ADB Distinguished Speaker Lecture

The Economics of Investing in COVID-19 Vaccines: Implications for the Asian Development Bank

Michael Kremer*

I. Introduction

I am pleased to have the opportunity to discuss the coronavirus disease (COVID-19) pandemic and the experience of the Accelerated Health Technology Group to draw some lessons in these pages. The group was created by colleagues on the Advance Market Commitment for pneumococcal vaccine, to which donors committed a $1.5 billion reward to help finance the purchase of an approved vaccine against pneumococcus strains prevalent in low- and middle-income countries.

As we began to engage with the case of COVID-19 vaccines, we realized there were some fundamental differences between the economics of COVID-19 vaccines and those of the pneumococcal vaccine. We therefore involved a wider group of economics, including experts in contract design, such as Susan Athey, to help think through the economic impacts.

---

* Michael Kremer (corresponding author): Kenneth C. Griffin Department of Economics, The University of Chicago, United States. This article is based on a lecture delivered at the Asian Development Bank Distinguished Speaker Program on 1 October 2020. It does not reflect more recent developments.

Asian Development Review, Vol. 40, No. 1, pp. 1–12
DOI: 10.1142/S0116110523500099

We are all aware of the significant human and economic costs of the pandemic. But it is still worth quantifying the impact of accelerating vaccine availability. The International Monetary Fund initially expected that the COVID-19 pandemic would reduce global gross domestic product (GDP) by $12 trillion over the 2020–2021 period. This implied that ending the economic losses of the pandemic a month earlier than forecasted would be worth $500 billion. This is based on back-of-the-envelope calculations, but it is likely a conservative estimate as it does not include averted mortality and health benefits. Therefore, in terms of COVID-19 vaccine development, speed has always been essential.

This urgency contrasts with the normal timeline for developing vaccines. Vaccines are different from drugs. Manufacturing is complex and requires the installation of unique factories and equipment. Each piece must be separately licensed and tested. There is a saying in the industry that the process is the product. Governments do not license vaccines but rather a particular production factory, and that means that it typically takes 3–4 years from the initial testing of a vaccine to widespread commercial availability.

Normally, because this is a complicated, expensive, and risky process, firms only install capacity at commercial scale after trials are complete. These trials typically take at least 6 months, which implies a long delay before vaccines are available. Moreover, even when firms do build capacity, they typically build somewhat limited capacity—that is, just enough to serve high-income markets. There is usually a long delay of a decade or more before low- and middle-income countries are served.

This makes commercial sense for a private company in normal times, but the COVID-19 pandemic was a case in which the social value of early, large-scale capacity investment greatly exceeded the private value. By early, I mean installing the capacity in parallel with the testing process rather than waiting until testing is completed.

There are several reasons for this. First, there are large epidemiological externalities from vaccination. If I take a vaccine, that does not just benefit me. It benefits many other people, and therefore the social value exceeds the private value. Second, governments do a lot of the purchasing of vaccines, which creates political constraints and puts limits on pricing, perhaps especially in a pandemic. This is not to say such limits are not warranted, but with limits on pricing, the social gains to expanding capacity will rapidly exceed the private gains. Third, if firms install large-scale capacity, that can put downward pressure on prices in high-income countries. All these factors tend to mean that the social value of early, large-scale investment greatly exceeds the private value to the manufacturer.

There are some effects that go in the opposite direction. There could be, hypothetically, a business-stealing effect, but that is likely much smaller. To get a sense of the relative magnitude, note that early in the epidemic when Moderna announced positive vaccine news, its stock price went up considerably. The value of the increase in terms of market capitalization was $5 billion. However, the total gains among the 500 companies included in the S&P Index reached $810 billion. This only captures the increases in the stock market in the United States (US), and while these types of calculations are always subject to debate, it suggests that a private company is only going to capture a tiny fraction of the overall social value of a vaccine.

If we think that the social value of developing a vaccine early in the course of a pandemic is huge but the private value might not be as large, then it may make sense for governments to make a deal in advance to purchase the vaccine even before the testing is complete in order to induce manufacturers to install large-scale production capacity before they have trial results. *Ex ante*, it may or may not be worth it as it is a quantitative question. Our analysis below suggests that, in the case of the COVID-19 pandemic, it was worth it.

