

NBER WORKING PAPER SERIES

PREPARING FOR A PANDEMIC:
ACCELERATING VACCINE AVAILABILITY

Amrita Ahuja
Susan Athey
Arthur Baker
Eric Budish
Juan Camilo Castillo
Rachel Glennerster
Scott Duke Kominers
Michael Kremer
Jean Nahrae Lee
Canice Prendergast
Christopher M. Snyder
Alex Tabarrok
Brandon Joel Tan
Witold Więcek

Working Paper 28492
<http://www.nber.org/papers/w28492>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2021

Athey acknowledges a generous gift from Schmidt Futures. Ahuja and Kremer acknowledge generous funding from the Wellspring Philanthropic Fund. Baker and Więcek acknowledge generous funding from Emergent Ventures. Kominers acknowledges generous funding from the Washington Center for Equitable Growth. Kremer is a Senior Advisor to the World Bank on vaccines; he is also a co-founder and scientific director of USAID's Development Innovation Ventures. Lee is a senior economist at the World Bank and has some involvement in vaccine financing. Athey, Kremer, and Snyder served on a pro-bono advisory group on the design of COVAX. Budish, Kominers, and Prendergast serve on a pro-bono expert group advising on the design of the COVAX exchange mechanism. For contributing input and expertise, we are indebted to numerous experts in vaccine procurement from international organizations and governments. Whit Athey, Owen Barder, Tasneem Chipty, Lisa Danzig, Greg Larson, Jonathan Levin, Jessica Pickett, and Duc Tran provided invaluable advice. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Juan Camilo Castillo, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Jean Nahrae Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon Joel Tan, and Witold Więcek. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Preparing for a Pandemic: Accelerating Vaccine Availability

Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Juan Camilo Castillo, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Jean Nahrae Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon Joel Tan, and Witold Wiecek

NBER Working Paper No. 28492

February 2021

JEL No. D02,H12,I11,I18,O19,O31

ABSTRACT

Vaccinating the world's population quickly in a pandemic has enormous health and economic benefits. We analyze the problem faced by governments in determining the scale and structure of procurement for vaccines. We analyze alternative approaches to procurement. We find that if the goal is to accelerate the vaccine delivery timetable, buyers should directly fund manufacturing capacity and shoulder most of the risk of failure, while maintaining some direct incentives for speed. We analyzed the optimal portfolio of vaccine investments for countries with different characteristics as well as the implications for international cooperation. Our analysis, considered in light of the experience of 2020, suggests lessons for future pandemics.

Amrita Ahuja
Douglas B. Marshall Jr. Family Foundation
600 Jefferson Street, Suite 310
Houston, Texas 77002
amrita.b.ahuja@gmail.com

Susan Athey
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
athey@stanford.edu

Arthur Baker
Harvard University
arthur.w.baker@gmail.com

Eric Budish
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
eric.budish@chicagobooth.edu

Juan Camilo Castillo
Department of Economics
University of Pennsylvania
jccast@sas.upenn.edu

Rachel Glennerster
Foreign, Commonwealth and
Development Office
22 Whitehall
London, SW1A 2EG
United Kingdom
rachel.glennerster@fcdo.gov.uk

Scott Duke Kominers
Harvard Business School &
Department of Economics
Rock Center 219
Soldiers Field
Boston, MA 02163
kominers@fas.harvard.edu

Michael Kremer
University of Chicago
Kenneth C. Griffin Department of Economics
1126 E. 59th St.
Chicago, IL 60637
and NBER
kremermr@uchicago.edu

Jean Nahrae Lee
World Bank
1818 H St NW
Washington, DC 20433
Jlee20@worldbank.org

Canice Prendergast
Graduate School of Business
The University of Chicago
1101 E. 58th Street
Chicago, IL 60637
canice.prendergast@ChicagoBooth.edu

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

Alex Tabarrok
Department of Economics
George Mason University
Tabarrok@gmu.edu

Brandon Joel Tan
Harvard University
btan@g.harvard.edu

Witold Więcek
WAW Statistical Consulting Ltd.
witold.wiecek@gmail.com

I. Introduction

Each month, the COVID pandemic kills 300,000 people (as of January 2021) and reduces global GDP by approximately \$500 billion. The full cost, including losses to health and human capital is likely much larger. Cutler and Summers (2020) estimate total losses of \$16 trillion (around \$800 billion per month of the pandemic) for the US alone.

Beyond the epidemiological externalities that motivate governments to play a central role in vaccination programs, new issues arise in a pandemic. Vaccine candidates face substantial risk of failure or delay in proving safety and efficacy. Normally firms wait to resolve this uncertainty before scaling up manufacturing, a risky and time-consuming process requiring specialized facilities and specific investments. However, in a pandemic the benefit of speed makes it socially valuable to invest in manufacturing capacity in parallel with testing.

Social and political limits on vaccine prices during pandemics mean the social value far exceeds the commercial returns to vaccine manufacturers from installing capacity. In a companion piece (Castillo et. al 2020), we estimate that increasing the total supply of vaccine capacity available in January 2021 from 2 billion to 3 billion courses per year generated \$1.75 trillion in social value, while additional firm revenue was closer to \$30 billion assuming a price of \$15 per dose. Additional profits would be smaller still. Since unprecedented acceleration of manufacturing involves substantial, risky investment, carefully crafted public intervention is needed to align social and private benefits. What magnitude and structure of public intervention is best?

Early in the pandemic, we analyzed alternative approaches to procurement, arguing that buyers should directly fund manufacturing capacity and shoulder most of the risk of failure in exchange for the right to buy doses at close to marginal cost should the vaccine candidate be successful, while maintaining some direct incentives for speed. We analyzed the optimal portfolio of vaccine investments for countries with different characteristics as well as the implications for international cooperation. Our analysis, considered in light of the experience of 2020, suggests lessons for future pandemics.

II. Design of Procurement Contracts

Since many aspects of the vaccine development process are beyond the full control of vaccine developers (e.g., the rate at which trial participants will be infected or the occurrence of adverse

events that pause trials), bonus/penalty clauses for speed large enough to reflect its value would involve unacceptable levels of risk for firms and could have unintended consequences. Capacity installation is more predictable and contractible.

Buyers can incentivize early installation of capacity by reimbursing firms for the cost of installation (before testing and regulatory approval are complete), thus transferring risk of failure from firms to themselves (push contracts); or they might commit to purchase vaccines on specified terms, but only if regulatory approval is completed successfully, leaving the risk with firms (pull contracts). We analyze the costs and benefits of these alternatives in Appendix A.4. We consider a stylized environment in which firms have private information on their chance of success, buyers want many firms to invest—including those with only a modest chance of success—and the cost of capacity is observable.

When pull contracts are used, a large pull payment is required to induce at-risk capacity construction for candidates with a low chance of success. When it is not possible to observe and condition payments on the probability of success, governments wishing to induce a diverse set of firms to make at-risk investments must design pull contracts to offer all firms a price high enough to induce the marginal firms to invest. This price structure generates substantial rents for firms with a high probability of success and is therefore expensive compared to cost reimbursement. Our analysis implies that buyers should contract directly on capacity, relying primarily on at-risk cost reimbursement (push funding). However, in practice, buyers have some information on probability of success, do not perfectly observe costs, and—critically—want to incentivize speed. We therefore recommend that push payments cover less than the full cost, giving firms “skin in the game” and deterring those with no realistic chance of success from accepting push funding. A pull component can be calibrated to induce the marginal firm to participate and structured to incentivize speed.

We contrast the situation we study with the classic case for an advance market commitment, in which resources are severely limited and intended to address a substantial research and development challenge. It is harder to estimate the cost of research and development than capacity installation, and more difficult to judge which activities or firms should be funded to give the best chance of success. In the classic case, pull funding induces investment only from firms that are likely to succeed, thus aligning incentives.

How do our recommendations compare with what happened in 2020? While full contracts are not public, many deals in 2020 incorporated push payments by governments, covering the costs of late-stage trials and scaling up of vaccine production, including investments in inputs such as syringes and vials. In exchange, firms committed to deliver a specified number of doses within a certain time frame. In contrast, Pfizer's contract included only pull funding. Buyers committed to purchase a given quantity by a given date. Pfizer built capacity in advance of clinical approval at its own risk, but not enough to serve the world in 2021.

III. Selecting a Portfolio

How should a buyer decide how many vaccine candidates to support and how much capacity to procure prior to regulatory approval? We consider a model in which the buyer accelerates delivery of a successful candidate by 3 to 6 months by choosing the capacity for each candidate (courses per month) that will be installed at the buyer's risk.

Based on expert opinions, historical data, and a database of the vaccine pipeline (WHO, 2020), we constructed estimates of the probability of success for each vaccine, updated throughout 2020; Appendix A.1 presents the base case probabilities as of August 2020. We made assumptions about the correlation of success among vaccines using similar technologies and calculated the expected number of successful courses across vaccines. Diversification across candidates and platforms increases the probability of success. Using country-specific estimates of mortality and GDP loss, we estimate the total health and economic benefits a country can expect from investing in different vaccine portfolios. One critical and uncertain parameter is the elasticity of supply of vaccine capacity. Many experts in industry and in international organizations believed there were hard limits on how much supply could be created in the relevant time. A standard economic analysis would suggest that prices would reflect the value of obtaining early doses of vaccine for the marginal country, which would be immense, and that if supply is inelastic and prices are subject to a ceiling, shortages ensue. To illustrate this, Appendix A.2 considers a range of assumptions about supply elasticity.

Across scenarios, we find early at-risk investments in vaccine manufacturing capacity would have had, in expectation, large net benefits for countries for all levels of income (Appendix Table A.1). However, with elastic supply, global net benefits are roughly two times larger. The portfolio that maximizes net benefits varies substantially across countries. Higher income

countries (HICs) find it worthwhile to purchase more courses per capita and a larger portfolio of candidates, while lower-income candidates invest only in candidates most likely to succeed.

How does this model compare with actual deals made during 2020? Many countries made bilateral deals. HICs, especially the U.S. and U.K., invested billions of dollars at risk, contracting for large numbers of courses across multiple candidates. These investments will accelerate the end of the pandemic. Upper middle-income countries made deals for a smaller number of doses and candidates (Duke Global Health Innovation Center, 2020). While most lower income countries (LICs) did not purchase doses, they anticipated receiving them from donors through COVAX. The unprecedented investments that countries made were still smaller than our model suggested.

Because our model assumed a relatively low chance of success for each vaccine candidate and only a modest correlation between these probabilities, it suggested that it was optimal to invest in significant capacity for each of many candidates, including some early-stage candidates. Ex post, multiple candidates turned out to be successful, including the never-before used mRNA platforms. Perhaps this was a lucky realization, or perhaps we underestimated the chance of success or correlation. However, investing in the amount of capacity recommended by our model (as of August 2020) would have allowed the U.S. to complete vaccination by March 2021 rather than waiting until summer as is currently expected, and the world to complete vaccination by October 2021 rather than June 2022. Using a more conservative approach to computing economic harm from the pandemic than Cutler and Summers (2020), and also accounting for diminishing returns from vaccination in a conservative way, we estimate that this acceleration would have had benefits of \$167 billion for the U.S. and \$1.14 trillion for the world (Table 1), substantially higher than likely costs.

Table 1: Advance capacity expansion effects

Advance capacity investment	At-Risk Capacity of Approved Vaccines (bn courses)	Benefits relative to Zero At-Risk (\$bn)	Vaccination complete by
<i>Panel A, U.S.</i>			
Recommended	1.05	556.9	Mar 2021
Actual	0.45	389.9	Jul 2021
Zero	—	—	Oct 2021
<i>Panel B, World</i>			
Recommended	7.12	2748.7	Oct 2021
Actual	3.75	1606.4	Jun 2022
Zero	—	—	Sep 2022

See Appendix A.3 for full calculation.

In a companion piece (Castillo et. al), we estimate that even at this late stage, investment to expand manufacturing capacity would have large benefits. This could mean repurposing facilities, adding lines to existing ones or increasing throughput along the supply chain. There have also been interesting proposals to wring more output out of existing factories, for example by producing lower-dose vaccines. Governments could elicit bids from firms to identify ways to produce more doses.

IV. International Considerations

Many observers voiced concern about “vaccine nationalism” during 2020. This included fears that countries would prevent export of vaccines, and some producer countries such as India reportedly considered the idea, according to the CEO of Serum Institute, the largest vaccine producer in India. As a result of the risk, countries have distorted investment disproportionately towards domestic vaccine candidates. To date, most international export contracts have been honored. This fact has allowed countries to diversify their portfolios, which our analysis suggests is extremely valuable. Given that many vaccine candidates have proved successful, diversification has so far not been critical, but this could not have been predicted and may not be the case in future pandemics. Reinforcing norms of honoring contracts and not banning exports would be valuable.

