Vaccine Supply Chain Resilience and International Cooperation

Chad P. Bown
Peterson Institute for International Economics
Center for Economic and Policy Research (CEPR)

Christopher M. Snyder
Dartmouth College Department of Economics
National Bureau of Economic Research (NBER)

Robert W. Staiger
Dartmouth College Department of Economics
National Bureau of Economic Research (NBER)

July 2022
Preliminary and Incomplete

Abstract: We construct a theoretical model in which a country's government procures a vaccine from a firm to protect its citizens against pandemic harm. The firm can accelerate the vaccine's availability by investing before regulatory approval, but this risks stranding investment if approval fails. The government can encourage such at-risk investment by adding capacity subsidies on top of its procurement price. We analyze the optimal procurement mechanism under asymmetric information about the firm's cost and study how this mechanism changes when part of the vaccine supply chain is located offshore. In the absence of international cooperation on vaccine supply-chain policies, offshoring leads the government to reduce the subsidy to avoid information rents leaking to foreign firms whose profits it does not internalize. A second distortion arising with offshored vaccine inputs is that the foreign government may restrict exports in a crisis, holding up the capacity subsidy. Calibrations using estimates of the costs and benefits of a Covid-19 vaccine suggests that the identified distortions can be substantial and that international cooperation aimed at reducing them can result in a many-fold improvement in program net benefits.

Journal of Economic Literature codes: F53 (International Agreements and Observance), I18 (Government Health Policy), L65 (Biotechnology), D82 (Mechanism Design)

Keywords: Vaccines, Covid-19, supply chain, at-risk investment, government procurement, subsidy, international cooperation

Acknowledgements: The authors are grateful to Arnaud Costinot, Teresa Fort, Matthew Grant, Sam Kortum, Marc Melitz, Nathan Zorzi, and seminar participants in the Dartmouth Trade Workshop and Princeton IES Summer Trade Workshop for helpful feedback.

1. Introduction

Accelerating the pace of vaccine development and manufacturing was crucial during the Covid-19 pandemic. In record time, multiple Covid-19 vaccines were invented, proceeded successfully through clinical trials, and began to be produced at commercial scale. Yet, an open question is whether elements of that sequential process, especially investment in production capacity, could have been started earlier so that more vaccine doses would have arrived sooner. Potentially millions of lives, as well as trillions of dollars of economic activity, could have been saved (Cutler and Summers, 2020; Agarwal and Gopinath, 2021). Why governments did not intervene more to accelerate and expand vaccine production capacity remains a puzzle.

Heading into the pandemic, much of the pharmaceutical industry was characterized by a fragmented production process, relying on outsourcing to contract manufacturers, often also with offshoring to firms located abroad. A plant to manufacture the drug substance for a vaccine, for example, might be located in one country, while a second country would host a second and different type of facility to then formulate that vaccine and place it into vials for distribution. Furthermore, specialized inputs for the industry may only have been available through imports. A vaccine manufacturer in India or Germany might require variable inputs only supplied, at least in the very short run, by a firm at arm's length located in the United States or United Kingdom.

This paper provides a theoretical framework to investigate a number of potential implications of these issues. It provides one justification for why most of the early vaccine supply chains to emerge during the Covid-19 pandemic were national and concentrated in only a few countries, despite the possibility for much greater geographic diversification of that fragmented production process. The framework also allows for an examination into whether an industrial structure characterized by offshoring and imported inputs, where it did arise, may have contributed to the failure of governments to align private and social incentives for early investment in vaccine manufacturing. Finally, we use the framework to investigate the possibility of potentially novel motivations for international economic policy cooperation—in the form of multilateral commitments to maintain expansionary subsidy policies along the vaccine supply chain and to avoid export restrictions—to ensure the resilience of that supply chain during a pandemic.

We develop a simple theoretical model of vaccine procurement in a pandemic reflecting key

insights from of a series of articles by the Accelerating Health Technologies team (Snyder et al., 2020; Ahuja et al., 2021; Castillo et al., 2021) summarized in Athey et al. (2022). We characterize a pandemic as imposing per-period harm on individuals that can be mitigated with a vaccine. The model's firm is endowed with a promising vaccine technology and faces the decision over when to invest in costly production capacity, capacity that takes time to install. We focus on the possible divergence between private and social incentives to undertake "at-risk" investment in capacity for vaccine production, as emphasized by Athey et al. (2022) and Bown (2022b). Following these authors, by at-risk investment we mean investment in vaccine production capacity during the period in which clinical trials for the vaccine are in progress, and hence before it is known whether the vaccine is safe and effective and will be granted approval for use to inoculate the population against the virus. When the installation of capacity takes time, the benefit of at-risk investment—as opposed to the alternative of investment in capacity only once the vaccine has been approved for use—is the accelerated availability of the vaccine, if it turns out to be successful, to the population; the cost of at-risk investment is that this investment in capacity may turn out to be wasted if the vaccine is not approved for use and the capacity cannot be repurposed.

We start by analyzing a closed-economy model in which the government procures the vaccine from an integrated firm located domestically. We allow the government a full gamut of contractual instruments beyond simple linear price per dose including capacity subsidies and bonuses for early delivery. In the absence of other frictions, those instruments allow the government to obtain the first best. The fact that the firm only internalizes its profits, not the full social benefit of capacity or its acceleration via at-risk investment, which might lead to underinvestment with linear vaccine payments, can be corrected by direct government capacity subsidies.

The compounding of two standard frictions adapted from the literature on procurement and regulation (Laffont and Tirole, 1993) leads the optimal procurement contract for the government to fall short of the first best. One friction is a social cost of public funds, reflecting the deadweight loss of taxation needed to fund public expenditure. Estimates of the deadweight loss of taxation by Snow and Warren (1995) suggest that the social cost of public funds is substantial even for developed countries. This friction leads the government to curtail the range of parameters for which it induces any capacity investment and the smaller range of parameters for which it induces at-risk investment in what we label the second-best outcome (the government's optimum with only the friction of the

social cost of public funds).

The second friction introduced in the model is that the firm's capacity cost is private information which the firm observes but the government only knows the distribution of. Asymmetric cost information is not only the canonical friction in theoretical analyses of procurement (Laffont and Tirole, 1993), it is also of practical relevance in vaccine procurement contracts, thought to be a key friction in major recent vaccine procurement programs including the advance market commitment piloted by GAVI (formerly the Global Alliance for Vaccines and Immunizations) for pneumococcal vaccine conducted over the past decade (Snyder et al., 2011; Kremer, Levin, and Snyder, 2020, 2022) as well as procurement of Covid-19 vaccines in the pandemic (Snyder et al., 2020). In what we label the third-best outcome (the optimal contract for the government with both a social cost of public funds and asymmetric information), the government further curtails the thresholds for both at-risk and any investment in capacity to reduce information rents flowing to the firm. In the absence of a social cost of public funds, those rents would be transfers having no effect on the surplus of a benevolent government internalizing domestic firm profit. In the presence of a social cost of public funds, however, information rents are distortionary, leading the government to economize on them by reducing capacity subsidies and incentivizing less investment.

With this setup in hand, we turn to the core questions of the role of offshoring and imported inputs. To do so, we adapt our model to a setting in which the domestic vaccine manufacturer needs to source inputs via imports from a foreign supplier. We assume that the contracting firms engage in efficient Nash bargaining over all investment and supply decisions, allowing them to achieve the integrated outcome (thus abstracting from hold up between contracting firms introduced in Antras and Staiger (2021)). The foreign input supplier obtains its bargaining share of integrated profits. Assuming the domestic government internalizes only profits earned by domestic not foreign firms, the offshoring of vaccine inputs leads the goverment to curtail the thresholds for at-risk and any investment in capacity yet further relative to the third-best levels with domestic input supply. Leakage of information rents to foreign firms is even more costly to it than leakage to domestic firms.

We also model a second distortion from foreign input supply. We posit a probability that the foreign country restricts export of essential vaccine inputs, effectively preventing the firm from fulfilling its contract to supply the domestic government, holding up the domestic government's

capacity subsidy and the firm's investment. The prospect of export restrictions leads the domestic government to yet further curtail the thresholds for at-risk and any capacity investment. Of course the direct loss of vaccine supply to the domestic government is a direct loss of global surplus caused by export restrictions (to the extent not offset by increased supply to citizens of the foreign country).

