Relationship Between Epidemic Intensity Economic Losses in Historical Pandemics. Notes: Data points from Table 1. Regression line given by equation (*4*).

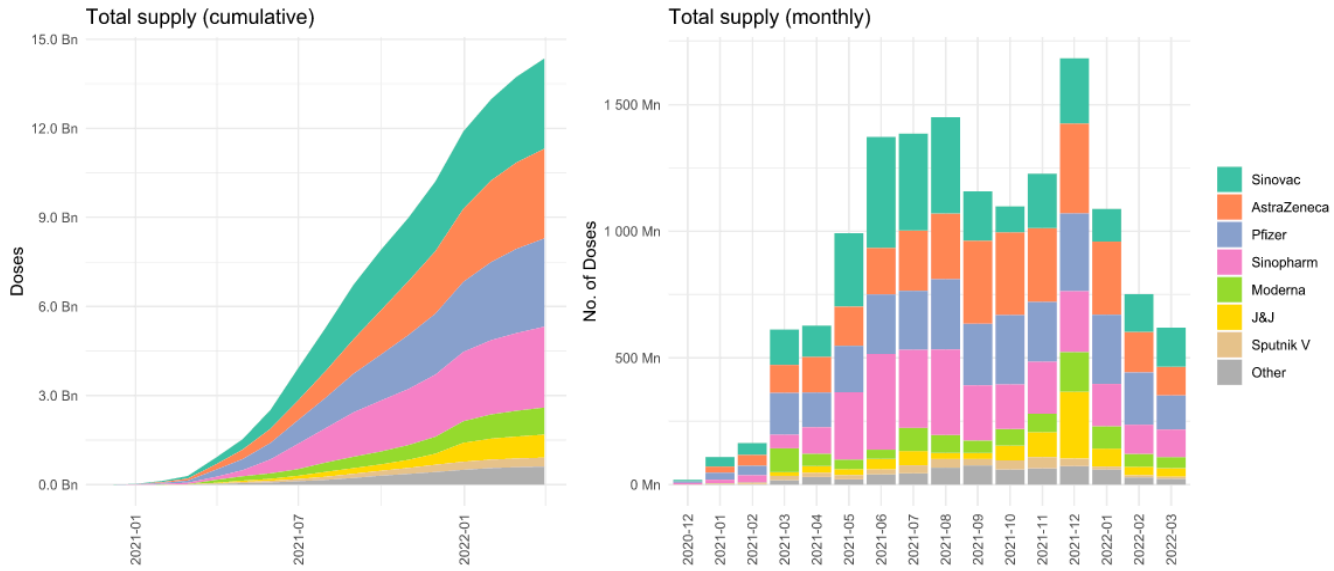

Fig. 3 Covid-19 Vaccine Supplies Over Time