The US paid AstraZeneca $1.2 billion for 300 million doses, so roughly $4 per dose. The US was an early mover in this respect, but COVAX had two programs: one that would be funded to cover low-income countries and another in which middle-income countries would have to purchase vaccines themselves. The World Bank also initially committed $12 billion for developing countries to finance the purchase and distribution of COVID-19 vaccines, tests, and treatments for their citizens in October 2020. Were these multiple policy moves worth it? The potential benefit of financing capacity installation in parallel with testing, rather than waiting for the testing to be done, is that the vaccine might become available earlier. What is the potential cost? Governments might waste money on a vaccine that fails. Most candidate vaccines fail. It is very difficult to predict, based on early trials, whether a vaccine will succeed, so there is a considerable risk. Let me present some back-of-the-envelope calculations to assess whether these bets were worth it.

Let us take that US deal as an example. Suppose *ex-ante* there is only a 10% chance that government-guaranteed purchases would accelerate a vaccine's availability by 6 months. (Though, in retrospect the success rate was higher, which only strengthens the case for investment.) Working with that 10% number, the benefit–cost ratio is 45:1, so it is worth it with even a very low chance of success given the large GDP losses from the pandemic continuing without mitigation from vaccination. Another way to put this is that a $1.2 billion investment would be worth it if we knew for sure it would accelerate the vaccine by just 10 hours.

Here is another calculation. Imagine that all vaccines had a 10% chance of success. For the sake of simplicity, suppose that the vaccine trials are independent draws. How many independent developers would you want to engage? It turns out you would want

37 “shots on goal.” The actual calculation is considerably more complicated, but I think the intuition that you want to invest in a lot of different vaccines is correct.

I want to emphasize that our group’s focus was on the question of vaccine capacity. As researchers, we focus on a limited number of issues. In contrast, a multilateral development bank must confront a whole set of issues including research and development trials and vaccine delivery plans. As we saw in the case of COVID-19, vaccine delivery and uptake is an important challenge that requires careful thought, but the focus of this piece is on vaccine capacity.

II. Vaccine Procurement

In the introduction, I argued that it is worth investing at risk. But a key lesson from economics generally and our experience with COVID-19 and pneumococcal vaccines is that the details of contract design matter immensely. One lesson from our experience is that both economics and legal expertise in contract design are critical. While I am a development economist, some of my colleagues came from an industrial organization and contract design perspective, and I will draw on their insights below.

The first principle, in a pandemic context, is that doses are much more valuable if delivered early. That may sound trivial. But we often sign contracts and need something delivered soon but do not get it delivered soon. I moved house recently, so I have been dealing with housing contractors. You can sign a contract with the contractor saying that the house will be finished on 1 November. But what do you do if the contractor later says the house will not be ready in time? You are unlikely to simply walk away from the deal and say: “You didn’t deliver on time, forget it. We don’t want the house after all.” You need that house. Similarly, we are unlikely to walk away from vaccine developers when they face delays.

How can contracts be designed to avoid this problem? One way is to include a penalty clause to say that if the house is not ready on time, there is an extra payment. Another approach is to say that if you are more than 1 month late, you pay an extra $10,000 for every day you are late. This is a standard contract design issue, which is a trade-off between incentives and risk-bearing.

Imagine we wanted to get the incentives exactly right for vaccine development. Then we would need to charge the pharmaceutical companies $500 billion for every month they are late. That imposes a huge risk on these companies. No company would accept that risk, or if they did, they would have to charge huge amounts for something that could bankrupt them, especially because there are legitimate events outside their control that they cannot do anything about. Some clinical trials have had to be paused due to adverse events. Therefore, an approach based on a penalty clause will not achieve a first-best result—the risks are too high for the firm to bear.

But one thing governments can do to mitigate the risk of delays is to make the contract specify that the company has to install sufficient manufacturing capacity. Otherwise, the firm may install a small factory and governments risk getting the doses 2 years late. One lesson from the Advanced Market Commitment for the pneumococcal vaccine work and COVID-19 vaccine modeling is that to accelerate vaccine capacity, one should directly incentivize it. One way to structure that is through a contract in which the company is paid to install capacity, and in exchange, the company gives the right to purchase doses that come from that capacity so the purchaser can get those doses when they are ready.