International cooperation can serve to eliminate these distortions and return the outcome to the third best with domestic input supply. To the extent that foreign input supply would have been preferred because that it is cheaper there—due either to comparative advantage or perhaps historic location advantages—international cooperation can even improve on this third-best outcome. In that case, the potential benefits of international cooperation could be hidden. The world would observe the domestic government offering capacity subsidies to incentivize investment by domestic input suppliers. The outcome seems efficient because indeed the subsidies are second-best efficient conditional on the supplier's location. It would be harder to ascertain that the off-equilibrium distortions to capacity subsidies when the input is supplied by a foreign firm might be deterring the firm from locating input supplies abroad, where costs would have been lower.

To assess whether any of the identified distortions are large enough to be of practical concern and, if so, to gauge their relative importance, we calibrate the performance of a hypothetical US procurement program for Covid-19 vaccines. We draw on survey data of Covid-19 vaccine manufacturers from Snyder et al. (2020) to provide an estimate of the distribution of vaccine costs. That source also provides an estimate of the per-capita harm from the Covid-19 pandemic potentially relieved by a vaccine. We find that the social cost of public funds and asymmetric information can combine to reduce vaccine program net benefits by over 30% and cut the probability of at-risk investment nearly in half. Assuming a 50% chance that export restrictions are imposed by the foreign country, in the absence of international cooperation over vaccine supply-chain policies, net program benefits are 60% lower when vaccine inputs are offshored compared to onshored, about a third of this loss coming from the subsidy-leakage effect and two thirds from the possibility of export restrictions. At-risk investment virtually disappears in the calibration with offshored input supply. Compared to the third best with foreign input supply but no international cooperation, international cooperation over vaccine capacity subsidies and export restrictions leads to an 1.5-fold gain.

As far as we are aware, our paper is the first in the economics literature to provide a formal analysis of government supply chain policies directed to at-risk investment, and to consider the

role of international cooperation in the design of such polices. Grossman, Helpman, and Lhuillier (2021) is related, but their focus is on supply chain resilience more generally rather than attaining the optimal level of at-risk investment in pandemic situations, and they do not consider the possibility of international policy cooperation; and given their different focus, their modeling approach is of course quite different as well. In the operations research literature, Sun, Toyasaki, and Sigala (2021) is closer to our paper in terms of focus, but their analytical approach is quite different and they do not consider international policy cooperation either.

The rest of the paper proceeds as follows. Section 2 provides the institutional background, describing the key empirical patterns of Covid-19 vaccine supply chains and the policies that emerged during the pandemic. Section 3 introduces our basic model of vaccine manufacturing in an economy that is experiencing a pandemic, and derives the essential timing of the at-risk and not-at-risk capacity investment decisions that will form the focus of our ensuing analysis. Section 4 considers optimal government policies for the case of domestic input supply, first in a frictionless (first-best) world and then when there is a cost of public funds (second best) and when the vaccine producing firm possesses private information (third best). Section 5 considers how optimal government policy is altered if input supply comes from a foreign country, both in the case where governments choose their vaccine supply chain policies noncooperatively and in the case where they are able to cooperate internationally over these policies. Section 6 turns to a rough calibration of the model, in order to assess the quantitative importance of the distortions to vaccine supply chain policy associated with the frictions that the model considers, and to gauge the magnitude of the potential gains from international cooperation over vaccine supply chain policy. Section 7 concludes. The appendix contains proofs, omitted from the text for brevity, of some lemmas and propositions.

2. Institutional Background

Covid-19 vaccines were completely new products. This section draws from Bown and Bollyky (2022), which tracks both the emergence of the supply chains (from scratch) needed to manufacture the new vaccines and the government policies announced over 2020–21 that potentially impacted their formation as well as subsequent production. That research catalogs the vaccine supply chains to arise from Pfizer-BioNTech, Moderna, AstraZeneca-Oxford, Johnson & Johnson (Janssen), No-

vavax and CureVac. Through 2021, the only other Covid-19 vaccines with sizeable production were from China (Sinovac and Sinopharm), Russia (Sputnik V, from Gamaleya Research Institute of Epidemiology and Microbiology), and an additional one in India (Bharat Biotech).[1] In addition to these vaccine candidates, governments provided support for a number of others over 2020–21, including some that were not authorized for distribution to the general public.

2.1. Vaccine Supply Chains

This section briefly summarizes five key features of the Covid-19 vaccine supply chains that emerged over 2020–21.

First, there was geographic diversity at the national level regarding where these vaccines were invented. The Moderna and Novavax vaccines, for example, originated in the United States (US); the Johnson & Johnson vaccine was co-invented between scientists at its Janssen lab in the Netherlands and scientists at a hospital in Boston. The AstraZeneca-Oxford vaccine was invented in the United Kingdom (UK), and vaccines from BioNTech and CureVac originated in Germany.

Second, the vast majority of vaccine production would ultimately be conducted via subsequent outsourcing arrangements in which the vaccine sponsor contracted at arms length with third parties to handle manufacturing. In part, this was because the inventors of the vaccine technology were often either a biotech firm (BioNTech, Moderna, Novavax, CureVac) or a university (Oxford), and thus without access to their own production facilities, at least at the onset of the pandemic. However, even Johnson & Johnson hired a number of so-called contract development manufacturing organizations (CDMOs) for its production needs, as did AstraZeneca (despite having its own production facilities globally) when it was hired to coordinate the manufacturing and global distribution of the Oxford vaccine. The primary exception was Pfizer, which retrofitted its own plants in the US to manufacture the BioNTech vaccine. (The European supply chain for the Pfizer-BioNTech vaccine involved a network of plants from Pfizer, plants that BioNTech purchased that eventually came online, as well as other CDMOs.)

Third, the manufacturing process for each of these Covid-19 vaccines exhibited considerable

[1] India's main Covid-19 vaccine output was the AstraZeneca-Oxford vaccine produced locally by the Serum Institute of India (SII). Even though Novavax and CureVac announced full supply chains, Novavax had minimal global production through 2021 as it was slow in being authorized by regulators, and CureVac's initial Covid-19 vaccine candidate reported poor phase 3 trial results in June 2021 and was subsequently abandoned.

fragmentation, with its core being two distinct plants. The vaccine drug substance was manufactured in one facility, and almost always a separate plant would receive that drug substance, add more ingredients to formulate it into the vaccine's drug product, and then "fill and finish" it into tens of thousands of sterile glass vials, assembly-line style, for distribution.[2] For Covid-19, each of these plants had to be retrofitted with specialized capital equipment, and their manufacturing processes would then also become subjected to strict regulatory oversight.

Fourth, those drug substance and fill-and-finish plants for any given vaccine's supply chain were typically located in the same country (or in the case of the European Union (EU), the same union of countries).[3] This arose even with considerable opportunity for further geographic diversification, especially given that the fragmentation of the manufacturing often resulted in two different CDMOs handling the two different production steps, even for the same vaccine. Furthermore, to sell to different markets, most vaccine sponsors chose to set up parallel supply chains, with drug substance and fill-and-finish plants rarely in separate countries. With the exception of CureVac, each of the major vaccines set up parallel supply chains in at least the US and EU. (AstraZeneca and Novavax set up additional parallel supply chains in other regions of the world.) One implication of such a choice is that vaccine sponsors did not set up an alternative "hub and spoke" supply chain structure with, say, one upstream plant manufacturing all of the drug substance for its vaccine at greater scale, the output of which could then be sent to multiple downstream plants in different countries to be formulated, filled, and finished.

Fifth, virtually all firms complained of input shortages. In addition to insufficient availability of (the capital embodied in) entire fill-and-finish facilities, for example, firms also complained about access to variable inputs from their other, arms-length suppliers. Examples ranged from the lipid nanoparticles that were essential for the new mRNA vaccine technology platforms used by Pfizer-BioNTech and Moderna, to the bioreactor bags, filters and other "consumable" inputs used up in their part of the production process at the plants making vaccines for other firms.

[2]The main exception was the complex mRNA vaccine of Pfizer-BioNTech in which there were multiple plants involved in an even more fragmented process of manufacturing the drug substance.

[3]One notable exception was Moderna's European supply chain where the drug substance facility was in Switzerland, whereas the fill-and-finish plants were in Spain and France.

2.2. Government Policies Impacting Vaccine Manufacturing

Governments pursued a variety of contracting approaches with the vaccine companies beginning in 2020. However, relatively few wrote at-risk (push or pull) contracts to accelerate vaccine manufacturing, and fewer still subsidized any firms beyond the initial vaccine sponsor—i.e., there were few subsidies allocated directly to input providers elsewhere along the supply chain.