Vaccines delivered earlier are more valuable than those delivered later. This distinction creates challenges if countries sign contracts believing that they are entitled to doses soon after a vaccine is approved but are in fact placed in a “queue” behind other countries for the output of the same capacity. International norms that encourage firms to commit to construct sufficient manufacturing capacity to meet promised delivery timelines might increase capacity creation while reducing uncertainty.

Many commentators have emphasized that there may be negative externalities on other countries from ordering vaccines if the supply of manufacturing capacity is inelastic, leading to high prices or, if there are constraints on pricing, shortages. However, our analysis suggests that, to the extent that supply is elastic, there will also be positive externalities by expanding global capacity, allowing manufacturers to serve other countries faster than they would have otherwise. If supply is perfectly elastic, the negative externality is eliminated, while the world as a whole benefits from faster vaccination. The fact that actual capacity available at the end of 2020 exceeds initial forecasts suggests substantial elasticity even in the short run. Indeed, Pfizer argued that earlier investments in its supply chain could have accelerated capacity expansion (LaFraniere and Thomas, 2020).

For future pandemics, elasticity of supply of vaccine capacity is a policy variable. It can be increased through investment in supply chains, either by stockpiling or by building extra manufacturing capacity for key intermediate inputs, such as bioreactors, delivery devices, or adjuvants. These actions would not only be directly valuable, they would reduce incentives for governments to use emergency authority to prohibit exports, or race to secure supplies in ways that could create negative externalities for other countries. If we anticipate social constraints on pricing during a pandemic, reaching the efficient level of input capacity or stockpiling will require public subsidy.

There have been some attempts to centralize global vaccine procurement. This is not necessary to diversify vaccine portfolios, since this can be achieved through trade, but it may have other benefits. First, monopsony power can be used to hold down prices. However, during the current pandemic, prices have not skyrocketed. Second, a central procurement vehicle can be used to coordinate donations to lower income countries, which is desirable on humanitarian grounds and to prevent disease spread from these countries. In 2020, the COVAX facility played this role. Third, there can be economies of scale in contracting, planning and supply chain investment.

Finally, centralization can facilitate efficient prioritization. From a global health perspective, the allocation should respond to local conditions as well as target individuals by vulnerability or their externalities.

International agreements require agreement about size, scope, and allocation. In Appendix A.3, we analyze the incentives of countries to participate in agreements with alternative configurations of countries as well as allocation schemes. Our analysis shows that if a centralized arrangement allocates vaccines to countries in proportion to population, then HICs prefer to purchase bilaterally instead. Linking allocation to contributions more closely strengthens incentives for HICs to invest. Moreover, as the optimal portfolio differs across countries, centralized agreements could allow flexibility in how much to invest and in what candidates.

How does this reasoning match the events of 2020? Most deals have been bilateral. Early in 2020, attempts at enlisting HIC countries in international cooperation faltered, other than through the E.U., which invested conservatively and serves a group of countries with similar incomes and similar needs. COVAX now allows countries to buy different quantities of vaccine, and most self-financing countries opted for an “optional purchase agreement”, giving them flexibility about which vaccine to purchase.

V. Conclusions

During a pandemic, expanding manufacturing capacity for a wide portfolio of vaccine candidates has large benefits. It is efficient to contract on capacity rather than doses and to primarily use push funding. Investing in capacity for or stockpiles of intermediate inputs used in vaccine production could enable faster and cheaper capacity installation during future pandemics, yielding significant benefits for the global economy and global health.

References

Castillo, Juan C, Susan Athey, Arthur Baker, Eric Budish, Rachel Glennerster, Scott Duke Kominers, Michael Kremer et al. 2020. *Accelerating Broad Availability of a COVID-19 Vaccine*. Unpublished.

Cutler, David M., and Lawrence H. Summers. 2020. *The COVID-19 pandemic and the \$16 trillion virus*. *Jama* 324, no. 15 (2020): 1495-1496.

Duke Global Health Innovation Center. 2020. "Launch and Scale Speedometer." Accessed Dec. 29, 2020. <https://launchandscalefaster.org/covid-19>.

LaFraniere, Sharon, and Katie Thomas. 2020. "Pfizer Nears Deal with Trump Administration to Provide More Vaccine Doses." *The New York Times*, December 22, 2020, sec. U.S.

WHO. 2020. "Draft Landscape of COVID-19 Candidate Vaccines." July 31, 2020.

Appendix

This appendix outlines our methodology. We focus on COVID-19, but a similar approach can be applied to future pandemics. We made various assumptions and abstractions based on available data and information, and we encourage future work to refine and extend our framework.¹ Section A.1 considers the selection of an optimal portfolio of vaccine candidates in partial equilibrium. Section A.2 considers global demand and supply for vaccine manufacturing capacity in general equilibrium. Section A.3 analyzes the incentives for international cooperation. Section A.4 discusses the structuring of procurement contracts.

A.1 Optimal Portfolio of Vaccine Candidates

This section considers the problem facing a country of selecting the optimal portfolio of vaccine candidates, taking as given the price of installing capacity for those candidates. In our baseline analysis, we took the optimum to be the portfolio that maximizes social welfare (social benefit from vaccine consumption net of capacity and production costs). The focus of this analysis was not to guide specific vaccine purchases where the characteristics of each individual candidate must be carefully considered, but to provide a framework for making decisions on the size and scope of a vaccine portfolio.

In Section A.1.1, we set up the model. Section A.1.2 describes the benefits calculation, and Section A.1.3 outlines the portfolio maximization problem. We present baseline results in Section A.1.4, and results for alternative parameters in Section A.1.5.

A.1.1 Setup

In the model, the country invests in vaccine candidates in its portfolio “at risk,” that is, in parallel with clinical trials. While this strategy carries the risk of delivering no return on investment for candidates that do not succeed in clinical trials, it has the advantage of accelerating the delivery of successful vaccines to the buyer.

We modeled the probability that each vaccine “succeeds,” where that probability is considered from the perspective of a buyer investing in advance—i.e., while clinical trials are

¹ The model described in this appendix is part of our on-going effort to study vaccine policy, which we have updated throughout 2020, to provide guidance to decision-makers based on the latest available information.

ongoing. Success from the perspective of our model includes not just successful clinical trials and regulatory approval, but also being able to establish vaccine manufacturing capacity² and associated supply chains as well as meet approval for manufacturing. For vaccines, manufacturing is typically a multi-step process. Each step of the process (e.g. the step where bioreactors are used to grow cells) must be tested, and there can be setbacks if processes do not perform as intended, as occurred in practice for Johnson and Johnson (Zimmer et al. 2021). In addition, each facility must be separately approved, and changes in suppliers lead to a requirement to recertify the production process (Plotkin et al. 2017). Different production processes have different bottlenecks or stress points. Pfizer's mRNA vaccine, for example, is technically easier to manufacture than traditional vaccines, but as the technology is new, Pfizer had to build new production facilities entirely under its own management. The AstraZeneca vaccine is easier to manufacture, and AstraZeneca wrote licenses for production in several facilities around the world. However, differences between facilities resulted in errors in the clinical trials which delayed approvals (Walsh 2020).

We first considered the question of how much acceleration can be gained if an investment is made at a point in time. Suppose that it takes M months to scale up manufacturing and that approval of the vaccine is M' months away. Then if $M' > M$, early investment will speed up availability by M' months. On the other hand, if $M' < M$, then early investment will speed up availability by M months. Let T denote the number of months availability is accelerated. Our calibrations considered cases $T = 3$ or $T = 6$. Early in 2020, approval was projected to be more than 6 months away for all candidates, and industry estimates suggested that 6 months was an aggressive timeline (indeed, it was unprecedented), so 6 months of acceleration was a reasonable estimate of acceleration (corresponding to repurposing existing facilities, certifying at-scale manufacturing processes, and solving distribution problems specific to the candidate). Later in 2020, approval was near for several candidates, and so 3 months of acceleration was a more reasonable assumption for leading candidates. In practice, during 2020, firms began manufacturing doses prior to approval while they were in the process of scaling up capacity, creating a stockpile of millions of doses at the time of approval.³ Given slippages in manufacturing

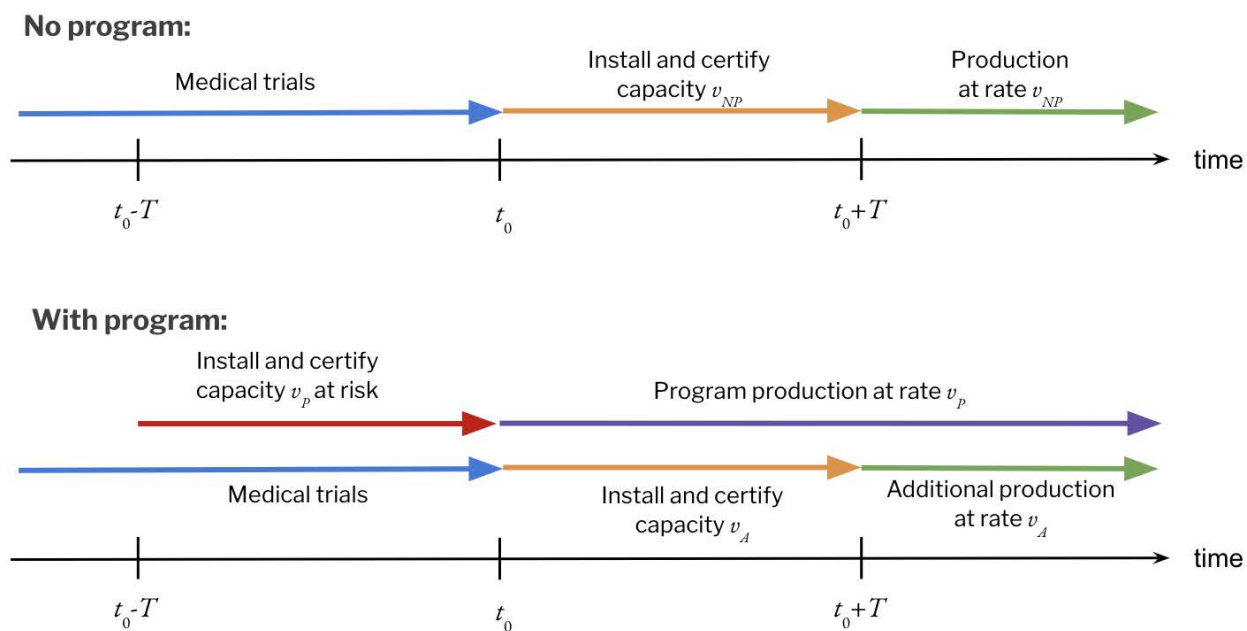
² In capacity we included the factory, securing a reliable supply of inputs, and establishing necessary quality control procedures.

³ We abstract away from stockpiling.

and the huge orders to fulfill, however, most manufacturing for the COVID-19 pandemic will take place after approvals (Goldhill 2020), in line with the model.

Figure A.1 illustrates the timeline of capacity investment with and without advance investment. Without the investment program, firms will proceed sequentially from trials to installing and certifying capacity and then to production (for reasons described in the paper). With the program, firms install and certify capacity at risk in parallel with clinical trials which allows for production to begin T months earlier which is immediately after approval. As illustrated in Figure A.1, additional capacity may be constructed after approval as well.⁴

Figure A.1: Illustration of Accelerating Vaccine Manufacturing Capacity



There are J vaccine candidates. Country i purchases a portfolio of manufacturing capacities (courses / month) from different candidates. Our goal was to compute the optimal portfolio for the country.

We assumed that besides paying for installing vaccine capacity, countries must also pay a price for each manufactured course. We assumed that eventually the same fraction of the population is vaccinated with and without investment, but to the extent the portfolio chosen by the

⁴ We abstract away from differential timing of vaccines. A richer model might take into account learning from earlier vaccine candidates before choosing to invest in later candidates.

government is successful, a higher fraction of the population will be vaccinated earlier than otherwise. Additionally, we assumed that investing early does not affect the price per vaccine course. These two assumptions allowed us to abstract away from the price per vaccine course and focus only on the price of capacity installation, since the total price paid for vaccines is the same regardless of the portfolio chosen by the country.

Note that in deals that were actually signed in 2020, terms were typically described in news reports as relating to a number of doses delivered by a particular date at a particular price per dose. Some deals relate to particular manufacturing facilities, e.g. deals signed by European manufacturers were in some cases associated with European facilities. Although cost details are not publicly available, prices appeared to be relatively close to public estimates of average cost, with some profit margin. The details of signed contracts are more complex than in our model. We study contracting in more detail in Online Appendix Section A.4; here we focus on the capacity investment problem.