I do not want to claim that is the only way to structure a contract. There are probably multiple ways, but what is key is to create incentives to actually build the capacity. This is a basic lesson, but it takes a lot of work to figure out how to implement.

Another issue is whether to pay a reward only for a successful product or to pay up front before knowing whether the product will be successful? The pneumococcal Advanced Market Commitment was structured as a $1.5 billion reward for a successful product—so the company only got paid if their vaccine was successful—and there are circumstances under which that type of approach makes sense.

In the COVID-19 vaccine case, one thing that comes out of the full analysis is that it makes sense to have many shots on goal. It is optimal to incentivize many different vaccine developers to invest in a vaccine, even if some of them have a low subjective probability of success. In the simplified example in the introduction, everybody knew that all the candidate vaccines had a 10% chance of success. Obviously, the real world is not like that. Some candidates are more likely to succeed, others are less likely to succeed.

However, this is an asymmetric information problem. The optimal procurement design is going to depend on the information structure and the ability to differentiate prices between different producers. The buyer does not have the full information and there is uncertainty about many variables. One thing buyers probably have particularly bad information about is the firm's estimate of the probability of success. There is also uncertainty about the cost of capacity installation, but that is probably much more observable. If the firm says it needs a particular piece of machinery for its factory, they can show the invoice for the cost of that machinery. Maybe they can inflate the cost a little bit, but there are limits to that.

I shall offer an illustrative example. It is an artificial case, but in our model, we try to use actual numbers to gauge if these effects are important. One type of funding, which is what was used for the Advanced Market Commitment for the pneumococcal vaccine, was "pull" funding. Pull funding is only given for a successful vaccine. The alternative is "push" funding, which is up-front financing that provides money to cover costs (e.g., for a factory). The buyer knows that each unit of capacity costs $4 to install, but he does not know the chance of success for the producer and that the producer might believe that they have a 20% chance that their vaccine will succeed, or perhaps they are pessimistic and they think there is only a 10% or 5% chance of success.

How much push funding and how much pull funding is needed to incentivize the construction of the factory? If it is push funding, governments just pay for the capacity. That would be true in any of these three scenarios (i.e., 5%, 10%, or 20% chance of success). The chance of success does not affect the cost of building the factory. So, it would cost $4 per dose to build this.

There is another way to pay — using pull funding. Governments can say that they are only going to pay if the firm is successful. To illustrate, let us start in the top row of Table 1.

Imagine that buyers knew there was a 20% chance of success. If the firm is risk-neutral, which it is presumably not, and estimates a 20% chance of success, then they would need to be paid five times the $4 if they succeed to compensate them for the additional risk, so the buyer would need to pay them $20. The buyer pays the same amount in expectation of a successful vaccine.

But now, suppose that buyers do not know the chance of success, or buyers do not know what the company believes the chance of success to be. It might only be 10%. For 10%, the company would need $40 to justify this investment; or it might only be 5%, in which case they would need $80 for a successful vaccine to justify the investment. That puts the buyer in a very difficult situation, because this is a case where a 5% chance of success is big enough that it is worth investing in the vaccine. The buyer can say that they are just going to offer only $20 for success. That is fine if it turns out that the company believes there is a 20% chance of success, but if it turns out the company does not believe that and is not willing to invest, then the buyer may risk losing out on a successful vaccine. Alternatively, buyers might wind up paying

Table 1. **Push and Pull Funding Example**

Probability of Success	Necessary Push	Necessary Pull
20%	$4	$4/0.20 ~ $20
10%	$4	$4/0.10 ~ $40
5%	$4	$4/0.05 ~ $80

Source: Author's calculation.

the full $80, but then buyers are paying much more than you would have needed to in these other possible scenarios in which the company was more optimistic.

Governments can solve this problem by just paying up front for the capacity (i.e., push funding). However, there are downsides to this approach as well. Governments could say that they will conduct a scientific assessment of which vaccines are worth investing in and then pay for the factory capacity accordingly. I think that should be a fairly large element of any incentive scheme.