The US was the major exception. Through its Operation Warp Speed initiative begun shortly after the onset of the pandemic in early 2020, the federal government initially supported seven vaccine sponsors and provided the largest amount of total funding. It started by funding a number of clinical trials, including the lengthy, costly, and pivotal phase-3 trials. Then, in the summer and fall of 2020, it contracted with five different vaccine sponsors to accelerate vaccine production in advance of (still ongoing) phase-3 trials and committed to purchase 100 million (or more) doses upon emergency use authorization (EUA) from the US Food and Drug Administration (FDA). While the public versions of the contracts with these companies contain considerable redactions, according to the US Government Accountability Office (GAO), the agency tasked under the CARES Act with oversight of Congressional funding for Covid-19, the contracts included at least some push (guaranteed funding for inputs, regardless if the vaccine received regulatory approval) and pull (bonus payments for early delivery of doses) incentives to accelerate the production capacity capable of producing at least 100 million doses.

On push funding, for example, the government agreed to pay Moderna "incrementally for meeting certain milestones without requiring Moderna to first obtain an EUA"; furthermore, the government funding assisted the company "with cash flow by providing interim payments." Each of the five contracts also included the right for the government to terminate the agreement and only have to pay "for work performed in accordance with the agreement terms" (GAO, 2021, p. 19). Given that the government had the right to terminate the contracts, the total amount of funding each firm would receive was unknown. Finally, on pull funding, the initial contract with Moderna had a firm-fixed price of $12.50 per dose for the first 100 million doses, but it also included a bonus payment of $3.00 per dose for meeting a regulatory authorization deadline of January 31, 2021. (The FDA authorized Moderna's vaccine for emergency use in the US on December 18, 2020.)

The US government had a completely different, procurement-only contract with Pfizer. The US negotiated a $1.95 billion contract in July 2020 for 100 million doses of its vaccine, thus providing

it a relatively higher price ($19.50 per dose). However, the payment would only be made upon regulatory approval and delivery of doses. Thus, in this case, Pfizer retained the risk of failure.[4]

Finally, the US also provided contracts directly to a number of input providers beginning in the summer of 2020 so that they could install additional capacity at risk. This included at least one contract with a fill-and-finish facility to be used as part of a Covid-19 vaccine manufacturing process, as well as capacity-enhancing contracts to other firms providing variable inputs essential to vaccine production—e.g., bioreactors, bioreactor bags, cellular material—and vaccine dose delivery—e.g., vials, needles, syringes (Bown and Bollyky, 2022, Table 4).

No other government matched the scale or scope of the US at-risk contracts to vaccine sponsors or to input providers. In Europe, the UK was closest, committing a mostly nonrefundable £914 million in five contracts with vaccine sponsors "prior to any vaccine being approved by the regulator. . . to start manufacturing and to support clinical trials" (NAO, 2020, pp. 24–25). The UK government payments were also to be credited against future purchases of any vaccines authorized by regulators. Germany and the EU (via the European Investment Bank) provided at-risk financing to BioNTech and CureVac to allow them to expand production capacity. The Coalition for Epidemic Preparedness Innovations (CEPI), a foundation dedicated to financing independent research projects to develop vaccines against emerging infectious diseases, provided some at-risk funding to the Serum Institute of India to manufacture the AstraZeneca vaccine to be allocated to COVAX. COVAX, an acronym for Covid-19 Vaccines Global Access, is a facility aiming to procure and distribute Covid-19 vaccines to low-income countries globally (Bown and Bollyky, 2022, Tables 6 and 7).

Even in major vaccine manufacturing economies, some governments took a different approach. For example, aside from minor at-risk subsidies to help BioNTech and CureVac access additional manufacturing capacity, the EU focused on negotiating procurement agreements collectively for the EU member states, at low prices, and also on ensuring that EU member states did not fight over supplies once they became available.[5] India, the country with the largest vaccine manufacturing company (Serum Institute of India) heading into the pandemic, did not offer subsidies to get its vaccine companies to expand capacity until April 2021, eight to ten months later than many of the contracts were agreed in the US, for example.

[4] According to the GAO, "To minimize financial risk to the government, the parties agreed that the government would pay Pfizer only after its vaccine received authorization" (GAO, 2021, p. 18).

[5] This may have been motivated in part by EU member states imposing export controls on personal protective equipment trade with each other in March 2020 (Bown, 2022a).

Not surprisingly, at a broad level the vaccines that received early regulatory approval (after good phase-3 trial results) with manufacturing facilities supported by at-risk funding tended to be the "success" stories. Upon regulatory approval, they got more vaccine output from their plants than the manufacturing plants in other countries, or the manufacturing plants in the same countries where contracts were not provided at risk (Bown, 2022b,d). Finally, and as expected, some of the at-risk investments were lost because some vaccines did not make it out of phase-3 trials and thus were not authorized for use by regulators.[6]

A final important policy issue to arise during the pandemic was the possibility that countries might impose export restrictions on inputs critical to vaccine manufacturers located abroad. CEOs of downstream vaccine manufacturers in India and Europe, for example, accused the US government of limiting exports of such critical supplies; the political issue escalated to the level of French President Emmanuel Macron accusing the Biden administration of banning such exports in May 2021 (Bown and Rogers, 2021; Bollyky and Bown, 2021). The US government denied the allegations, indicating that the lack of sufficient exports was due to an input shortage that had forced it to ration supplies under the Defense Production Act. On the other hand, the threat of the UK imposing export restrictions on lipid nanoparticles, a critical input destined for the Pfizer-BioNTech plants in the EU, reportedly helped stop a trade war from erupting between the UK and EU, including the possibility that the EU might limit shipments of finished AstraZeneca vaccines back to the UK (Bown and Bollyky, 2022, pp. 477–479).

3. Model

A country with a continuum of citizens with a mass normalized to 1 faces a viral outbreak. Absent a vaccine, the virus circulates in the population for t_e periods, at which point the outbreak ends. The end of the outbreak could reflect the emergence of a highly contagious but benign variant of the virus after period t_e which dominates other variants, precluding serious harm from the disease after that point. Alternatively, period t_e could reflect the date after which a repurposed generic drug is discovered to eliminate harm from the virus at very low cost. We take the outbreak to be of moderate enough duration that it is reasonable to abstract from discounting, setting the discount rate

[6]For US regulators and the vaccine candidates that received at-risk funding under Operation Warp Speed, this included AstraZeneca, Novavax and Sanofi-GSK.

for all players to 0.

We adopt a very simple model of the disease's epidemiology. Each period during the outbreak, an unvaccinated individual receives an independent and identically distributed (iid) draw of a Bernoulli random variable indicating their infection status, with $\beta \in (0,1]$ denoting the probability they are infected. An infected person is sick during the period and recovers by the period's end. Let $h > 0$ denote the social harm from sickness. This factor reflects the individual's physical harm from illness as well as possible wage and education losses. It may also reflect broader social costs including the provision of medical services and economic output losses beyond the wage. To avoid the complication of changing population size over time, we assume (somewhat counterfactually for Covid-19) that infection does not result in death. Our model abstracts from the dynamics of contagion, assuming an individual's chance of infection is independent of time and disease prevalence. Recovery does not confer immunity: a recovered person can be reinfected the next period, possibly capturing the circulation of new variants that escape antibodies generated by previous infections.

A government G seeks to maximize social welfare in the country by optimizing the procurement of a vaccine. The government deals with a single, domestic, profit-maximizing firm, which develops a promising vaccine candidate in period 0. We initially analyze the case of integrated domestic supply, shown in the bottom left panel of Figure 1. In this case, the firm produces the vaccine with domestically sourced inputs produced completely within the firm boundaries. Later, we will analyze cases, shown in the other panels of the figure, in which the vaccine producer (D) obtains inputs from a separate firm (U), located in the same or a different country depending on the case analyzed.

Clinical trials for the vaccine take time, spanning periods $t = 1$ through period $t_a - 1$. We abstract from clinical-trial costs, reflecting the reality that they are eclipsed by capacity and production costs for vaccines supplied at pandemic scales (see cost estimates in Snyder et al. (2020)). After receiving the data on safety and efficacy upon completion of clinical trials, at the beginning in period t_a, a regulator determines whether to approve the vaccine for use. The vaccine cannot be used unless and until it is approved. The regulator is non-strategic, basing its decision on an objective analysis of the clinical-trial data. Let $s \in (0,1]$ be the probability that the vaccine candidate succeeds in clinical trials and is approved.