A.1.2 Benefits

We model the benefits for the country in two steps. First, we model the distribution of total effective capacity V_i for each country i , given the country purchases a portfolio v_i , a list of courses from each candidate j . By total effective capacity we mean the number of vaccine courses with regulatory approval that the portfolio delivers per month. This is the ex-post quantity measured after clinical trials finish, at which point it is clear which candidates are successful. Second, we translate the effective capacity V into health and economic benefits from vaccination.

Effective capacity

The effective capacity for country i depends, first, on the success of individual candidates. Let y_i be a dummy variable for whether candidate j is successful. For each candidate j , we assign a platform p and subcategory s following public sources (WHO 2020 and updates, Le et al. 2020). For example, the Moderna vaccine was an RNA vaccine under subcategory LNP-encapsulated mRNA. We assumed that the candidate is successful if all the following events happen:

- No overall problem prevents feasibility of a vaccine (denoted by $x_o = 1$, with prob. q_o)
- No problem prevents success at the platform level (denoted by $x_p = 1$, with prob. q_p)

- No problem prevents success at the subcategory level (denoted by $x_s = 1$, with prob. q_s)
- No problem prevents success at the individual vaccine level (denoted by $x_j = 1$, with prob. q_j)

Candidate j is successful if $y_j = x_o x_p x_s x_j = 1$, and $\Pr(y_j=1) = q_o q_p q_s q_j$.

This structure implicitly introduces correlations between different candidates. There is an overall correlation across all vaccines through x_o (overall vaccine feasibility). There is a stronger correlation between vaccines that belong to the same platform—through x_p —and an even stronger correlation between candidates in the same subcategory—through x_s .

In order to use the model to estimate the social benefits of vaccination, it was necessary to assign values to success probabilities, recalling that success means reaching a substantial manufacturing scale in a short period of time. Early in the pandemic we made initial assessments of these probabilities, and then we updated them over time as information became available. Our assessments of appropriate values of q_o , q_p , q_s , and q_j were based on discussions with experts as well as historical data (DiMasi et al. 2010, Pronker et al. 2013, WHO 2020 and updates, Le et al. 2020, Hodgson 2020, Kalorama Information 2020, Lurie et al. 2020). We assigned higher probabilities of success to candidates using vaccine technologies that had been used before, where processes were better established for both the science and manufacturing, and where supply chains were better established; we assigned a lower probability of success to mRNA candidates than to inactivated virus candidates. We further assigned higher probabilities of success as candidates successfully moved through clinical trial phases. We allowed for correlations across platforms and also for the virus, i.e. there was some probability that properties of the virus would make it difficult for any vaccine to work successfully. The upshot of our modelling was that in summer 2020 no single vaccine candidate had a greater than 50% chance of success. It followed that a large, diversified portfolio was valuable.

In our model as of August, 2020, we assumed that $q_o = 0.9$ and $q_s = 0.8$. We assumed that q_p varies by the past record of each platform. Viral vector, inactivated, attenuated, and protein subunit vaccines, which had been extensively used in the past, were assigned $q_p = 0.8$. RNA vaccines, a promising technology that had never been approved for human use before the COVID-

19 pandemic, was assigned $q_p = 0.6$. For DNA vaccines, a more experimental technology, we set $q_p = 0.4$.

Finally, q_j varies by stage, or how advanced trials are. We assigned these probabilities as a function of the clinical phase a candidate was currently in (preclinical, phase 1, phase 2, phase 3). These probabilities took into account that there was still uncertainty even for a candidate in phase 3. There might be challenges with finding an effective dosing regime that does not induce too many side effects, as well as problems in scaling manufacturing or procuring the inputs needed for manufacturing or distribution. We set q_j at 0.5 for vaccines in phase 3 clinical trials, at 0.32 for vaccines in phase 2 clinical trials, at 0.23 for vaccines in phase 1 clinical trials, and at 0.14 for vaccines in preclinical trials.

Our code (available for download at <https://github.com/jc-castillo/vaccineEarlyInvest>) allows users to input alternative parameters for the success probabilities.

If country i chooses some portfolio v_i , the total effective capacity it obtains is the sum of the installed capacities over all successful candidates:

$$V_i = \sum_j y_j v_{ij}$$

Benefits as a function of capacity

Benefits come both from economic and health benefits, taking into account country-specific characteristics: GDP losses and mortality due to COVID-19, as well as the fraction of the population that is high-risk--i.e., the elderly and healthcare workers. We discounted benefits to take into account that an effective treatment might be developed, that non-pharmaceutical interventions may stop the pandemic or that herd immunity might be achieved before any capacity becomes available.

Benefits vary as a function of the number of people who are vaccinated at any point in time. Let $\lambda_i(t)$ be the fraction of the population of country i that has been vaccinated at a time t . Suppose the country does not obtain any vaccines from early capacity. In this scenario, we assumed that no people were vaccinated before $t = t_0$, the time when vaccine production starts if there was no early capacity investment. At time T , country i starts receiving V_{NP} courses of vaccine per month. The fraction of its population that is vaccinated is

$$\lambda_{NP}(t) = \begin{cases} 0 & \text{if } t \leq t_0 \\ \frac{V_{NP}}{P}(t - t_0) & \text{if } t > t_0 \end{cases}$$

If, instead, country i obtains an effective capacity of V , it starts vaccinating V people per month at time $t = t_o - T$. T measures how much earlier capacity is available with early capacity investment. At time t_0 capacity ramps up to V_{NP} . Thus, the fraction of the population that is vaccinated is

$$\lambda(t; V) = \begin{cases} 0 & \text{if } t \leq t_0 - T \\ \frac{V}{P}(t - t_0 + T) & \text{if } t_0 - T < t \leq t_0 \\ \frac{V}{P}T + \frac{V_{NP}}{P}(t - t_0) & \text{if } t > t_0 \end{cases}$$

We then translated the fraction vaccinated $\lambda_i(t)$ into benefits per unit of time. Let H_i be the monthly health and economic harm due to COVID-19 for country i . It is the sum of total economic and health (mortality) harm. We based our estimates of economic harm on estimates from the World Bank (5%-20% of GDP, World Banks Global Economic Prospects 2020).⁵ Our estimate of health benefits was based on estimates of the mortality rate, statistical value of a life, and years lost per COVID-19 death. Health harm is the product of (a) mortality, which we assumed was 200,000 per month distributed across the world in proportion to population, (b) the value of a statistical life, which is proportional to GDP per capita and is \$7 million for the US, and (c) the fraction of one life that is lost on average due to COVID-19 deaths, $\frac{10}{71}$, which assumes that each death implies a loss of 10 years and that there is a life expectancy of 71 years.

We assumed that the health and economic benefits of vaccination at time t , relative to no vaccination, are given by $\delta H_i f_i(\lambda_i(t))$. $\delta \in [0,1]$ is a factor by which we discounted benefits because of the possibility of improved treatments, mitigation strategies such as effective contact tracing, etc. that would preclude an important share of the benefits from vaccination. In our model as of August 2020, we assumed that by the time the vaccine arrived, 50% harm would have been mitigated due to these other factors. In retrospect, that estimate was too optimistic for the situation as it stood in late 2020 when the first vaccines received approval;⁶ as such our results understate

⁵ These estimates are more conservative than in Cutler and Summers (2020), which calculated that the costs of the COVID-19 pandemic in the US alone may reach \$16 trillion, which represents about 18% of world GDP in US dollars in 2019 (World Bank Data, 2020).

⁶ As of July 2020, it appeared that many areas had “flattened the curve,” economic re-opening had occurred in many parts of the world, and treatment protocols had improved, improving patient outcomes. But by December 2020, when mass vaccinations began in multiple high-income countries, infection rates were at an all-time high, many

the benefits of investments to accelerate vaccine success. The function $f_i(\cdot)$ measures the fraction of harm that is avoided. It is a continuous function satisfying (a) no benefit is obtained if no one is vaccinated and (b) all economic harms are relieved if all are successfully immunized.

Countries generally first distribute doses to high-priority populations (especially the elderly) since this results in the greatest reduction in mortality for a limited vaccine supply, as can be shown by using simple epidemiological models (Bubar et al. 2021). The same models predict that the reductions in infections are roughly linear in the proportion vaccinated between 0 and the point where fraction vaccinated approaches herd immunity. If the vaccinations are prioritized according to age only, the reductions in infections are slightly convex (Bubar et al. 2021, Fig 3B); however, since high-risk populations also include individuals at high risk of transmission (such as health care workers), we hypothesize that the reductions in infections will be linear. It was unknown whether reductions in economic losses from COVID-19 will more closely track reductions in mortality or reductions in infections. It was also unknown whether the efficacy in preventing severe infection shown by vaccines that were in development would translate into efficacy in preventing transmission. To accommodate this uncertainty, we specified

$$f_i(\cdot) = \rho f_i^l(\cdot) + (1 - \rho) f_i^{nl}(\cdot),$$

a weighted mean of two functions: $f_i^l(\cdot)$, which is a simple linear function of $\lambda_i(t)$, and $f_i^{nl}(\cdot)$, which is a nonlinear function capturing averted mortality. Weight ρ can take on any value in the unit interval, but for simplicity we set $\rho = 0.5$.

A country uses its initial vaccines for its high-priority population, providing θ_i times the benefit of vaccinating a non-priority person, where $\theta_i \in [5,10]$ is a parameter that is linear in the per-capita GDP in country i , varying between 5 for the lowest-income country (Burundi) and 10 for the highest-income country (Monaco). We based these values on epidemiological models of mortality reduction, which suggest that over 80% of mortality reductions are obtained from vaccinating the first 20% of the population (Bubar et al. 2021), consistent with empirical data on age-specific mortality rates in developed countries. The lower value for lower-income countries reflects the fact that relative mortality risk for older vs younger individuals is less steep in these countries, by as much as a factor of three (Demombynes 2020). However, the situation in January 2021 in many low-income countries, with lower overall prevalence levels, may lead the optimal

countries and US states imposed new restrictions, and hospital systems became overwhelmed. Debate remained about the effectiveness of antibody treatments, and these were not in widespread use.

policy to reduce mortality to depart from vaccinating the elderly first to perhaps vaccinating working-age adults first (Hogan et al. 2020).


We assumed that the benefit function has a kink at the threshold λ' at which all high-priority people have been vaccinated and the vaccine begins to be distributed to others. We defined two higher kink points in the benefit function: $\lambda'' = 0.4$, and $\lambda''' = 0.7$. We assumed that at λ'' , the slope of the benefit function falls in half. Between λ'' and λ''' , benefits increase linearly at the lower rate until the threshold for herd immunity, λ''' , is reached. Finally, we assumed that all harm is averted at this threshold and higher levels of vaccination: $f_i(\lambda) = 1$ for all $\lambda > \lambda'''$. Of course, the slope of the benefit function equals zero above threshold λ''' . To understand the rationale for incorporating the additional kink points, note that a simple epidemiological model puts the threshold for herd immunity at 60%. However, a number of factors suggest that full benefits may be obtained above or below this theoretical threshold. Factors pushing the threshold down include (a) preexisting immunity or lower susceptibility in younger individuals, obviating a need for them to be vaccinated (Davies et al. 2020); (b) high levels of acquired immunity, especially in high-income countries; (c) heterogeneity in spread, leading herd immunity to be reached earlier than a simple epidemiological model with homogeneous agents would predict (Britton et al. 2020).

The net benefit from effective capacity V is given by

$$b_i(V) = \delta H_i \int f_i(\lambda(t; V)) dt - \delta H_i \int f_i(\lambda_{NP}(t)) dt$$


Figure A.2 illustrates our assumptions about flow benefits as a fraction of the population vaccinated, where λ is the fraction vaccinated and $f_i(\lambda)$ is the flow benefit. We assumed that the function increases steeply as the initial priority groups are vaccinated (e.g. health care workers, since health system capacity is a limiting factor for opening economies) and then increases more slowly thereafter. Figure A.3 illustrates the vaccination schedule with and without advance capacity investment. Without it, vaccination starts at rate V_{NP} at time t_0 . With advance investment, vaccination starts at time $t_0 - T$. The initial vaccination rate is V , the total effective capacity from the portfolio, and ramps up to rate V_{NP} at time t_0 . Figure A.4 illustrates how these two vaccination schedules translate into flow benefits. The net benefits from advance capacity investment are shown as the area between both curves, which is shaded in yellow.