In practice, the model suggests that these types of forces are important and that it makes sense to have a big push element. It probably does make sense to have some pull element as well to incentivize speed and capacity. If there are some manufacturers who know that they really have a one-in-a-thousand chance of success, governments may not want to invest in them or pay for their factory. Having them put some skin in the game probably makes sense. The extreme conclusion that you go with 100% push funding is probably not correct.

III. Investing at Risk

This ought to give a sense of what the structure of a contract might look like. The optimal structure of a contract includes paying for capacity with substantial up-front funding. If that is the form of the contract, is it worth investing at risk and how much do you want to invest? That is what we turn to now.

Regarding the details that went into the model, we looked at the economic harm caused by COVID-19 based on estimates from the International Monetary Fund, which was typically a contraction in GDP of 5%–10%. That was also true for most member economies of the Asian Development Bank, but there was obviously variation. We tried to include the health benefits of the vaccine, but these results were driven by the economic benefits. We cut all benefits in half since another treatment or mitigation strategy could have been developed before the vaccine became widely available. That was a crude approach, but we were trying to be conservative. A large share of the health benefits came from vaccinating health-care workers and the elderly, so we assumed that delivering the first set of doses is the most valuable. We assumed that investing early in parallel with testing accelerates vaccine availability by 3 months.

On the supply side, we looked at more than 100 vaccine candidates. Based on the historical data on chance of success at different stages and expert opinion, we assigned probabilities of success to each candidate. We also made assumptions on the correlation in probability between different vaccine candidates and, to some extent, the differences in cost due to different technological approaches.

One thing the analysis brought out is that there were particularly high returns to the initial candidates and the initial units of capacity per candidate. This is because it is optimal to start with the candidates that have a high ratio of probability of success to cost and then to move on to less desirable candidates. The probability of success is capped at 1. That approach is obviously slightly artificial. In the real world, vaccines have different characteristics. There are benefits to having multiple vaccines, so this is a rough approximation.

Figure 1 illustrates the very high returns—immense amounts of money that are saved and immense amounts of lives—from a successful vaccine. There are very high returns for the first candidates, and then it flattens out a bit.

Figure 2 shows the benefits in terms of months saved from adding capacity. Initially, building capacity moves the development of a successful vaccination forward in time. But eventually, when adding capacity, the amount of time saved starts to flatten out. This is because when doubling capacity, the cost doubles, but the time to vaccinate only halves. With twice as much capacity, it takes half as long to vaccinate. Doubling again doubles the cost, but the time halves from an already lowered level, so the absolute gain is smaller.

These very high returns for the first vaccine candidates and initial capacity imply that all countries—not just high-income ones—could have invested in vaccine development. Yes, the US and the United Kingdom were among the first to invest in

Figure 1. **Probability of Success as a Function of Vaccine Candidates**

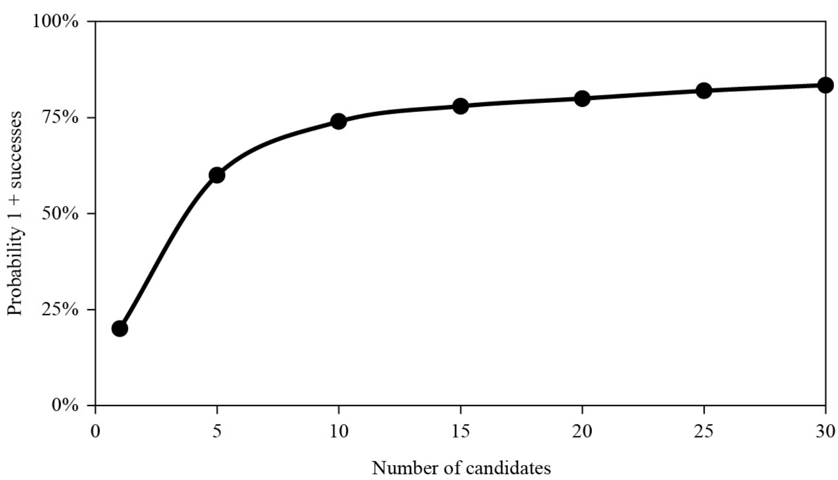

Note: The chart illustrates the probability of at least one success as a function of the number of vaccine candidates.
Source: Author's illustration.