The vaccine offers only non-durable protection: a dose only protects a person vaccinated at the beginning of the period against infection for the duration of that period. One interpretation, reflect-

ing the evolution of recent Covid-19 variants, is that continual protection requires an individual to receive a booster each period. A vaccinated person receives a draw of a Bernoulli random variable, iid across people and periods, indicating the efficacy of the dose, with $\theta \in (0,1]$ denoting the probability that the dose is effective.[7]

We model vaccine production as a two-step process. An input (say the vaccine's active ingredient) produced in an upstream facility U is shipped to downstream facility D that fills a vial with the input and other additives constituting a finished dose. As mentioned, we begin by assuming the two facilities are located domestically and integrated in the same firm. In any period t, input production cannot exceed U's capacity Q_t. Capacity costs k per unit to install. After it is installed, one unit of capacity can produce the input requirement for one vaccine dose each period thereafter. Capacity is sunk and cannot be repurposed to manufacture other pharmaceuticals or products.[8] Besides the capacity cost, each unit of input involves a marginal production cost of c. In sum, input production involves linear capacity and production costs.[9] We abstract from capacity and production costs involved in transforming the input into a final vaccine dose, assuming that is done costlessly and is not capacity constrained.

The firm can choose to install any amount of capacity in any period in as many tranches as it chooses. Without further assumptions, it is evident that it is both profitable and socially efficient to hold off investing until after regulatory approval to avoid sinking capacity investment for a candidate that ends up failing. However, we assume that installing capacity takes time: in particular, capacity installation begun in a certain period is not completed and cannot begin producing until t_ℓ periods later. The lag of t_ℓ periods provides a rationale for investing before period t_a, while clinical trials are still underway, referred to as "at-risk" investment. At-risk investment trades off of wasted capacity investment if the vaccine candidate fails against earlier availability of capacity allowing production of doses with less of a lag to if the candidate succeeds. Once the installation process is underway, we abstract from any option value from abandoning it; the k capacity cost per unit of the capacity

[7] While the model allows θ to be a free parameter, in practice the approval process may lead to restrictions on its value, for example only granting approval if θ exceeds a certain threshold.

[8] An equivalent alternative assumption is that capacity is fungible but only imperfectly so, with k capturing the value lost in repurposing.

[9] The assumption that capacity and production costs are linear was also adopted by the series of articles (Ahuja et al., 2021; Castillo et al., 2021; Athey et al., 2022) motivating our work. The basic results are not sensitive to the functional form of costs. A previous version of the paper, available from the authors on request, assumed a general convex function for capacity costs, obtaining similar basic results. The tractability of linear costs allows us to push the analysis further in this paper, among other results deriving the optimal mechanism under private cost information.

installation begun then is expended in any event. It turns out to be convenient to adopt the accounting convention associating capacity investment with the period in which it is started. Thus, we let x_t denote new capacity investment started in period t, which will not be productive until at least period $t+t_\ell$ (later if approval does not come until later).

To avoid a taxonomy of uninteresting cases, assume that t_ℓ, t_a, and t_e are natural numbers ordered as

$$t_\ell \leq t_a < t_a + t_\ell < t_e. \tag{1}$$

The first inequality implies that capacity can be installed at least as quickly as clinical trials can be completed, opening up the possibility of starting at-risk capacity installation early enough that it is ready for production immediately upon vaccine approval in period t_a. The last inequality implies that the pandemic lasts long enough after the approval date that there will be positive demand in some periods for capacity installed on or after that date, which we will refer to as capacity installed "not at risk."

If at-risk capacity provides too much of an advantage over not-at-risk capacity, the analysis becomes uninteresting because the firm's private incentives will favor at-risk capacity without any need for external inducement. The following condition turns out to be necessary and sufficient to rule out that uninteresting case:

$$1 - s > \frac{t_\ell}{t_e - t_a}. \tag{2}$$

In words, the condition says that the probability of failure must exceed the proportion of the post-approval period that can only be served by at-risk capacity.

Letting X_r denote total at-risk capacity installed and X_n total not-at-risk capacity installed, we have $X_r = \sum_{t=1}^{t_a-1} x_t$ and $X_n = \sum_{t=t_a}^{t_e} x_t$. Since there is no discounting and capacity cost is linear, it is immediate that, rather than spreading each of X_r and X_n out in tranches, a capacity strategy that is both weakly profitable and weakly socially efficient is to undertake each investment all together as early as possible. This strategy starts the installation of X_r in period 1 and the installation of X_n in period t_a. By definition, installation of X_n cannot start before t_a or else it would be classified as at-risk capacity.

The firm can pursue one of two strategies for utilizing at-risk capacity X_r. It can delay utilization until period t_a when it learns whether the vaccine has been approved, or it can utilize X_r while

clinical trials are being run to build a stockpile ready to augment supplies rolled out after period t_a. The advantage of the delay strategy is that it avoids wasting production costs for the stockpile if the vaccine candidate fails to be approved. On the other hand, stockpiling allows a given quantity of vaccine to be produced using less capacity. We acknowledge that stockpiling is a theoretically interesting strategy and had the potential to improve global surplus if pursued more vigorously in the Covid-19 pandemic. To simplify the analysis and reflect the actual amount of stockpiling during the pandemic, however, we abstract from stockpiling, assuming either that the vaccine has a limited shelf life, precluding the storage of a stockpile, or that the lag between the start of capacity investment and its completion is roughly the same as that between a vaccine's development and regulatory approval, i.e., $t_\ell \approx t_a$. Either assumption would preclude utilization of at-risk capacity before period t_a to build a stockpile.

The simplicity of the capacity-investment strategy allows for a simple depiction of the model's timing, shown in Figure 2. The vaccine is developed in period 0. Clinical trials last from period 1 to period $t_a - 1$. The regulator makes the approval decision in period t_a. At risk capacity X_r that was installed in period 1 is the only capacity available from period t_a to $t_a + t_\ell - 1$. In period $t_a + t_\ell$, the lag between when installation of not-at-risk capacity X_n starts, in period t_a, and when it becomes available for production is over. From period $t_a + t_\ell$ to the last pandemic period $t_e - 1$, combined capacity $X_r + X_n$ is available.

The amount of available capacity is the only difference among the periods after regulatory approval. Given that infected individuals recover by the end of the period and immunizations only provide one period of protection, the consumer population looks identical at the start over every period between t_a and t_e. Let Q_1 denote the output produced with capacity X_r each period from time t_a to $t_a + t_\ell - 1$ (a time interval indicated by the first shaded box in Figure 2). Let Q_2 denote the output produced each period with capacity $X_r + X_n$ from period $t_a + t_\ell$ to $t_e - 1$ (a time interval indicated by the second shaded box in the figure). Since supply and demand conditions are identical across periods within each shaded interval, specifying a constant output per period in each shaded interval is without loss of generality.

To avoid the trivial (and counterfactual) outcome in which the government is able to obtain the first best across a variety of different configurations with quite simple contracts, we introduce frictions in subsequent subsections. The frictions are standard in the industrial-organization literature

on procurement (see e.g., Laffont and Tirole (1993)). Section 4.2 introduces the social cost of public funds, denoted $\lambda \geq 0$, reflecting the deadweight loss of distortionary taxation needed to raise revenue for public expenditures.

Section 4.3 introduces an asymmetric-information friction. In particular, following textbook treatments of government procurement (Laffont and Tirole, 1993), we will suppose that the private information regards firms' costs. As discussed in the introduction, the existence of asymmetric cost information is not only the canonical friction in theory, it was also of practical relevance in the procurement of Covid-19 vaccines as well as earlier programs. While firms might have private information regarding capacity cost, production cost, or both, for simplicity, we will assume that capacity cost k alone is the source of private information. Assume k has probability density function (pdf) $f(k)$, cumulative distribution function (cdf) $F(k)$, and support given by non-negative real numbers. Assume $F(k)$ is logconcave, a property exhibited by most commonly used distributions in practice (Bagnoli and Bergstrom, 2005). In first-best and second-best scenarios, we will assume that k is observable to both the firm and the government. In third-best scenarios, we will assume that the firm observes the realization of k but the government only knows the distribution.