Figure A.2 Flow Benefits to Vaccination




Note: This plot shows the function that measures the fraction of benefits as a function of the fraction of the population that has been vaccinated, for a country with GDP per capita of \$17,000 and 15% high-risk population. The slope is highest before λ' because the country vaccinates the high-priority population, obtaining large benefits per person. Benefits increase more slowly as the rest of the population is vaccinated.

Figure A.3 Vaccination schedules with and without acceleration



Note: The green line shows the fraction of the population vaccinated as a function of time if vaccination is not accelerated. Vaccination starts at a constant rate at time t_0 . The red line shows what happens with acceleration. Vaccination starts at an earlier time, $t_0 - T$, at a rate determined by the outcome of the vaccine portfolio. Then vaccination ramps up at time t_0 to the same rate that would take place without acceleration.

Figure A.4 Vaccination benefits over time with and without acceleration



Note: This figure shows the benefits per unit time that are obtained from the vaccination schedules in Figure A.3 above. The net benefits from the portfolio are equal to the shaded area between the green and red curves.

A.1.3 Countries' optimal portfolio problem

The benefits country i gets if it chooses a portfolio v_i are given by

$$B_i(v_i) = \mathbb{E}[b_i(V)|v_i]$$

The benefit $b_i(V)$ that arises if the effective capacity is V is integrated over the distribution of effective capacity that is generated by portfolio v_i .

On the cost side, there are a variety of potential cost functions for capacity.

$$\max_{v_i} B_i(v_i) - C(v_i)$$

For simplicity in our base case, we assumed there was a single price p per unit of capacity across all candidates;⁷ in Online Appendix A.3 we consider convex capacity costs. Then, the country's problem is

$$\max_{v_i} B_i(v_i) - p \sum_j v_{ij}$$

We make available our code for portfolio optimization given inputs so that the computation can be updated over time, and so that the code could be used as a starting point in a future pandemic. Our results below illustrate optimal investments using the assumptions we made as of August 2020.

⁷ This assumption is likely applicable to the problem for a small open economy.

A.1.3 Main Results

We compute the optimal portfolios for each country in the world for a constant price of \$10 per vaccine course per year.^{8,9} We found that early investments in vaccine manufacturing capacity would have large net benefits for high-, middle-, and even low-income countries (Table A.1). Summing over individual countries, the optimal portfolio for the world consisted of a total investment of sufficient capacity for 2.3 billion courses per month, of which in expectation, 0.5 billion would have been successful. The expected benefit for the world was \$137 per capita, while the cost was \$37 per capita. The portfolio which maximized net benefits varied substantially between countries. It was optimal for higher-income countries to purchase more doses per capita and invest in a larger portfolio of candidates. This result is explained by the fact that a large portion of the estimated benefits of a vaccine come from averting losses to GDP. For example, the optimal portfolio for the US included investment in 27 candidates at a total of 462 million courses per month, while Chile's optimal portfolio included investment in 12 candidates at a three-times lower level of capacity investment per capita. It was optimal for each country to invest in differing amounts across candidates to maximize expected successful capacity by investing more in candidates with higher probabilities of success and by diversifying across platforms.

Table A.1: Baseline Optimal Portfolio

Country	Mean Number of Candidates	Total Capacity (mn. courses / mth)	Expected Effective Capacity (mn. courses / mth)	Total Capacity (courses / mth per 1000 pop.)	Expected Benefits (per cap.)	Total Cost (per cap.)
World	8.82	2290.05	538.87	304.40	137.41	36.53
High Income	18.26	1418.03	307.97	1196.54	699.25	143.58
Middle Income	6.73	906.88	239.27	170.21	40.71	20.43
Low Income	1.26	2.33	0.61	2.18	0.58	0.26
United States	27.00	462.30	97.97	1415.06	923.36	169.81
E.U.	17.00	477.58	105.12	1093.85	603.46	131.26

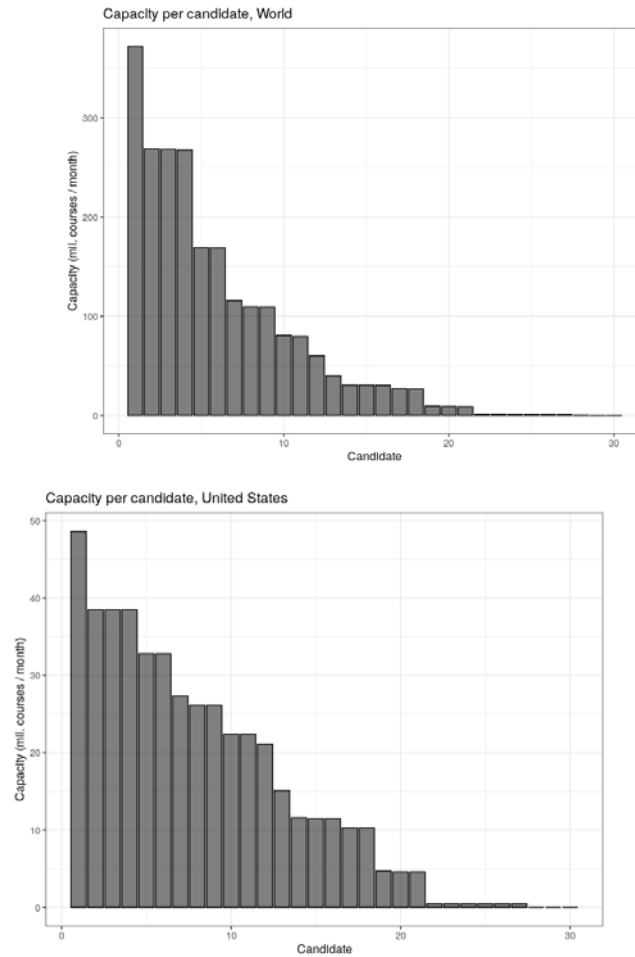
⁸ The AstraZeneca deal with the US was for 300 million doses for a total price of \$1.2 billion. Suppose that capacity costs are 75% of the price and zero profit. We can calculate that the marginal cost of production is \$1 per dose and the cost of production is \$3 per dose per year assuming that all doses will be produced over 2021. Then the cost of production for AstraZeneca would be $(\$1 \text{ per dose}) \times (300 \text{ m doses}) + (\$3 \text{ per dose/year}) \times (300 \text{ m doses/year}) = \1.2 bn , consistent with the deal they signed. At 2 doses per course, that calculation implies a total cost of production of \$6 per course / year. Some part of the \$1.2 bn actually went to fund the clinical trials, so the cost of production could be even lower. If we take AstraZeneca to be one of the cheapest vaccines (by a factor of three) to produce, then it seems reasonable to assume that the cost of production for other vaccine candidates likely ranges between \$4-\$20 per course per year.


⁹ Kis et al. (2021) estimate that building production capacity for enough doses to vaccinate the entire world within a year for an mRNA vaccine would cost less than \$4 billion.

Germany	21.00	113.30	24.22	1366.61	855.50	163.99
U.K.	21.00	85.30	18.41	1283.47	763.14	154.02
Canada	21.00	45.90	9.93	1238.61	719.27	148.63
New Zealand	18.00	5.80	1.27	1198.10	670.71	143.77
Australia	21.00	34.50	7.37	1380.96	879.99	165.71
Chile	12.00	10.90	2.67	581.98	183.03	69.84
Israel	19.00	10.20	2.24	1148.29	633.38	137.79
Hong Kong	20.00	9.40	2.02	1261.58	740.41	151.39
Japan	18.00	129.50	28.96	1023.72	494.81	122.85

Note: This table presents the optimal investment portfolio of various countries/coalitions. We assumed \$10 per vaccine course per year, vaccine availability is accelerated by 3 months, and baseline vaccine success probabilities.

Figure A.5: Distribution of investment across candidates





Note: This figure presents the optimal investment by candidate for the World, the United States and Chile. We assumed \$10 per vaccine course per year, vaccine availability is accelerated by 3 months, and baseline vaccine success probabilities.

A large portion of the gains to a country come from the first candidates and doses; we estimated that spending half the ideal budget would have produced about three quarters of the benefits (Table A.2). Specifically, the expected benefits for the world were \$109 per capita with half the budget versus \$137 per capita with the full budget. There are two reasons behind this result. First, effective capacity has diminishing marginal returns. The time to vaccinate the whole population is $Time = Population / Capacity$. Doubling capacity reduces the time to vaccinate the whole population by one half. Doubling it again reduces the time to vaccinate only by one fourth of the original time. Second, the likelihood of at least one successful vaccine increases much more with the first few candidates than later candidates (Figure A.6). Diminishing returns in the capacity of each successful vaccine implies that countries should have invested in more than candidates, instead of greater capacity for a given candidate. Because the returns to the first candidates and doses were so high, even poorer countries would benefit from purchasing at least some.


Table A.2: Portfolio spending half the optimal budget

Country	Mean Number of Candidates	Total Capacity (mn. courses / mth)	Expected Effective Capacity (mn. courses / mth)	Total Capacity (courses / mth per 1000 pop.)	Expected Benefits (per cap.)	Total Cost (per cap.)
World	5.41	1145.02	267.94	145.43	109.27	18.27
High Income	11.15	709.01	172.53	602.49	579.38	71.79
Middle Income	4.40	453.44	130.40	85.10	28.52	10.22
Low Income	1.19	1.16	0.30	1.09	0.42	0.13
United States	18.00	231.15	54.74	707.53	774.92	84.91
E.U.	10.00	238.79	57.53	546.92	491.08	65.63

Germany	13.00	56.65	13.47	683.31	710.96	82.00
U.K.	12.00	42.65	10.10	641.74	625.85	77.01
Canada	12.00	22.95	5.42	619.30	585.21	74.32
New Zealand	11.00	2.90	0.69	599.05	541.87	71.89
Australia	14.00	17.25	4.08	684.47	734.65	82.86
Chile	11.00	5.45	1.71	290.99	154.48	34.92
Israel	11.00	5.10	1.25	574.14	515.48	68.90
Hong Kong	12.00	4.70	1.16	644.21	615.90	75.70
Japan	11.00	64.75	15.74	498.02	389.15	61.43

Note: This table presents the investment portfolios spending half the optimal budget for various countries/coalitions. We assumed \$10 per vaccine course per year, vaccine availability is accelerated by 3 months, and baseline vaccine success probabilities.

Figure A.6: Probability of at least one successful vaccine



Note: This figure presents the probability of at least one successful vaccine by the number of candidates a country/coalition chooses to invest in.

A.1.5 Comparative Statics

We now consider how our results change under alternative assumptions.

Number of Months of Acceleration

We first considered variation in the number of months the early investment program saves. Our baseline analysis considers a three-month acceleration in vaccine production capacity. Early in the pandemic, experts expressed substantial skepticism and uncertainty about how quickly manufacturing could be brought to scale (Dunn 2020; Khamsi 2020; Thompson 2020). Accelerating by six months was unprecedented, even though some vaccines used relatively

standard processes and had access to contract manufacturing facilities, and there was expected to be at least a six-month wait for approval. In addition, our model suggested investments in acceleration on a world scale of hundreds of billions of dollars, more than had ever been considered. Thus, we estimated that early investment might result in a six-month acceleration. In Table A.3, we show how the optimal portfolio at a price of \$10 per vaccine course per year changes when the acceleration is six months rather than three. We found that all countries invest in a greater number of candidates and a larger total amount of capacity. This is because the expected benefits per unit of capacity were much greater.

Table A.3: Optimal Portfolio - 6 months Acceleration

Country	Mean Number of Candidates	Total Capacity (mn. courses / mth)	Expected Effective Capacity (mn. courses / mth)	Total Capacity (courses / mth per 1000 pop.)	Expected Benefits (per cap.)	Total Cost (per cap.)
World	12.85	2742.96	610.87	364.6	224.96	43.75
High Income	23.54	1464.53	295.41	1235.77	1071.64	148.29
Middle Income	11.18	1288.95	316.16	241.92	82.07	29.03
Low Income	3.25	28.47	7.94	26.76	5.18	3.21
United States	30.00	476.00	93.69	1456.99	1407.00	174.84
E.U.	22.19	495.71	101.26	1135.35	928.17	136.24
Germany	30.00	117.00	23.19	1411.24	1298.85	169.35
U.K.	30.00	87.90	17.57	1322.59	1167.3	158.71
Canada	28.00	47.30	9.49	1276.39	1105.31	153.17
New Zealand	25.00	6.10	1.23	1260.07	1034.79	151.21
Australia	30.00	35.50	7.00	1420.98	1343.33	170.52
Chile	18.00	11.80	2.64	630.03	305.95	75.60
Israel	28.00	10.60	2.13	1193.32	981.86	143.20
Hong Kong	28.00	9.60	1.92	1288.42	1134.87	154.61
Japan	28.00	133.80	27.84	1057.71	764.60	126.92

Note: This table presents the optimal investment portfolio of various countries/coalitions. We assumed \$10 per vaccine course per year, vaccine availability is accelerated by 6 months, and baseline vaccine success probabilities.