Figure 2. **Months to Vaccination as a Function of Capacity Added**

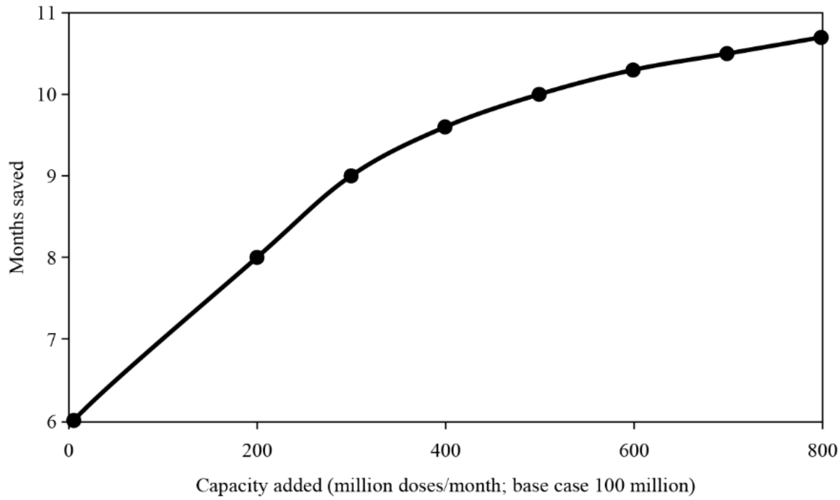

Note: The chart illustrates the number of months to vaccinate 1.2 billion individuals accelerated by added capacity.
Source: Author's illustration.

Table 2. **Comparison of Low- and Middle-Income Country Investment Decisions**

	One Vaccine Candidate			**Three Vaccine Candidates**		
Country Group	**Expected Economic Benefits ($ billion)**	**Cost ($ billion)**	**Benefit–Cost Ratio**	**Expected Economic Benefits ($ billion)**	**Cost ($ billion)**	**Marginal Benefit–Cost Ratio**
Low income	5	1	5.00	10	5	1.25
Middle income	87	4	21.75	174	18	6.21

Source: Author's calculations.

a large quantity of vaccines, but it actually made sense for even low-income countries to invest in some vaccines. Suppose that the following assumptions held, we compare low-income countries and middle-income countries in Table 2:

1. First candidate has a 25% chance of success. Three candidates have a 50% chance of at least one success.
2. First candidate sells each course of vaccines for $4; other candidates sell for $6.
3. Enough vaccines will be purchased to vaccinate 20% of the population.

Investing in just one vaccine candidate costs $1 billion for low-income countries but has $5 billion in benefits, so low-income countries should do it. For middle-income countries, the benefit–cost ratio is even greater at almost 22:1. Expanding out to three candidates, it is still worth it for the low-income country in this run of our model. They get an additional $5 billion of benefits, while it costs them an additional $4 billion. Therefore, the marginal benefit–cost ratio is just 1.25. It would have been optimal for low-income countries to order a few candidates, but not a large number.

On the other hand, for middle-income countries it was a very high return on investment to seek additional candidates. If we extended this and showed estimates for high-income countries, it would have made sense for them to order numerous vaccine candidates.

What are the implications for optimal investment? All countries should invest at risk. Obviously, for humanitarian and equity reasons, it makes sense for all countries to have access to a vaccine, but this calculation shows these were also good investments based only on the economic benefits and assuming that countries would spend their own money. Of course, the higher-income countries should invest in more candidates and more capacity. This also has one important application, which I will come back to, for international cooperation. As we have seen with the US and the United Kingdom, the optimal policy for a high-income country was to take a lot of shots on goal and to put in a lot of capacity for each of those shots.

Let us turn to middle-income countries. Latin America was hit hard by the pandemic, and some middle-income countries in Asia did not see as many deaths but bore large economic costs. For middle-income countries, it makes sense to invest a fair amount in vaccines, but not quite as much as high-income countries. The optimal investment program differs for individual countries. But if everybody were completely internationally minded and focused on equity and global health, countries would have spent a combined $100 billion on an international effort to develop a vaccine. Unfortunately, we did not see that. Instead, we saw vaccine nationalism—that is, countries seeking out vaccine suppliers and making purchases for themselves.