4. Domestic Input Supply

In this section, we analyze the case in which the input supplier is a domestic firm, adhering to the model introduced in the previous section, corresponding to the structure depicted in the lower left panel of Figure 1. We begin by characterizing the first best in which the government has full control over firms' operations and faces no other frictions, whether a social cost of public funds or private information about capacity costs. We then study the second best which is the same as the first best except the government now faces a social cost of public funds added to any expenditures. We then study the third best, the optimal government contract when it faces both a social cost of public funds and asymmetric information.

All the analysis assumes establishments U and D are integrated in the same firm. The section concludes with a note that the analysis is identical if the domestic establishments are unintegrated, operating as separate firms. The distinction between integration and not will only start to matter in a later section when one or the other establishments is located abroad.

4.1. First Best

To characterize the first best, we will solve for the optimum chosen by a government that internalizes consumer and producer surplus, that faces no social cost of public funds, and that controls all the operations of the integrated firm. The government faces no asymmetric information, able to observe the random draw of capacity cost k and condition policy on that realization.

The government controls all firm operations, including the firm's capacity decisions, which, as argued in Section 3, can be reduced to two choices: the amount of at-risk capacity X_r and not-at-risk capacity X_n. The government also controls the firm's output path over the post-regulatory-approval time interval. In general, this output path could be quite complicated; but as we saw in Section 3, simplifying features of the model reduced it to two quantities, Q_1, produced from period t_a to $t_a+t_\ell-1$ using at-risk capacity X_r, and Q_2, produced from period t_a+t_ℓ to t_e using combined capacity X_r+X_n.

Expected consumer harm over the pandemic is given by

$$\sum_{t=1}^{t_a-1}\beta h+\sum_{t=t_a}^{t_a+t_\ell-1}\beta h(1-s\theta Q_1)+\sum_{t=t_a+t_\ell}^{t_e-1}\beta h(1-s\theta Q_2) \tag{3}$$

$$=\beta h t_e-s\beta h\theta Q_1 t_\ell-s\beta h\theta Q_2(t_e-t_a-t_\ell). \tag{4}$$

The first term in equation (4) reflects the disease harms experienced over the pandemic in the absence of a vacine, equal to the mass of infected consumers each period β, times the harm h, times the pandemic duration t_e. The next term reflects the expected reduction in harm from the Q_1 doses supplied each period when only at-risk capacity is available times the duration of that interval. The benefit of those doses are scaled by the probability of success s and efficacy θ. The last term reflects the expected reduction in harm from the Q_2 doses supplied each period after not-at-risk capacity comes online to supplement at-risk capacity. We will focus not on expected consumer harm but the reduction in that harm due to the vaccine, embodied in the last two terms in (4):

$$s\beta h\theta Q_1 t_\ell+s\beta h\theta Q_2(t_e-t_a-t_\ell). \tag{5}$$

Expected total production cost over the pandemic is given by

$$\sum_{t=t_a}^{t_a+t_\ell-1} scQ_1 + \sum_{t=t_a+t_\ell}^{t_e-1} scQ_2 = scQ_1 t_\ell + scQ_2(t_e - t_a - t_\ell). \tag{6}$$

Production costs are only expended if the vaccine is approved, so all terms in (6) are scaled by the probability of success s. Expected total capacity cost is given by

$$kX_r + skX_n. \tag{7}$$

The investment cost associated with at-risk capacity is expended with certainty, but that associated with not-at-risk capacity is only expended conditional on successful approval, with probability s.

The first best maximizes expected vaccine benefits (5) net of the costs in (6) and (7):

$$\max_{X_r, X_n, Q_1, Q_2 \geq 0} \left[sQ_1(\theta\beta h - c)t_\ell + sQ_2(\theta\beta h - c)(t_e - t_a - t_\ell) - kX_r - skX_n \right], \tag{8}$$

subject to

$$Q_1, Q_2 \leq 1 \tag{9}$$

$$Q_1 \leq X_r \tag{10}$$

$$Q_2 \leq X_r + X_n. \tag{11}$$

Constraint (9) ensures that there is no benefit from vaccinating more than the total population mass, normalized to 1. Constraints (10) and (11) ensure output cannot exceed available capacity in the relevant periods.

The maximization problem in equations (8)–(11) is a linear program. Linear programs typically involve corner solutions, and that is the case here. Both optimal capacity and optimal output are corner solutions. If the optimum involves positive output, it can be implemented soley with at-risk capacity, or if not, soley with not-at-risk investment. Sufficient capacity is installed to vaccinate the entire population each period. Optimal output fully utilizes available capacity. Formally, we have the following lemma, where a star added as a superscript indicates the first-best value of that endogenous variable.

Lemma 1. *The maximization problem in* (8)–(11) *has the trivial solution* $X_r^* = X_n^* = Q_1^* = Q_2^* = 0$ *if and only if* $k \geq \bar{k}^*$*, where*

$$\bar{k}^* = (\theta\beta h - c)(t_e - t_a - t_\ell). \tag{12}$$

For $k < \bar{k}^*$*, if* $X_r^* = Q_1^* = 0$ *and* $X_n^* = Q_2^* = 1$ *does not solve the maximization problem, then the solution is* $X_r^* = Q_1^* = Q_2^* = 1$ *and* $X_n^* = 0$*.*

The proof of the lemma, provided in Appendix A, is based on the Kuhn-Tucker conditions from the linear program (8)–(11).

The lemma streamlines the search for the first-best solution when $k < \bar{k}^*$ down to a comparison of two alternatives, investing in population-level capacity at risk or investing in that capacity level not at risk. Substituting $X_r^* = Q_1^* = Q_2^* = 1$ and $X_n^* = 0$ into (8), the government's expected surplus from fully investing at risk is

$$-\beta h t_e + s(\theta\beta h - c)(t_e - t_a) - k. \tag{13}$$

Substituting $X_r^* = Q_1^* = 0$ and $X_n^* = Q_2^* = 1$ into (8), the government's expected surplus from fully investing not at risk is

$$-\beta h t_e + s(\theta\beta h - c)(t_e - t_a - t_\ell) - sk. \tag{14}$$

Setting (13) equal to (14) and rearranging shows that the government is indifferent between the two investment strategies if and only if

$$s(\theta\beta h - c)t_\ell = (1-s)k. \tag{15}$$

The left-hand side reflects the advantage from shifting a unit of capacity from the not-at-risk earlier to the at-risk tranche. The shift allows another individual to be vaccinated during the lag t_ℓ it takes for not-at-risk capacity to start producing after initial approval. The advantage only materializes if the vaccine candidate is approved, with probability s. The right-hand side of (15) refects the option value of waiting to install the unit of capacity. If the vaccine fails to be approved, with probability $1-s$, the firm can save the social cost k of investing in that unit of capacity.

Rearranging equation (15) gives a threshold capacity cost in the first best

$$\hat{k}^* = \frac{s}{1-s}(\theta\beta h - c)t_\ell, \tag{16}$$

such that the government strictly prefers at-risk investment if and only $k < \hat{k}^*$. The following propo-

sition summarizes the preceding analysis.

Proposition 1. *Under maintained assumptions, the following policy is first best. The firm invests in capacity at risk for all $k \in [0, \hat{k}^*]$, invests in capacity not at risk for all $k \in (\hat{k}^*, \bar{k}^*]$, and does not invest for all $k > \bar{k}^*$. Sufficient capacity is installed to cover the population, and the firm utilizes all available capacity in production every period until the pandemic's end after t_e.*

Figure 3 provides a schematic diagram of first-best capacity investment as k varies. The interval of at-risk investing has positive measure if and only if

$$\theta\beta h > c, \tag{17}$$

ensuring that a dose's social value exceeds its production cost, a minimal condition to justify any capacity investment.[10] The interval of not-at-risk investing is nonempty if and only if, in addition to condition (17), condition (2) holds.

Comparative statics effects of parameters on the first-best investment policies can be inferred from the thresholds $\hat{k}^*$ and $\bar{k}^*$. The first best is more likely to involve at-risk capacity investment the higher is $\hat{k}^*$, which is increasing in the probability of success s, efficacy θ, disease transmissibility and harm β and h, and capacity-installation lag t_ℓ, and decreasing in production cost c. The first best is more likely to involve capacity investment of some sort (whether at risk or not) the higher is $\bar{k}^*$, which is increasing in θ, β, h, and t_e, and decreasing in c, t_a, and t_ℓ. Since not-at-risk investment only happens conditional on successful approval, the probability of any sort of investment in the first best is also increasing in s.