Higher Success Probabilities

We next consider how optimal portfolios change with assumptions about success probabilities. Evaluating the success probabilities with the benefit of hindsight in January, 2021, the model looks modestly pessimistic on the probabilities of success. For example, our model predicted multiple failures were likely, but only one vaccine that received substantial investment has definitely failed. The Australian government had contracted (at government risk) \$1 billion dollars for 51 million doses of the University of Queensland vaccine (Johnson and Whitley 2020). The Queensland vaccine was abandoned after entering clinical trials because a side-effect of the vaccine was the production of antibodies that looked like HIV antibodies, leading to incorrectly signaling on

standard tests that people inoculated with the vaccine had HIV. We expected multiple failures of this kind. On the other hand, most vaccines will not be scaled up to be worthwhile in the pandemic and thus can be said to have failed. Even the vaccines that have been successful at establishing safety and efficacy have had to scale back planned deliveries. For example, Pfizer originally projected that it could deliver 40 million doses in 2020 but only delivered half that, and it shut down a European manufacturing facility for several weeks in order to prepare for a capacity expansion, delaying deliveries it had promised to European countries (Peel, Milne and Mancini, 2021). Johnson and Johnson's vaccine is yet to be approved, but if approved was contracted to deliver 12 million doses by the end of February. However, manufacturing problems will make scaling up slower (Zimmer et al. 2021). Our model predicted failure even with very substantial advance investment, but what we observed was failure with less advance investment. It is hard to assess what difference investment would have made for some of these candidates, and in particular, whether the lack of investment was due to poor prospects of efficacy, or whether additional investment, e.g. in supporting clinical trials, would have led to more viable candidates succeeding and scaling. It is difficult to believe, however, that additional funding of substantial magnitude, say \$10 billion, would not have made a difference. Moreover, it is clear that funding in complementary infrastructure such as vaccination clinics would have been very valuable. Although we focus on manufacturing capacity, the high value we estimated for vaccines suggests the value of investing everywhere along the supply chain. Only successfully delivered and vaccinated doses produce social value.

In Table A.4, we analyzed the optimal portfolio for the case where success probabilities were higher. Specifically we considered the case where the vaccines were half as likely to fail at the candidate-level. We found that all countries should have invested in more candidates, and that middle- and lower-income countries invest in more capacity (because of higher expected capacity from investing), while high-income countries invest less (because of diminishing returns to successful capacity). Specifically, the high-income countries invest in 1066 mn. courses per month, down from 1418 mn. courses per month. On the other hand, middle income countries invest in 1061 mn. courses per month, up from 907 mn. courses per month. It is clear that the relationship between total vaccine capacity investment and probability of success is non-monotonic. At a probability of success of zero for all vaccines, countries would invest nothing. At a probability of success of one for all vaccines, countries at most would invest exactly enough to vaccinate their

whole population in the first month. At probabilities between zero and one, countries (for example the US in Table A.4) may invest in more capacity than enough to vaccinate their whole population in the first month because expected effective capacity is much lower. The expected benefits were higher across countries from high success probabilities regardless of income level.

Table A.4: Optimal Portfolio - High Probability of Success for all candidates

Country	Mean Number of Candidates	Total Capacity (mn. courses / mth)	Expected Effective Capacity (mn. courses / mth)	Total Capacity (courses / mth per 1000 pop.)	Expected Benefits (per cap.)	Total Cost (per cap.)
World	12.84	2107.36	803.72	280.12	162.08	33.61
High Income	22.24	1066.33	394.71	899.77	778.45	107.97
Middle Income	13.01	1061.10	415.84	199.15	58.14	23.90
Low Income	2.35	12.59	5.34	11.83	2.10	1.42
United States	25.00	336.50	124.25	1030.00	1015.81	123.60
E.U.	21.42	366.26	135.78	838.87	676.75	100.66
Germany	25.00	83.10	30.71	1002.34	944.90	120.28
U.K.	25.00	63.10	23.33	949.44	846.95	113.93
Canada	25.00	34.30	12.69	925.58	801.07	111.07
New Zealand	24.00	4.40	1.64	908.90	750.53	109.07
Australia	25.00	25.20	9.30	1008.70	970.01	121.04
Chile	24.00	10.00	3.75	533.93	227.47	64.07
Israel	24.00	7.80	2.90	878.10	711.52	105.37
Hong Kong	24.00	7.00	2.59	939.47	823.23	112.74
Japan	24.00	100.50	37.29	794.47	562.85	95.34

Note: This table presents the optimal investment portfolio of various countries/coalitions with higher vaccine success probabilities. We assumed that vaccines are half as likely to fail at the candidate-level than according to our baseline probabilities. We assumed \$10 per vaccine course per year and that vaccine availability is accelerated by 3 months.

By fall or winter 2020, the uncertainty had diminished on safety and efficacy for leading candidates, but risks remained on manufacturing, particularly for the mRNA candidates. One leading candidate, the Oxford/AstraZeneca vaccine, was also showing promising results, but still significant risks in timing; in particular, it faced setbacks in its trial and uncertainty about the timing of approval (Robbins and Mueller 2020). We now know that the mRNA platform was successful with safety and efficacy and partially successful with scaling manufacturing. Here, we considered a scenario where we assumed a 0.8 probability of RNA platform success instead of 0.6 in our base case. We found very similar results, with more investment into mRNA vaccines and slightly higher expected benefits (Table A.5).

Table A.5: Optimal Portfolio - High Probability of Success for RNA

Country	Mean Number of Candidates	Total Capacity (mn. courses / mth)	Expected Effective Capacity (mn. courses / mth)	Total Capacity (courses / mth per 1000 pop.)	Expected Benefits (per cap.)	Total Cost (per cap.)
World	9.00	2287.77	576.87	304.10	141.34	36.49

High Income	18.51	1375.50	321.55	1160.65	711.87	139.28
Middle Income	6.79	946.12	264.13	177.57	43.52	21.31
Low Income	1.53	2.53	0.69	2.37	0.62	0.28
United States	24.00	447.50	102.04	1369.76	938.61	164.37
E.U.	16.85	464.31	110.03	1063.45	615.07	127.61
Germany	24.00	109.70	25.20	1323.19	870.02	158.78
U.K.	21.00	82.60	19.13	1242.85	776.52	149.14
Canada	21.00	44.50	10.34	1200.83	732.28	144.10
New Zealand	21.00	5.60	1.31	1156.79	682.64	138.81
Australia	24.00	33.30	7.63	1332.92	894.26	159.95
Chile	11.00	10.70	2.83	571.30	189.06	68.56
Israel	21.00	10.00	2.37	1125.77	646.75	135.09
Hong Kong	21.00	9.10	2.11	1221.31	753.42	146.56
Japan	21.00	125.90	30.27	995.26	505.13	119.43

Note: This table presents the optimal investment portfolio of various countries/coalitions with higher vaccine success probabilities for RNA candidates. We assumed a 0.8 probability of RNA platform success instead of 0.6 in our baseline. We assumed \$10 per vaccine course per year and that vaccine availability is accelerated by 3 months.

Higher Correlation

Another important parameter for our portfolio optimization problem is the correlation among outcomes for different candidates. In Table A.6, we analyze how the portfolio choices change when we increase the correlations within sub-platform by doubling the sub-platform probability of failure at the candidate level, while halving the probability of failure at the candidate level to keep overall probabilities constant. We found that countries increase the number of candidates they invest in to diversify their portfolios.

Table A.6: Optimal Portfolio - High Sub-Platform Correlation

Country	Mean Number of Candidates	Total Capacity (mn. courses / mth)	Expected Effective Capacity (mn. courses / mth)	Total Capacity (courses / mth per 1000 pop.)	Expected Benefits (per cap.)	Total Cost (per cap.)
World	13.57	2292.76	636.43	304.76	151.37	36.57
High Income	24.78	1278.61	344.03	1078.90	750.71	129.47
Middle Income	13.14	1049.27	301.99	196.93	49.28	23.63
Low Income	1.89	2.77	0.82	2.60	0.69	0.31
United States	30.00	406.50	108.74	1244.26	983.97	149.31
E.U.	23.88	437.22	118.02	1001.40	650.90	120.17
Germany	30.00	100.20	26.86	1208.60	913.93	145.03
U.K.	27.00	76.00	20.44	1143.54	817.76	137.22
Canada	27.00	41.20	11.09	1111.78	772.53	133.41
New Zealand	26.00	5.30	1.42	1094.82	723.35	131.38
Australia	29.00	30.40	8.13	1216.84	938.80	146.02
Chile	23.00	11.40	3.17	608.68	210.37	73.04
Israel	27.00	9.40	2.53	1058.22	685.34	126.99
Hong Kong	27.00	8.40	2.25	1127.37	794.00	135.28
Japan	27.00	119.90	32.45	947.83	538.66	113.74

Note: This table presents the optimal investment portfolio of various countries/coalitions with higher vaccine success correlation within sub-platform. We double the sub-platform probability of failure at the candidate level, and half the probability of failure at the candidate level to keep overall probabilities constant. We assumed \$10 per vaccine course per year and that vaccine availability is accelerated by 3 months.

Different Price Levels

Finally, we tested the sensitivity of our results to the assumed price level. In our base case, we assumed a constant price of \$10 per course/year. In Table A.7, we present our results for prices between \$4 and \$40. As prices increase, all countries would have invested in fewer candidates and in less capacity. However, it still made sense to invest in capacity at-risk for most countries even at much higher prices, especially high- and middle-income countries. Low-income countries would have invested much less at higher prices.

Table A.7: Optimal Portfolio - Different Price Levels

Country	Mean Number of Candidates	Total Capacity (mn. courses / mth)	Expected Effective Capacity (mn. courses / mth)	Total Capacity (courses / mth per 1000 pop.)	Expected Benefits (per cap.)	Total Cost (per cap.)
Panel A: Price = \$4						
World	13.91	4771.07	1055.71	634.19	162.10	30.44
High Income	25.19	2447.74	484.35	2065.41	765.21	99.14
Middle Income	12.31	2334.47	570.45	438.15	60.65	21.03
Low Income	3.77	55.78	15.53	52.42	4.04	2.52
United States	30.00	790.30	152.82	2419.04	999.83	116.11
E.U.	23.88	830.38	166.39	1901.87	664.80	91.29
Germany	30.00	194.30	37.76	2343.62	929.75	112.49
U.K.	30.00	146.80	28.77	2208.84	833.14	106.02
Canada	30.00	79.10	15.57	2134.51	787.26	102.46
New Zealand	28.00	10.00	1.97	2065.69	736.45	99.15
Australia	30.00	58.90	11.43	2357.63	954.32	113.17
Chile	18.00	20.30	4.52	1083.87	221.27	52.03
Israel	28.00	17.80	3.52	2003.87	698.64	96.19
Hong Kong	28.00	16.30	3.20	2187.63	810.02	105.01
Japan	28.00	225.50	45.79	1782.61	551.98	85.57
Panel B: Price = \$10						
World	8.82	2290.05	538.87	304.40	137.41	36.53
High Income	18.26	1418.03	307.97	1196.54	699.25	143.58
Middle Income	6.73	906.88	239.27	170.21	40.71	20.43
Low Income	1.26	2.33	0.61	2.18	0.58	0.26
United States	27.00	462.30	97.97	1415.06	923.36	169.81
E.U.	17.00	477.58	105.12	1093.85	603.46	131.26
Germany	21.00	113.30	24.22	1366.61	855.50	163.99
U.K.	21.00	85.30	18.41	1283.47	763.14	154.02
Canada	21.00	45.90	9.93	1238.61	719.27	148.63
New Zealand	18.00	5.80	1.27	1198.10	670.71	143.77
Australia	21.00	34.50	7.37	1380.96	879.99	165.71
Chile	12.00	10.90	2.67	581.98	183.03	69.84
Israel	19.00	10.20	2.24	1148.29	633.38	137.79
Hong Kong	20.00	9.40	2.02	1261.58	740.41	151.39
Japan	18.00	129.50	28.96	1023.72	494.81	122.85
Panel C: Price = \$20						
World	5.73	1212.82	295.95	161.21	113.24	38.69
High Income	13.56	910.41	211.76	768.20	626.47	184.37
Middle Income	3.40	322.99	89.58	60.62	22.32	14.55
Low Income	0.39	0.64	0.16	0.60	0.33	0.14
United States	18.00	303.50	68.67	928.99	840.96	222.96
E.U.	12.85	302.67	71.36	693.22	535.26	166.37
Germany	18.00	74.20	16.93	894.99	775.51	214.80