IV. Policy Implications for the Next Pandemic

The implications for domestic policy makers and regional development banks are clear. In contexts such as the COVID-19 pandemic, you should invest early in capacity at risk.

From a global perspective, it makes sense to support low- and middle-income countries to vaccinate the priority populations early—health workers, other essential workers, vulnerable populations—which will benefit the country and the world as well.

What are the implications for the Asian Development Bank? Early investment at risk makes sense. If countries face financing challenges, it makes sense to provide financing, particularly for the vaccination of priority populations. It is also important to work with partners to structure contracts to expand their capacity.

Several issues merit further discussion. One is exchange mechanisms, as different vaccines have different characteristics. We have mostly abstracted from that, but some vaccines are easier to deliver than others. Some might be one dose, while some might be multiple doses. Some vaccines might be available earlier than others. Some vaccines might be 50% effective, while others might be 80% effective. Countries may have different preferences for vaccines depending on their needs with respect to delivery infrastructure and containment situations. If countries put different weights on vaccine characteristics, then having some ability to exchange vaccines probably makes sense. A mechanism like COVAX could work better if it included a mechanism for exchanges to be made.

I will briefly mention supply chains. Suppose that to increase the supply of vaccines we need more bioreactors, or more glass vials, or more adjuvants. If we are trying to immunize the whole world quickly, we need a huge supply of those things. We may need to expand the factory capacity for those inputs. For producers of inputs facing a temporary increase in demand that might not last a long time, it is unclear whether it makes sense to build a huge factory that is going to last for 20 years given the risk it will end up idle. That is a tough decision to make.

To make the decision to build up capacity, producers may need a very high price. If there are political constraints on pricing, firms may not want to build that factory at all. One of the findings from our analysis is that just as there is a case for building support for investing in vaccine capacity, there is a case for investing in capacity for the whole supply chain. Since much of the vaccine supply chain is in Asia, the Asian Development Bank could potentially contribute by making not just sovereign loans to purchase vaccines, but also loans to firms to increase production capacity—both for vaccines and for the supply chain capacity.

V. Conclusion

Optimal incentives for vaccine development in a time-sensitive pandemic context include a large push component with up-front payments for the installation of production capacity in exchange for the option to purchase the vaccine. That is probably the most efficient way to do it, but it does not have to be done that way.

The main thing is just to make sure we invest and do it in a way that is actually going to accelerate the capacity installation.

The optimal number of vaccine candidates and production capacity differ among countries, so building flexibility into whatever system is used makes sense. I think that the COVAX self-pay program does some of that, and probably could go a bit further. Allowing exchange mechanisms could also be valuable, and it is worth investing in supply chains. The Asian Development Bank could have a major role to play—both by lending to countries to buy vaccines (especially for priority populations) and to build supply chain capacity.

Getting contracts right is tricky. If governments buy doses but do not specify when those doses are going to come, then they may just get the doses at the back of the queue 2 years later and not actually advance the delivery of vaccines. That is a situation to avoid in the future: Some low- and middle-income countries were effectively at the back of the queue during this pandemic. There are also ways to design contracts so that governments are getting as much value for money as possible by avoiding paying excessive rents to inframarginal firms. That probably involves paying for capacity, rather than just for a successful vaccine.

References

Ahuja, Amrita, Susan Athey, Arthur Baker, Eric Budish, Juan Camilo Castillo, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Jean Nahrae Lee, Candice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon Joel Tan, and Witold Więcek. 2021. "Preparing for a Pandemic: Accelerating Vaccine Availability." *AEA Papers and Proceedings* 111: 331–35.

Castillo, Juan Camilo, Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Tasneem Chipty, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Greg Larson, Jean Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon Joel Tan, and Witold Więcek. 2021. "Market Design to Accelerate COVID-19 Vaccine Supply." *Science* 371 (6534): 1107–09.

Kremer, Michael, Jonathan Levin, and Christopher M. Snyder. 2020. "Advance Market Commitments: Insights from Theory and Experience." *AEA Papers and Proceedings* 110: 269–73.

Kremer, Michael, and Christopher Snyder. 2020. "Strengthening Incentives for Vaccine Development." *NBER Reporter* 4: 7–11.