4.2. Adding Social Cost of Public Funds (Second Best)

We next turn to the optimal vaccine procurement policy when the government must finance its expenditures with distortionary taxation. Let $\lambda > 0$ denote the social cost of public funds. Each dollar of government spending reduces its surplus by $1 + \lambda$ dollars. If that dollar is spent on a transfer to an agent whose surplus the government internalizes, the agent's surplus gain is credited back to the government, but λ is still lost per dollar transferred. The government thus prefers limiting transfers even to domestic agents to avoid this transaction cost.

[10] We will not introduce (17) as a maintained assumption yet, waiting to introduce a stronger version of the condition (integrating consideration of the social cost of public funds, to be defined) in the next subsection.

We continue to suppose that the government controls all the firm's operations including capacity and output decisions. We continue to suppose that the government faces no asymmetric information, able to observe the random draw of capacity cost k and condition policy on that realization. The government's objective function is the same as in the first-best problem (8) except that everywhere cost parameters k or c appear, they need to be scaled by $1+\lambda$ since the government now directly covers those costs with tax revenue. We will refer to this outcome as the second best, and distinguish values of variables that are optimal in this setting by two stars.

It is straightforward to see that identical analysis to the previous subsection applies here except with cost parameters c and k multiplied by $1+\lambda$. For example, a lemma equivalent to Lemma 1 holds, except that the cutoff capacity cost above which there is no investment is now given by

$$(1+\lambda)\bar{k}^{**} = [\theta\beta h-(1+\lambda)c](t_e-t_a-t_\ell). \tag{18}$$

Dividing through by $1+\lambda$ yields an equivalent expression

$$\bar{k}^{**} = \left(\frac{\theta\beta h}{1+\lambda}-c\right)(t_e-t_a-t_\ell). \tag{19}$$

This is identical to the threshold for a social planner facing no social cost of public funds except that the marginal social benefit of a vaccine dose $\theta\beta h$ has to be discounted by $1+\lambda$ to reflect the fact that that surplus is generated with distortionary taxation.

The same logic suggests that the threshold capacity cost determining whether investment is at-risk or not-at-risk is given by

$$\hat{k}^{**} = \frac{s}{1-s}\left(\frac{\theta\beta h}{1+\lambda}-c\right)t_\ell. \tag{20}$$

For $k<\hat{k}^{**}$, the government orders full investment at risk, providing expected net program benefits

$$s[\theta\beta h-(1+\lambda)c](t_e-t_a)-(1+\lambda)k. \tag{21}$$

For $k>\hat{k}^{**}$, the government orders full investment not at risk, providing expected net program benefits

$$s[\theta\beta h-(1+\lambda)c](t_e-t_a-t_\ell)-s(1+\lambda)k. \tag{22}$$

As expected, these expressions for expected net program benefits in the second best are identical those for the first best in (13)–(14) except that all cost parameters have been scaled up by $1+\lambda$. The following proposition summarizes the preceding analysis.

Proposition 2. *Under maintained assumptions, the following policy is second best, the optimum when the government can observe the firm's capacity cost and control its operations but faces a social cost of public funds. The firm invests in capacity at risk for all $k \in [0, \hat{k}^{**}]$, invests in capacity not at risk for all $k \in (\hat{k}^{**}, \bar{k}^{**}]$, and does not invest for all $k > \bar{k}^{**}$. Sufficient capacity is installed to cover the population, and the firm utilizes all available capacity in production every period until the end of the pandemic.*

It is immediate that the first best is recovered from the second best in the limit as $\lambda \downarrow 0$. The second best inherits most of the comparative-static properties of the first best. In addition, one can see from its appearance in denominators in the formulas for $\hat{k}^{**}$ and $\bar{k}^{**}$ that the likelihood of both at-risk investment and investment of any sort are both decreasing in λ.

Figure 3 provides a visual comparison of first-best and second-best investment policies. We see that the second-best threshold $\hat{k}^{**}$ between investing at risk and investing not at risk has been shifted down relative to its first-best analogue $\hat{k}$, shrinking the interval of at-risk investment. The second-best cutoff $\bar{k}^{**}$ above which there is no investment has also been shifted down relative to $\bar{k}^{*}$, enlarging the no-investment interval. An increase in λ would exaggerate the shifts shown. As the figure makes clear, on average across the distribution of k, there is too little at-risk investment and too little investment overall in the second best compared to the first best for all $\lambda > 0$. It is also clear from inspecting the axis labels in the figure that the interval in which there is at-risk investment in the second best is nonempty if and only if

$$\theta\beta h > (1+\lambda)c. \tag{23}$$

This condition says that the marginal social value of a dose justifies the production cost even if that cost has to be paid with funds raised with distortionary taxation. We will maintain assumption (23) throughout the remainder of the paper to avoid a taxonomy of trivial cases.

The government can obtain the same outcome via simple contracts without having to control the firm's operations directly. The contracts need only pair a per-dose price with capacity subsidies. The per-dose price can be set at c to exactly cover production costs. The only potential source of firm profit then is the capacity subsidy, which can be tied to the realization of k, since the government

observes this, and in turn tied to the government's preferred investment timing conditional on k. For $k \in [0, \hat{k}^{**}]$, capacity subsidy k can be offered only for capacity installed at risk. For $k \in (\hat{k}^{**}, \bar{k}^{**}]$, capacity subsidy k can be offered only for capacity installed not at risk or, equivalently, only conditional on vaccine approval.

4.3. Adding Private Information (Third Best)

We next turn to the optimal vaccine procurement policy when the government continues to face a social cost of public funds but now also faces asymmetric information, no longer observing the firm's capacity cost. To avoid possible inconsistencies that might arise if players were allowed to choose strategies contingent on information they do not have, we no longer assume that the government controls the firm's operations directly but now must induce actions indirectly via contracts.

We continue to suppose unit capacity cost k is a random variable drawn from a distribution with pdf $f(k)$, logconcave cdf $F(k)$, and support given by the nonnegative real numbers. We now suppose that the draw of k is privately observed by the integrated firm. The government moves first, offering a procurement contract. Appealing to the revelation principle, we will look for the optimal contract in the set of revelation contracts, under which the firm announces its type $\tilde{k}$ and receives the provisions specified in the initial contract offer for that announcement. As an analytical approach, we will allow rich contractual provisions some of which we will not end up needing in the optimum. This approach allows us to claim that the quite simple form taken by the optimal contract is not due to some exogenous restriction but is without loss of generality.

To that end, consider the following contractual provisions as functions of the firm's announced capacity cost: at-risk and not-at-risk capacity required to install $X_r(\tilde{k})$ and $X_n(\tilde{k})$, output from those respective capacities $Q_1(\tilde{k})$ and $Q_2(\tilde{k})$, bonus price $p_1(\tilde{k})$ within t_ℓ periods after approval and sustained price $p_2(\tilde{k})$ afterwards, and per-unit subsidies for at-risk and not-at-risk capacity $\sigma_r(\tilde{k})$ and $\sigma_n(\tilde{k})$.

Despite the richness of potential contracts, a simple mechanism attains the optimum, as the proof of the next proposition shows. The contract specifies two cutoffs $\hat{k}$ and $\bar{k}$. For $k \leq \hat{k}$, the contract requires the firm invest in sufficient at-risk capacity to serve the whole population each period—and produces up to this capacity each period—in return for an up-front capacity subsidy of σ_r. For $k \in (\hat{k}, \bar{k}]$, the contract requires the firm to invest in sufficient not-at-risk capacity to serve the whole

population each period—and again produces up to this capacity each period—in return for an up-front capacity subsidy of σ_n. Assume the investment and subsidy payment in the not-at-risk case are only made conditional on approval. The firm receives price of c per unit for all doses delivered for either investment form.