U.K.	18.00	55.40	12.76	833.58	686.78	200.06
Canada	18.00	29.80	6.91	804.15	645.43	193.00
New Zealand	12.00	3.70	0.87	764.30	596.52	183.43
Australia	18.00	22.50	5.11	900.62	798.41	216.15
Chile	9.00	5.90	1.56	315.02	137.27	75.60
Israel	14.00	6.60	1.55	743.01	563.71	178.32
Hong Kong	17.00	6.10	1.39	818.68	665.34	196.48
Japan	12.00	81.00	19.64	640.32	429.55	153.68

Panel D: Price = \$40

World	3.40	589.33	149.82	78.34	85.93	37.60
High Income	8.97	536.70	134.71	452.87	519.25	217.38
Middle Income	1.27	60.79	17.46	11.41	6.82	5.48
Low Income	0.21	0.30	0.08	0.29	0.22	0.14
United States	12.00	186.90	45.86	572.08	719.74	274.60
E.U.	7.73	173.33	43.95	396.99	434.55	190.55
Germany	11.00	45.40	11.21	547.61	657.92	262.85
U.K.	11.00	33.50	8.35	504.06	575.20	241.95
Canada	11.00	18.00	4.50	485.73	538.17	233.15
New Zealand	9.00	2.20	0.56	454.45	489.64	218.14
Australia	11.00	13.70	3.38	548.38	677.86	263.22
Chile	4.00	2.10	0.60	112.12	69.12	53.82
Israel	11.00	3.90	0.99	439.05	460.58	210.74
Hong Kong	11.00	3.70	0.93	496.58	557.00	238.36
Japan	9.00	45.60	11.95	360.47	333.63	173.03

Note: This table presents the optimal investment portfolio of various countries/coalitions for various prices: \$4, \$10, \$20 and \$40. We assumed vaccine availability is accelerated by 3 months, and baseline vaccine success probabilities.

A.1.6 Benefits of Actual Versus Recommended Investments

Overall, we can compare the investments recommended by our baseline model as of August 2020 with the investment that had been undertaken by 2021, as well as to a counterfactual scenario where capacity was only created after regulatory approval. For the latter scenario, we assumed capacity would have been delayed three months, which is a very ambitious timeline for scaling capacity after approval (it would be consistent with a scenario where only the most expensive investments were delayed until after approval, or where investments began while a company had access to early trial data prior to approval).

We report results comparing these scenarios in Table 1 of the main paper, reproduced here for easy reference. Specifically, we compute the benefits of vaccination schedules for the US and for the world under three different scenarios:

1. *Actual advanced capacity investment:* The US gets 37.5 million vaccine courses per month, which is the total capacity it has contracted with Oxford/AstraZeneca, Pfizer/BioNTech, and Moderna (Duke Global Health Innovation Center, 2020), assuming those vaccines will be distributed over the course of one year. The world gets 312.5 million courses doses per month, which is roughly the sum of the capacity announced by Oxford/AstraZeneca

(AstraZeneca 2020), Pfizer/BioNTech (Pfizer 2020), Moderna (2020), Sinovac (Reuters 2021), and Sputnik V (Mullard 2020) for 2021, assuming it will be distributed evenly over 12 months. In both cases, vaccines start being delivered on January 1, 2021.

2. *Zero advanced capacity investment:* The US and the world get the same number of vaccines as in the actual capacity investment scenario, but vaccines start being delivered on April 1, 2021.
3. *Recommended advanced capacity investment:* Vaccines start being delivered on January 1, 2021, but the number of vaccines per month are based on the recommended portfolios from our model at a price of \$20 per vaccine course per year.¹⁰ For the US, we add the capacities our model recommends for Oxford/AstraZeneca, Pfizer/BioNTech and Moderna. For the world, we add the capacities our model recommends for Oxford/AstraZeneca, Pfizer/BioNTech, Moderna, Sinovac, and Sputnik V for every country in the world.

For each of the scenarios above, we compute the time when the US and the world will finish vaccinating 70% of the population. As described above, that assumes vaccination would have started on April 1, 2021 without at-risk investment, and on January 1, 2021 with at-risk investment.

Table 1: Advance capacity expansion effects

Advance capacity investment	At-Risk Capacity of Approved Vaccines (bn courses)	Benefits relative to Zero At-Risk (\$bn)	70% of Vaccination complete by
Panel A, US			
Recommended	1.05	556.9	Mar 2021
Actual	0.45	389.9	Jul 2021
Zero	—	—	Oct 2021
Panel B, World			
Recommended	7.12	2748.7	Oct 2021
Actual	3.75	1606.4	Jun 2022
Zero	—	—	Sep 2022

¹⁰ This price is higher than the price of \$10 per course per year from our baseline calculations from Section A.1.3, which give recommendations to individual countries facing an external price. If all countries in the world followed the recommendations from our model, marginal costs are likely to increase and prices rise (see the discussion in Section A.2). Thus, we use a higher price of \$20 per course per year.

We compute benefits generated in each scenario using the benefits function described in Section A.1.2 for each country. When computing benefits for the world, we assumed that in the “actual” and “zero” capacity investment scenarios 40% of capacity goes to high-income countries, 25% goes to upper-middle-income countries, and 25% goes to lower-middle-income countries. Within each income group, capacity is distributed by population, meaning that the largest country within the income group is the first to receive vaccines sufficient to serve 70% of its population, the remainder then goes to the next largest country, and so on. The remaining 10% is distributed evenly by population throughout the world (by COVAX, for instance). As soon as high-income countries vaccinate 70% of their population, their capacity is distributed by population across the world. Then, as soon as upper-middle-income countries vaccinate 70% of their population, their capacity is distributed by population throughout countries that have not vaccinated 70% of their population, and so forth. For the “recommended” capacity investment scenario, we assumed that 90% of the capacity is distributed across the four income groups (high, upper-middle, lower-middle, and low) according to the total demand for vaccines by group. Within each group, vaccines are distributed by population. The remaining 10% is distributed evenly by population throughout the world. As income groups finish vaccinating their population, their capacity is distributed to the rest of the world as described above for the “actual” and “zero” capacity investment scenarios.

The benefits shown in the table are measured relative to the “zero” capacity investment scenario. The “actual” capacity investment scenario has the same vaccination schedule, but it is shifted by three months. Thus, benefits are equal to the monthly harm from the pandemic multiplied by three. Relative to the “zero” scenario, the “recommended” scenario has greater benefits both because vaccination starts three months earlier and because capacity is larger, so people are vaccinated earlier relative to the time vaccination starts.

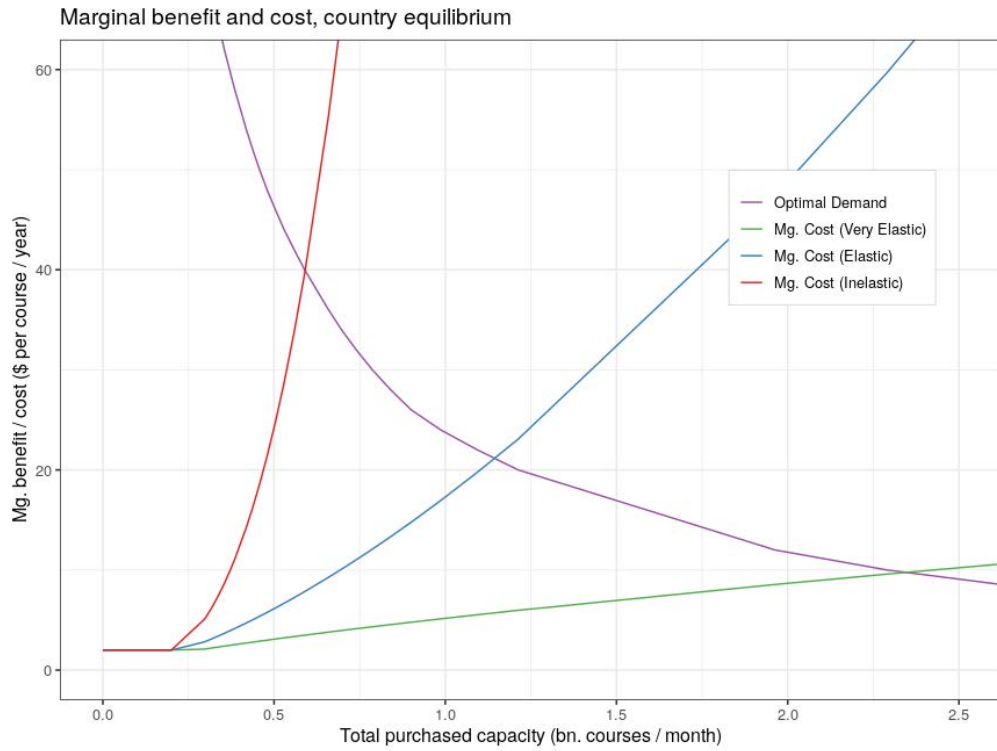
A.2 Global Supply and Demand

Assuming that countries invest in the optimal portfolio according to our model, we can calculate a global demand curve for vaccine capacity. For any given price, the demand curve gives the total monthly capacity demanded by all countries across all vaccine candidates. We can also compute a global supply curve, which specifies the amount of monthly capacity that would have been

supplied if firms were reimbursed at a given capacity cost. The global capacity for pandemic vaccine doses is large as considerable capacity can be repurposed from flu vaccines (McLean et al. 2016) but given the large upfront costs and challenges of installing capacity, we would expect the marginal cost to rise sharply above a certain level. As such, we modeled the global short-run vaccine supply as constant cost up to 200m doses a month, but relatively inelastic after that point (with elasticity of $1/3$, meaning that to induce a 1% increase in monthly capacity, the reimbursement must rise by 3%).


Under this approach, the market clearing price for capacity would have been around \$40 per vaccine course per year (Figure A.7). At prices of this magnitude, a large share of low-income countries would be priced out of the market (Table A.7) and low-cost firms would have accrued large rents. However, during 2020, vaccine deals were signed at prices substantially lower than this amount. So far, prices have remained lower than what our model indicates they would be in a market equilibrium. There may be political or ethical constraints which are holding prices down. In the presence of a price ceiling, there would be a large shortage in supply, and capacity would be allocated by speed of contracting or political influence, such as having manufacturing capability located in a particular country (Figure A.8). We also note that a price ceiling for vaccines will further translate into lower prices paid to suppliers, which in turn leads to shortages in the supply chain, unless countries use emergency powers to compel suppliers to produce critical inputs.

Figure A.7: Global Demand and Supply



Notes: This figure presents the global marginal benefits and costs curves for vaccine capacity. We assumed a marginal cost curve with constant cost up to 200m doses a month at \$2 per course/year then increasing after that point with an elasticity of 1/3 (red), 2/3 (blue), and 4/3 (green).

Figure A.8: Price Ceiling



Notes: This figure presents the global marginal benefits and costs curves for vaccine capacity with a price ceiling at \$20 per course per year. We assumed a marginal cost curve with constant cost up to 200m doses a month at \$2 per course/year then increasing after that point with an elasticity of 1/3.

This analysis implies that it would be valuable to establish capacity for supply chains for vaccines or a stockpile of vaccine manufacturing inputs (bioreactors, glass vials, adjuvants, etc.) going into the next pandemic to ensure that supply is elastic. At an elasticity of $\frac{2}{3}$ (meaning that to induce a 1% increase in monthly capacity, the reimbursement must rise by 1.5%), the market clearing price would be around \$20 per vaccine course per year. Under this elasticity, Table A.7 shows that global net benefits are approximately a third larger than in the inelastic supply scenario. At an elasticity of $\frac{4}{3}$ (meaning that to induce a 1% increase in monthly capacity, the reimbursement must rise by 0.75%), the market clearing price would be around \$10 per vaccine course per year with global net benefits approximately double that in the inelastic supply scenario.

A.3 International Aspects

In this section, we analyze the incentives countries or coalitions of countries have to participate in centralized vaccine procurement programs. As shown in Section A.1, nationally optimal investment differed dramatically across countries. Richer countries would have invested much more than poorer countries. Thus, centralized systems with vaccine allocation proportional to population or health need would not have been individually rational: rich countries would have been better off designing their own procurement strategies, all else equal, than joining a coalition characterized by redistribution, or where the portfolio size and scale did not align with their interest. To be individually rational, countries would have to set their own investment levels.

To understand better the incentives of countries to join coalitions, we analyzed six potential coalitions of countries: the whole world, the US, the EU, all high-income countries (those with GDP per capita above \$17,000 or in the EU), BRIC (Brazil, Russia, India, and China, which are lower-income countries with significant vaccine production capacity), and the rest of the world (i.e., all countries outside BRIC that are not high-income).

For each coalition, we analyzed six different scenarios. The first four correspond to global procurement programs in which countries contribute in proportion to their GDP. In scenario (1), vaccination capacity is distributed by population. In scenario (2), capacity was distributed by high priority population, which shifted distribution towards high-income countries relative to scenario (1). In scenario (3), capacity was distributed according to contribution (i.e., GDP), and was likely to resemble the contributions that would result if countries decided contributions voluntarily. In these three scenarios the global program could procure vaccine production capacity at a price of \$10 per course/year. Scenario (4) is the same as scenario (3), except that we assumed that the global program was able to exploit its larger bargaining power to obtain a price of production capacity of \$8 per course/year. In all these scenarios, the global problem chose the portfolio that maximizes net benefits to the world.

In the final two scenarios, we assumed each coalition set up its own vaccine procurement program. In scenario (5), coalitions are only able to procure vaccines from 5 candidates. This stylized assumption is meant to illustrate what could happen if there is autarky, and coalitions were not able to procure vaccines from candidates produced in countries outside the coalition (for example, because countries use emergency powers to prevent export of vaccines). In scenario (6),

coalitions set up their own procurement contracts under free trade, which means they would be able to procure from all candidates. In both scenarios (5) and (6), each country chooses the portfolio that maximizes net benefits at a price of \$10 per course/year.

Table A.8 shows the net benefits of procurement programs to different coalitions in every scenario. Comparing columns (1) and (5), it is clear that high-income coalitions would have been much worse off with a global program that allocates vaccines by population, even if setting their own program implies autarky: having to share vaccines with lower-income countries means giving away a substantial fraction of benefits. Column (2) shows that that is also true if the global program gives priority to high-risk populations, which benefits high-income countries (more elderly). In column (3), high-income countries only get slightly higher benefits than in column (5). That means that in order to make the global program incentive compatible (relative to autarky), most of the vaccine production must be distributed according to contributions to the program.

Column (6) shows what happens if coalitions set up their own programs with free trade, which is the optimal scenario (at a price of \$10 per course/year) in terms of total welfare. Relative to this scenario, it is hard to justify a worldwide program unless it is able to get lower prices, as shown in column (4). However, even a substantial decrease in prices only results in minor gains for high-income countries relative to the net benefits from the program, since benefits are an order of magnitude higher than costs. Thus, even if joining a global program leads to lower prices, those countries might prefer their own program if that better allows them to obtain a more appropriate portfolio.

Table A.8: International Incentives

	Net benefits of the procurement program under different scenarios (\$ bn)						
	Worldwide, by population	Worldwide, with priority	Worldwide, by contribution	Worldwide, high bargaining power	Own program, autarky	Own program, free trade	
	(1)	(2)	(3)	(4)	(5)	(6)	
World	490	642	725	781	—	725	
US	114	173	243	256	217	247	
EU	123	180	210	219	186	211	
High Income	341	499	648	682	573	651	
BRIC	72	82	43	56	46	49	
Rest	77	61	35	43	38	41	

Notes: This table presents the net benefits of procurement programs to different coalitions under various scenarios outlined in Section A.2.

There are some potential benefits to international cooperation. First, centralized procurement can use monopsony power to hold down vaccine prices. However, we have not observed prohibitively high prices in the COVID-19 pandemic as of January 2021. Second, cooperation could help address the needs of lower-income countries for humanitarian reasons or to manage infectious disease externalities. There may also be benefits from insurance, economies of scale in planning, and supply chain investment.

A.4 Structuring Procurement Contracts

In this section, we further develop our analysis of the structure of procurement contracts for vaccines. Advance contracts can be structured in two broad ways.

- **Push funding:** Upfront payments directly reimbursing manufacturers' costs associated with installing capacity in parallel with clinical trials. Under this structure, the firm's costs are reimbursed whether or not its product is successful.

- **Pull funding:** Payments to firms for successful vaccine production, where the government commits to a price per course, typically a premium over production cost, often specifying quantity guarantees for suppliers and/or options for buyers. Under this structure, a firm only gets paid if the vaccine is approved and delivered.

Pull contracts can come in a variety of forms. They may take the form of a simple, bilateral deal between the buyer and a specified firm, committing the buyer to buy a specified quantity (perhaps to be delivered at a specified date) at a specified price if the product meets a set of conditions (e.g., regulatory approval). Unlike typical contracts for future delivery of existing products, contracts for vaccines still under development are typically not considered breached if the product is never produced or does not meet the conditions.

As an alternative to a bilateral deal, the pull contract can be a broader commitment to a minimum purchase quantity aggregated across all qualified candidates, with the buyer reserving the right to allocate purchases among candidates. This is part of the design of advance market commitments (AMCs), proposed by Kremer and Glennerster (2004). As originally proposed, AMCs had other features of importance for endemic diseases, for example, specifying a tail price, capping prices after the committed subsidy fund runs out, keeping long-run prices near production

cost, mitigating deadweight loss. In a pandemic crisis, short-run considerations dominate, so long-run features such as tail prices may be of less relevance.

In the setting of COVID-19 vaccines, we will label any advance contract committing to aggregate rather than bilateral purchases an AMC. While our usage is loose, usage by agencies and in the popular press can be even looser, labeling any advance contract, even bilateral deals, as AMCs. For example, the COVAX facility has provisions for subsidizing the participation by 92 middle- and lower-income countries; this part of the program is labeled the COVAX AMC, perhaps echoing the pilot AMC used to incentivize distribution of pneumococcal vaccine targeted to these same countries. The COVAX AMC does not share the other features of the pneumococcal program, designed as it was to pilot Kremer and Glennerster's (2004) proposal.

The fact that AMCs commit to an aggregate quantity rather than specified amounts from individual firms can create incentives for firms to compete on quality or timeliness to be among the selected products. Winners may be difficult to identify early on, so a broadcast approach, open to any qualified firm, may be preferable (Jeppesen and Lakhani 2010). Some drawbacks of AMCs may include that they are less familiar to lawyers and procurement bureaucracies, requiring additional work that may be difficult to accomplish quickly in an emergency. With capacity possibly scarce given high country demand in a pandemic setting, AMCs may be vulnerable to having successful suppliers "poached" by countries offering bilateral deals involving slightly higher prices than the AMC commitment. The problem may have been less relevant for the pilot pneumococcal AMC because this was targeted to low- and middle-income countries with fewer resources to strike bilateral deals. Furthermore, pneumococcus was not a pandemic; country demand increased gradually rather than far outstripping available supply (Kremer, Levin and Snyder 2020), providing less incentive for a country to jump the queue with a bilateral deal then. Perhaps for these reasons, contracts used in the pandemic were primarily simple, bilateral advance purchase agreements coupled with direct funding of development. Apart from the COVAX AMC, which as explained is only nominally an AMC, there were no explicit AMCs for COVID-19. Some implicit market incentives may have been provided by firms' expectation of a future market for vaccines, after the pandemic subsides and COVID-19 possibly becomes endemic over the longer term.

Looking more closely at the bilateral deals signed during 2020, many of those advance commitments specified prices that were relatively close to expected average cost; given the

uncertainty and risks that firms faced to accelerate manufacturing in an unprecedented way, profit margins were most likely modest, although full cost information is not public (and may in any case be difficult to compute, given that some costs are opportunity costs for firms of repurposing inputs and human resources, and risks include intangible factors such as reputational considerations). Thus, the bilateral deals probably did not substantially increase the incentives for firms to invest, but they did reduce uncertainty. In the past, in response to potential pandemics governments promised research funding and firms invested, but the funding later dried up as the virus dissipated, leaving firms holding the bag. The cycle is so familiar it's been labelled the panic-neglect cycle (Yong, 2017). A contract made firm investments more secure. The advance purchase agreements may have had other benefits such as ensuring an orderly market, especially in an environment where firms may have expected that in the absence of advance contracts, governments might be tempted to use emergency powers to direct production to themselves or their allies. The contracts may also have helped companies secure financing as well as cooperation of suppliers, and they may have encouraged countries to assist in resolving supply chain challenges, consistent with a variety of news reports. On the other hand, many of the bilateral deals signed also included funding for scaling manufacturing and late-stage clinical trials (at the government's risk), in addition to promises of support for supply chains, with Pfizer's deal with the United States a notable exception. Although details of the deals are not publicly available, overall we would characterize the deals that were signed as primarily push funding.

We now turn to a more detailed analysis of pull funding versus push funding. Pull funding can be more expensive than push funding in certain settings, including when firms' costs are perfectly observable but their success prospects are not. This can be most easily seen in an example. Consider two firms, one with a vaccine candidate having a 20% chance of success and the other with a 5% chance of success. Suppose that either firm can build a factory for \$4 per unit of capacity. Thus, for example, \$400 million would be sufficient push funding to induce a firm to build a factory having capacity for 100 million doses. To incentivize the first firm to build the same factory with pull funding, the government must guarantee a price of at least \$20 per dose ($\$20 \times 20\% = \4). To incentivize the second firm, the government must guarantee a higher price, \$80 per dose ($\$80 \times 5\% = \4), to compensate for the firm's lower probability of success. If the government could distinguish the firms and offer them different contracts, push and pull funding would be equivalently expensive. However, if the government must offer the same price

to all qualified candidates, in order to induce both firms to invest, it must offer \$80 per dose, leaving the first firm with rents and raising total government expenditures above that of equivalent push funding.

While the example provides a simple illustration of why push funding can be cheaper, it left out important incentive benefits from pull funding. Pull funding mitigates several forms of the moral-hazard problem (providing incentives to avoid cost bloat, since firms are residual claimants of their unreimbursed costs, and providing incentives for firms to stop when prospects become unrealistic) and the adverse-selection problem (firms with unrealistic prospects are induced not to participate). If structured like an AMC, as mentioned, pull funding can also incentivize quality and speed.

As a heuristic, in mid-2020, we suggested that governments use push funding for 85% of total costs, while using a market-wide AMC to help align private and social incentives for speed. The market-wide AMC was proposed to take the form of a bonus for the first courses to be purchased by the government following approval, and only courses delivered in a specific time frame would be eligible, so that firms would be incentivized to install capacity prior to approvals. The size of the bonus and the number of courses subject to the AMC are design features of the program. In order to induce firms to make the remaining 15% investment at risk, pull funding must be sufficiently large.

As an example, we now compute how large the investment should be in order to obtain the optimal portfolio for the world in column (3) of Table A.8. That portfolio has a total capacity of 2.36 billion courses per month, with investment in 18 different candidates. Assuming a capacity cost of $C = \$10$ per course per year, the cost of installing that capacity is \$280 billion. A program that pays for 85% of this sum upfront would have to pay \$238 billion in push funding. In addition to that, pull funding is provided in the form of a price p per course delivered within the first $m = 3$ months. In order to compute how high that price must be, consider a vaccine manufacturer j that produces at a marginal cost of $c = \$1$ per course, and whose vaccine is successful with probability q_j . For simplicity, we assumed that all firms have the same costs C and c , but have different success probabilities. The program buys a yearly capacity x_j from the firm, which ends up producing and selling $q_j y x_j$ doses in expectation, where $y = \frac{m}{12}$ years. Its total expected cost is equal to $0.15C x_j + c q_j y x_j$, and its expected revenue is equal to $p q_j y x_j$.

Suppose that, due to the risky nature of the investment, firms choose to participate if the expected revenue is at least 1.5 times the expected cost, or, in other words, if $p > 1.5 \frac{0.15Cx_j + cq_jyx_j}{q_jyx_j}$. All firms will then participate if the price is greater than $p^* = 1.5 \frac{0.15C + cq'y}{q'y}$, where q' is the lowest success probability among all the candidates that are part of the optimal portfolio. Using the success model described in section A.1.2, the lowest such probability is 0.132, resulting in a price per course of $p^* = \$68.93$. The expected pull funding is then $p^* \sum_j q_jyx_j$, which in this case is equal to \$112 bn. Note that this is an average: total pull funding might be higher if many successful candidates are realized. The total expected cost of the program is \$350 bn.

If, instead, the program only pays for 60% of the upfront costs, a similar accounting results in push funding of \$196, expected pull funding of \$222, and a total expected cost of \$418. Relative to the previous funding structure, the advantage of shifting funding towards pull is that it could result in greater incentives for success, both in development and in meeting the contracted time frame.