The rather extensive proof, provided in the appendix, establishes that restricting the contract to this simple form is without loss of generality. Having done so, we can progress to deriving the optimal contract of this form. Let $\pi_r(k)$ and $\pi_n(k)$ denote, respectively, the firms' joint profit from investing at risk and not at risk. Given the per-dose price equals c, we have

$$\pi_r(k) = \sigma_r - k \tag{24}$$

$$\pi_n(k) = s(\sigma_n - k) \tag{25}$$

The government's expected net benefit from offering a contract of this form equals

$$\int_0^{\hat{k}} \Big\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a) - (1+\lambda)\sigma_r + \pi_r(k) \Big\} f(k)dk \\ \int_{\hat{k}}^{\bar{k}} \Big\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a - t_\ell) - s(1+\lambda)\sigma_n + \pi_n(k) \Big\} f(k)dk. \tag{26}$$

The contract must satisfy incentive-compatibility constraints, ensuring types that are supposed to invest at risk do so instead of investing not at risk, and vice versa:

$$\pi_r(k) \geq \pi_n(k) \;\; \forall k \in [0, \hat{k}] \tag{27}$$

$$\pi_n(k) \geq \pi_r(k) \;\; \forall k \in (\hat{k}, \bar{k}]. \tag{28}$$

The contract must also satisfy individual-rationality constraints, ensuring types earn nonnegative expected profit:

$$\pi_r(k) \geq 0 \;\; \forall k \in [0, \hat{k}] \tag{29}$$

$$\pi_n(k) \geq 0 \;\; \forall k \in (\hat{k}, \bar{k}]. \tag{30}$$

As usual, incentive compatibility binds only for the more not less productive types, so we can

ignore constraints (28). Also as usual, individual rationality binds only for the less not more productive types, so we can ignore constraints (29). In view of the profit expressions in (24)–(25), one can show that a necessary and sufficient condition for the suite of constraints to hold is that they hold at the upper boundary:

$$\pi_r(\hat{k}) \geq \pi_n(\hat{k}) \tag{31}$$

$$\pi_n(\bar{k}) \geq 0. \tag{32}$$

These constraints bind at an optimum. Treating the constraints as equalities, substituting from (24)–(25), and solving the system of equations for the capacity subsidies yields

$$\sigma_r = (1-s)\hat{k} + s\bar{k} \tag{33}$$

$$\sigma_n = \bar{k}. \tag{34}$$

Substituting from (33)–(34) into (26) and rearranging yields a new expression for the government's objective function

$$\int_0^{\hat{k}} \Big\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a) - \lambda[(1-s)\hat{k} + s\bar{k}] - k \Big\} f(k)dk + \int_{\hat{k}}^{\bar{k}} \Big\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a - t_\ell) - \lambda s\bar{k} - sk \Big\} f(k)dk. \tag{35}$$

Let $\hat{k}^d$ denote the optimal threshold between the types that invest at risk and not at risk, with the superscript indicating the third best in the case here under domestic input supply. Taking the first-order condition of (35) with respect to $\hat{k}$, rearranging, and substituting the definition of $\hat{k}^{**}$ from (20) yields

$$\hat{k}^d = \hat{k}^{**} - \left(\frac{\lambda}{1+\lambda} \right) \frac{F(\hat{k}^d)}{f(\hat{k}^d)}. \tag{36}$$

Similarly, let $\bar{k}^d$ denote the optimal threshold between the types that invest not at risk and do not invest. Taking the first-order condition of (35) with respect to $\bar{k}$, rearranging, and substituting the

definition of $\bar{k}^{**}$ from (19) yields

$$\bar{k}^d = \bar{k}^{**} - \left(\frac{\lambda}{1+\lambda}\right)\frac{F(\bar{k}^d)}{f(\bar{k}^d)}. \tag{37}$$

We have sketched the following proposition. Besides providing more rigorous derivations, the formal proof, provided in Appendix A, verifies that the simple contractual forms we posited as optimal are indeed optimal when an unrestricted set of contracts is allowed.

Proposition 3. *The following is a solution for the third best, the optimal procurement mechanism when the government faces a social cost of public funds and asymmetric information about capacity cost, under domestic input supply. For $k \in [0, \hat{k}^d]$, the firm invests in sufficient at-risk capacity to serve the whole population each period, produces up to available capacity each period until the end of the pandemic, and is paid a per-dose price c and a capacity subsidy $\sigma_r^d = (1-s)\hat{k}^d + s\bar{k}^d$. For $k \in (\hat{k}^d, \bar{k}^d]$, the firm invests in sufficient not-at-risk capacity to serve the whole population each period until the end of the pandemic, produces up to available capacity each period, and is paid a per-dose price c and a capacity subsidy $\sigma_n^d = \bar{k}^d$.*

It is straightforward to compare the the third-best thresholds for at-risk investment $\hat{k}^d$ and for any investment $\bar{k}^d$ to their second-best analogues, $\hat{k}^{**}$ and $\bar{k}^{**}$, since they differ only by the last term in (36) and (37), which has an intuitive form. Incomplete information leads the government to distort the thresholds downward, involving less at-risk investment and less investment overall. The distortion is increasing in the social cost of public funds λ. The distortion term is decreasing in the ratio $f(k)/F(k)$ at the respective thresholds, reflecting the relative likelihood of being near the threshold where the benefit of expanding the threshold is experienced versus being among the inframarginal firm types away from the threshold. An expanded threshold requires a greater capacity subsidy for the marginal type to break even, raising the rents earned by inframarginal types. Asymmetric information therefore reduces the likelihood of both at-risk and overall investment.

4.4. Unintegrated Input Supply

In this subsection, we consider an alternative organizational structure for the firm. Instead of taking input supplier U and vaccine manufacturer D to be two divisions within the same firm as done in the previous section, in this section we will take them to be separate firms, each maximizing their individual profits. The case of unintegrated (domestic) input supply is shown in the top left panel of Figure 1.

A variety of models have been proposed for transaction costs and/or contractual frictions between firms, and it is hard to say which is the leading one. To avoid having a specific transactions-cost model drive the results, perhaps limiting their generality, we will suppose that U and D engage in efficient contracting, arrived at via Nash bargaining. Assume bargaining takes place early enough ex ante to allow them to contract on the efficient capacity decision and the efficient production decision. In our partial-equilibrium setting, it is natural to focus on profits earned from sales to the government under consideration and abstract from external operations. This is accomplished in our Nash-bargaining protocol by setting each firm's threat points to 0. Let $\phi \in [0,1]$ denote U's bargaining share and $1-\phi$ denote D's. When we move to a structure in which U is located in a foreign country, ϕ will represent the foreign firm's bargaining share, a useful mnemonic device.

It is obvious that a move from an integrated to an unintegrated supplier does not change the first-best outcome if the government is allowed full control over both firm's operations. Nothing about the firm's production technology is changed, and the contractual relationship between firm divisions is irrelevant to government control. Similarly, the move from integrated to unintegrated supplier does not change the second-best outcome if the government is allowed full control over both firm's operations.

With a moment's reflection, one can see that the move from integrated to unintegrated supplier does not change the equilibrium with privatized firms. Whatever contract the government offers D will pass through to U via efficient bargaining to induce joint-profit-maximizing capacity decisions, the same as the integrated firm.

5. Foreign Input Supply

In this section, we continue to suppose that downstream firm D is located domestically but now move U to a foreign country. The structure is depicted in the right panel of Figure 1. The first subsection takes the firms' locations as exogenous and solves for the equilibrium without international cooperation. The second subsection solves for the outcome with international cooperation. The last subsection endogenizes the location of the input supplier.

5.1. Equilibrium Without International Cooperation

The previous section argued that changing the firm's organizational structure without changing its location has no direct effect on firm operations—by construction since we assumed firms bargain efficiently. Moving U to a foreign country has no direct effect on firms' operation because of efficient bargaining between them. However, we will uncover two indirect effects of the location change. First, the location change will reduce the government's willingness to let rents flow to the firm. Assuming that the government internalizes the surplus of domestic consumers and domestic producers only, transfers flowing to foreign firms will not be credited to the government's "surplus ledger." Rather than subtracting only the λ per dollar transferred, as the case with a domestic firm, the government subtracts off $1+\lambda$ for each dollar transferred to a foreign firm. This will more strongly incline the government to economize on transfers, leading it to incentivize less at-risk investment. We will call this the subsidy-leakage effect.