Pull funding is expensive due to the desire for a large portfolio of vaccines, and the fact that the marginal candidates perceive a low probability of success. Required pull funding could be reduced by increasing the share of funding given as push funding, perhaps targeted towards lower-probability candidates; and reducing the size of the portfolio for which at-risk investments are undertaken.

Another alternative to market-wide AMCs that our analysis suggests would be substantially cheaper is bilateral pull funding that directly incentivizes speed, but is targeted to individual firms making use of information about their success probabilities. That of course requires such information to be known by the buyer. Later in the development timeline, this information is more likely to be available to governments. We did not provide a full quantitative analysis of bilateral incentives for speed, and further exploring the tradeoffs and details of such contracts represents an area for future research.

As discussed above, actual contracts signed did not incorporate the type of market-wide pull funding we recommended, and the firm-specific pull funding provided only modest incentives given that firms should have rationally expected that they would be likely to be able to sell a safe and effective vaccine at similar or higher prices without the agreements. It appears that despite

the lack of substantial incentives, a number of vaccine candidates advanced to regulatory approval. However, it seems plausible that stronger incentives for speed might have incentivized firms to go farther in expanding capacity and investing in alleviating supply chain bottlenecks sooner. The delays experienced in manufacturing as of early 2021 highlight that substantial investment and redundancy are required to meet the ambitious goals of vaccinating at a rate that comes close to the socially beneficial one.

A.5 Table of Vaccine Candidates

Vaccine candidates in optimal selection order to maximize probability of at least one success (August 2020)

Platform	Subcategory	Phase	Cumulative Probability	Marginal Probability
Inactivated	Inactivated	Phase 3	0.288	0.288
Viral vector	Adenovirus (non-replicating)	Phase 3	0.483	0.195
RNA	LNP-encapsulated mRNA	Phase 3	0.583	0.099
Inactivated	Inactivated	Phase 3	0.658	0.074
Protein subunit	Recombinant protein	Phase 2	0.707	0.049
Protein subunit	S protein	Phase 2	0.744	0.036
Protein subunit	Recombinant protein	Phase 2	0.769	0.025
RNA	LNP-encapsulated mRNA	Phase 3	0.790	0.020
Inactivated	Inactivated	Phase 3	0.807	0.016
Viral vector	Adenovirus (non-replicating)	Phase 2	0.821	0.013
VLP	VLP	Phase 1	0.832	0.011
Viral vector	Adenovirus (non-replicating)	Phase 2	0.840	0.008
Viral vector	Measles (replicating)	Phase 1	0.847	0.006
Protein subunit	S protein	Phase 1	0.852	0.005
DNA	Electroporation	Phase 2	0.857	0.004
Protein subunit	S protein	Phase 1	0.861	0.003
Live attenuated virus	Live attenuated	Pre-clinical	0.865	0.003
DNA	Other DNA	Phase 2	0.868	0.002
Live attenuated virus	Live attenuated	Pre-clinical	0.870	0.002
Protein subunit	Recombinant protein	Phase 1	0.873	0.002
Live attenuated virus	Live attenuated	Pre-clinical	0.8752	0.002
Live attenuated virus	Live attenuated	Pre-clinical	0.877	0.001
Protein subunit	S protein	Phase 1	0.878	0.001
Live attenuated virus	Live attenuated	Pre-clinical	0.880	0.001
DNA	Plasmid + adjuvant	Phase 2	0.881	0.001
RNA	mRNA	Phase 1	0.882	0.001
Live attenuated virus	Live attenuated	Pre-clinical	0.883	0.001
Protein subunit	Recombinant protein	Phase 1	0.884	0.001
Viral vector	Horsepox (replicating)	Pre-clinical	0.885	0.001
Viral vector	Influenza (replicating)	Pre-clinical	0.886	0.001

Appendix References

- AstraZeneca. 2020. "AZD1222 vaccine met primary efficacy endpoint in preventing COVID-19". *Press release*.
- Bubar, Kate M., Kyle Reinholt, Stephen M. Kissler, Marc Lipsitch, Sarah Cobey, Yonatan H. Grad, and Daniel B. Larremore. 2020. "Model-Informed COVID-19 Vaccine Prioritization Strategies by Age and Serostatus," *MedRxiv*, January 202009.08.20190629, accessed at <https://doi.org/10.1101/2020.09.08.20190629>.
- Britton, Tom, Frank Ball, and Pieter Trapman. 2020. "A mathematical model reveals the influence of population heterogeneity on herd immunity to SARS-CoV-2." *Science* 369, no. 6505: 846-849.
- Cutler, David M., and Lawrence H. Summers. 2020. "The COVID-19 pandemic and the \$16 trillion virus," *Journal of the American Medical Association* 324, no. 15 (2020): 1495-1496.
- Davies, N. G., P. Klepac, Y. Liu, K. Prem, M. Jit, R. M. Eggo. 2020. "Age-Dependent Effects in the Transmission and Control of COVID-19 Epidemics," *Nature Medicine* 26: 1-7.
- Demombynes, Gabriel. 2020. "COVID-19 Age-Mortality Curves are Flatter in Developing Countries," *World Bank Policy Research Working Paper*.
- DiMasi, J. A., L. Feldman, A. Seckler, and A. Wilson. 2010. "Trends in Risks Associated With New Drug Development: Success Rates for Investigational Drugs," *Clinical Pharmacology & Therapeutics* 87 (3): 272-77, accessed at <https://doi.org/10.1038/clpt.2009.295>.
- Duke Global Health Innovation Center. 2020. "Launch and Scale Speedometer," Accessed at <https://launchandscalefaster.org/covid-19>.
- Dunn, Andrew. 2020. "Some Experts Are Deeply Skeptical We'll Have a Coronavirus Vaccine within 18 Months. Here's Why a Shot May Take Years, If We Can Succeed in Making One at All," *Business Insider*, April 25, 2020. accessed at <https://www.businessinsider.com/coronavirus-vaccine-timeline-prospects-can-we-make-one-2020-4>.
- Goldhill, Olivia. 2020. "COVID-19 Vaccine-Distribution Timeline Will Keep Slipping, Experts Say," *STAT+*, December 11, 2020, accessed at <https://www.statnews.com/2020/12/11/covid19-vaccine-timeline-keeps-slipping-experts-say-it-will-change-again/>.
- Hogan, A. B., P. Winskill, O. J. Watson, P. G. T. Walker, C. Whittaker, M. Baguelin, D. Haw, A. Løchen, K. A. M. Gaythorpe. 2020. Imperial College COVID-19 Response Team, F.

- Muhib, P. Smith, K. Hauck, N. M. Ferguson, A. C. Ghani. “Report 33: Modelling the Allocation and Impact of a COVID-19 Vaccine,” *Imperial College London*.
- Hodgson, John. 2020. “The Pandemic Pipeline,” *Nature Biotechnology* 38 (5): 523–32, accessed at <https://doi.org/10.1038/d41587-020-00005-z>.
- Jeppesen, Lars Bo, and Karim R. Lakhani. 2010. “Marginality and Problem-Solving Effectiveness in Broadcast Search,” *Organization Science*, 1–18, accessed at <https://doi.org/10.1287/orsc.1090.0491>.
- Johnson, Ed, and Angus Whitley. 2020. “Australia Cancels Order for CSL Vaccine as Trial Stumbles,” *Bloomberg.Com*, December 10, 2020, accessed at <https://www.bloomberg.com/news/articles/2020-12-10/australia-cancels-order-for-csl-vaccine-as-trial-stumbles>.
- Kalorama Information. 2020. “Vaccine Development and Production Trends, COVID-19 and Other Vaccines Pipeline,” accessed at <https://kaloramainformation.com/product/vaccine-development-and-production-trends-covid-19-and-other-vaccines-pipeline/>.
- Khamsi, Roxanne. 2020. “If a Coronavirus Vaccine Arrives, Can the World Make Enough?” *Nature* 580 (7805): 578–80, accessed at <https://doi.org/10.1038/d41586-020-01063-8>.
- Kis, Zoltan, Cleo Kontoravdi, Robin Shattock and Nilay Shah. 2021. “Resources, Production Scales and Time Required for Producing RNA Vaccines for the Global Pandemic Demand.”
- Kremer, Michael, Jonathan D Levin, and Christopher M Snyder. 2020. “Advance Market Commitments: Insights from Theory and Experience,” Working Paper Series, No. 26775, National Bureau of Economic Research, accessed at <https://doi.org/10.3386/w26775>.
- Le, Tung Thanh, Zacharias Andreadakis, Arun Kumar, Raúl Gómez Román, Stig Tollefsen, Melanie Saville, and Stephen Mayhew. 2020. “The COVID-19 Vaccine Development Landscape,” *Nature Reviews Drug Discovery* 19 (5): 305–6, accessed at <https://doi.org/10.1038/d41573-020-00073-5>.
- Lurie, Nicole, Melanie Saville, Richard Hatchett, and Jane Halton. 2020. “Developing COVID-19 Vaccines at Pandemic Speed,” *New England Journal of Medicine* 382 (21): 1969–73, accessed at <https://doi.org/10.1056/NEJMp2005630>.
- Mullard. 2020. “How COVID vaccines are being divvied up around the world,” *Nature*, November 30, 2020, accessed at <https://www.nature.com/articles/d41586-020-03370-6>.
- Peel, Michael, Richard Milne, and Donato Mancini. 2021. “EU countries decry ‘very short notice’ of delay in delivery of Pfizer vaccine,” *Financial Times*, January 15, 2021, accessed at <https://www.ft.com/content/e8177df6-04ae-4d20-8e62-ca76589c7653>

- Pfizer. 2020. "Pfizer and BioNTech to submit emergency use authorization request today to the US FDA for COVID-19 vaccine," *Press release*. Friday, November 20, 2020, accessed at <https://www.pfizer.com/news/press-release/press-release-detail/pfizer-and-biontech-submit-emergency-use-authorization>
- Plotkin, Stanley, James M. Robinson, Gerard Cunningham, Robyn Iqbal, and Shannon Larsen. 2017. "The Complexity and Cost of Vaccine Manufacturing – An Overview," *Vaccine* 35 (33): 4064–71, accessed at <https://doi.org/10.1016/j.vaccine.2017.06.003>.
- Pronker, Esther S., Tamar C. Weenen, Harry Commandeur, Eric H. J. H. M. Claassen, and Albertus D. M. E. Osterhaus. 2013. "Risk in Vaccine Research and Development Quantified," *PLoS ONE* 8 (3), accessed at <https://doi.org/10.1371/journal.pone.0057755>.
- Reuters. 2020. "China's Sinovac to double annual COVID-19 vaccine capacity to 1 billion doses," January 13 2021, accessed at <https://www.reuters.com/article/us-health-coronavirus-vaccine-sinovac/chinas-sinovac-to-double-annual-covid-19-vaccine-capacity-to-1-billion-doses-idUSKBN29I0YN>
- Robbins, Rebecca, and Benjamin Mueller. 2020. "After Admitting Mistake, AstraZeneca Faces Difficult Questions About Its Vaccine," *The New York Times*, November 25, 2020, sec. Business, accessed at <https://www.nytimes.com/2020/11/25/business/coronavirus-vaccine-astrazeneca-oxford.html>.
- Thompson, Stuart A. 2020. "Opinion | How Long Will a Vaccine Really Take?" *The New York Times*, April 30, 2020, sec. Opinion, accessed at <https://www.nytimes.com/interactive/2020/04/30/opinion/coronavirus-covid-vaccine.html>.
- Walsh, Fergus. 2020. "Oxford-AstraZeneca Vaccine: Bogus Reports, Accidental Finds - the Story of the Jab," *BBC News*, December 14, 2020, sec. Health, accessed at <https://www.bbc.com/news/health-55308216>.
- World Bank. 2020. "Global Economic Prospects, January 2020: Slow Growth, Policy Challenges," Washington, DC: World Bank. DOI: 10.1596/978-1-4648-1468-6.
- World Health Organization. 2020. "Landscape of COVID-19 Candidate Vaccines," accessed at <https://www.who.int/publications/m/item/draft-landscape-of-covid-19-candidate-vaccines>.
- Yong, Ed. 2017. "The Deadly Panic-Neglect Cycle in Pandemic Funding," *The Atlantic*. October 24, 2017, accessed at <https://www.theatlantic.com/science/archive/2017/10/panic-neglect-pandemic-funding/543696/>.
- Zimmer, Carl, Sharon LaFraniere, and Noah Weiland. 2021. "Johnson & Johnson Expects Vaccine Results Soon but Lags in Production," *The New York Times*, January 13, 2021, sec. Health, accessed at <https://www.nytimes.com/2021/01/13/health/covid-vaccine-johnson-johnson.html>.