A second effect is that the foreign government may restrict exports of essential pandemic commodities, possibly including vaccine inputs. Such controls are the foreign government's sovereign right, and thus can override any private contractual provisions. In effect, such export restrictions results in hold up à la Grout (1984) of advance capacity investment and subsidies. Let $\xi \in [0,1]$ be the exogenous probability that restrictions in the foreign country in which U is located prevent it from exporting vaccine inputs to D. Assume that imposition of export restrictions is a draw from a Bernoulli distribution with probability ξ realized at the same date t_a as vaccine approval. That the uncertainty around export restrictions is resolved exactly at date t_a is not important. All that matters is that at-risk investment is undertaken when there is still uncertainty about the export restriction, so can be subject to hold up, while not-at-risk investment is undertaken after that uncertainty is resolved, so can be conditioned on the export restriction and thus is not subject to hold up. The case in which the foreign government commits to export restrictions can be captured in the model by simply setting $\xi = 1$. Interior values $\xi \in (0,1)$ might reflect the foreign government's imperfect ability to commit not to impose export restrictions, depending on the strength of institutions respecting foreign contracts, the intensity of the pandemic in the foreign country, and the benefits from retaining essential pandemic commodities for use by its own citizens.

We will move straight to solving for the optimal procurement contract in the equivalent of the third-best environment in which the government faces a social cost of public funds and asymmet-

ric information, now with foreign input supply. With domestic input supply, the optimal contract reduced to specifying two cutoffs $\hat{k}$ and $\bar{k}$, such that U invests at risk up to the population level for $k \in [0,\hat{k}]$, invests not at risk up to the population level if $k \in (\hat{k},\bar{k}]$, and does not invest if $k > \bar{k}$. The contract sets the per-dose price at production cost c and provides unit capacity subsidies σ_r for at-risk investment and σ_n for not-at-risk investment. This contractual form remains optimal with foreign input supply. With domestic input supply, not-at-risk investment was only undertaken conditional on successful approval. With foreign input supply, that investment is also conditional on the foreign government's not imposing export restrictions.

The inclusion of the subsidy leakage and possible export restriction lead to only slight modifications of the government's objective function from (26), becoming

$$\int_0^{\hat{k}} \Big\{ s(1-\xi)[\theta\beta h-(1+\lambda)c](t_e-t_a)-(1+\lambda)\sigma_r+(1-\phi)\pi_r(k) \Big\} f(k)dk$$
$$\int_{\hat{k}}^{\bar{k}} \Big\{ s(1-\xi)[\theta\beta h-(1+\lambda)c](t_e-t_a-t_\ell)-s(1-\xi)(1+\lambda)\sigma_n+(1-\phi)\pi_n(k) \Big\} f(k)dk. \quad (38)$$

Since the government does not internalize the foreign firm's bargaining share ϕ, firm profits enter the government's objective function scaled by $1-\phi$. For capacity to generate successful output requires not just approval, which occurs with probability s, but no export restrictions to be imposed, which occurs with probability $1-\xi$. Hence, the other change from (26) in (38) is that everywhere probability s appeared with domestic supply, it now must be replaced by the joint probabilty $s(1-\xi)$. Expected firm profits are now given by

$$\pi_r(k) = \sigma_r - k \quad (39)$$
$$\pi_n(k) = s(1-\xi)(\sigma_n - k) \quad (40)$$

The analysis proceeds just as in Section 4.3. The only binding constraints are (31) and (32). Treating the constraints as equalities, substituting from (24)–(25), and solving the system of equations for the capacity subsidies yields

$$\sigma_r = [1-s(1-\xi)]\hat{k}+s(1-\xi)\bar{k} \quad (41)$$
$$\sigma_n = \bar{k}. \quad (42)$$

We proceed by substituting from (41)–(42) into (38) and taking the resulting first-order conditions with respect to the investment thresholds. Letting $\hat{k}^f$ and $\bar{k}^f$ denote the optimal investment thresholds in the third best with foreign input supply, after rearranging the first-order conditions, we obtain

$$\hat{k}^f = \hat{K}(c,\lambda,s(1-\xi)) - \left(\frac{\phi+\lambda}{1+\lambda}\right)\frac{F(\hat{k}^f)}{f(\hat{k}^f)} \tag{43}$$

$$\bar{k}^f = \bar{K}(c,\lambda) - \left(\frac{\phi+\lambda}{1+\lambda}\right)\frac{F(\bar{k}^f)}{f(\bar{k}^f)}, \tag{44}$$

where

$$\hat{K}(c,\lambda,s) = \frac{s}{1-s}\left(\frac{\theta\beta h}{1+\lambda} - c\right)t_\ell \tag{45}$$

$$\bar{K}(c,\lambda) = \left(\frac{\theta\beta h}{1+\lambda} - c\right)(t_e - t_a - t_\ell). \tag{46}$$

These expressions resemble the first-, second-, and third-best investment thresholds derived previously, indeed nesting them as follows: $\hat{k}^* = \hat{K}(c,0,s)$, $\hat{k}^{**} = \hat{k}^d = \hat{K}(c,\lambda,s)$, $\bar{k}^* = \hat{K}(c,0)$, and $\bar{k}^{**} = \bar{k}^d = \bar{K}(c,\lambda)$. We have the following proposition.

Proposition 4. *The following is a solution for the third best, the optimal procurement mechanism when the government faces a social cost of public funds and asymmetric information about capacity cost, under foreign input supply. For $k \in [0,\hat{k}^f]$, the firm invests in sufficient at-risk capacity to serve the whole population each period, produces up to available capacity each period until the end of the pandemic, and is paid a per-dose price c and a capacity subsidy $\sigma_r^f = [1-s(1-\xi)]\hat{k}^f + s(1-\xi)\bar{k}^f$. For $k \in (\hat{k}^f,\bar{k}^f]$, the firm invests in sufficient not-at-risk capacity to serve the whole population each period until the end of the pandemic, produces up to available capacity each period, and is paid a per-dose price c and a capacity subsidy $\sigma_n^f = \bar{k}^f$.*

The investment thresholds under domestic input supply, $\hat{k}^d$ and $\bar{k}^d$, are identical to those under foreign input supply, $\hat{k}^f$ and $\bar{k}^f$, except that factors ϕ and ξ have been zeroed out. Since the right-hand side of (43) and (44) are decreasing in ϕ and furthermore $\partial\hat{K}/\partial s > 0$ when (23) holds, we can see that $\hat{k}^f < \hat{k}_d$ and $\bar{k}^f < \bar{k}^d$, implying that both at-risk investment and investment of any sort are both less likely under foreign than domestic input supply. The domestic government's desire to limit leakage of the subsidy to a foreign firm and wastage of the subsidy on vaccine inputs that are barred from export reduce the capacity subsidy it offers, thus reducing U's investment incentives.

5.2. International Cooperation

We saw that moving the input supplier from the domestic to a foreign country introduced two sources of distortion, one due to the domestic government's desire to limit leakage of the subsidy to the foreign firm, another due to the threat and execution of export restrictions, reducing equilibrium capacity subsidies and curtailing vaccine supply in the domestic country if the restriction is enforced.

There is thus scope for international cooperation to improve global surplus by removing these distortions. Suppose that cooperation comes in the form of an agreement among countries in advance of any specific pandemic to principles that apply to all future pandemics. Countries reach the agreement behind the veil of ignorance, so come to an agreement that maximizes expected global surplus. An alternative interpretation is that countries sign a treaty at the outset of a pandemic already underway but early on, before investment is undertaken. One complication with that interpretation is that for negotiation among countries to generate a treaty that maximizes expected global surplus may require transfers between countries, which may introduce distortions due to the social cost of public funds. We can ignore this source of distortion if we consider the transferred money to reduce the tax burden in the recipient country, reducing the tax distortion there, offsetting the distortion in the country making the transfer.

A treaty that works to maximize expected global surplus has the domestic government agree to increase the subsidy it provides for capacity investment from the equilibrium levels absent international cooperation,

$$\sigma_r^f = [1-s(1-\xi)]\hat{k}^f + s(1-\xi)\bar{k}^f \tag{47}$$

$$\sigma_n^f = \bar{k}^f \tag{48}$$

up to the levels that would have been provided to a domestic input supplier in the absence of any threat of export restrictions,

$$\sigma_r^d = (1-s)\hat{k}^d + s\bar{k}^d \tag{49}$$

$$\sigma_n^d = \bar{k}^d. \tag{50}$$

The foreign government for its part agrees not to impose export restrictions.

It is immediate that this form of cooperation removes both sources of distortion due to foreign input supply. The treaty increases expected global surplus from its level in the third best under foreign input supply to its level in the third best under domestic input supply.

5.3. Endogenizing Location of Input Supplier

So far, the analysis has assumed that the location of the input supplier is exogenous. The input supplier happens to be located in a given country, perhaps because of historical accident or the availability of specialized inputs into its production process, and it cannot be easily moved during a pandemic. In this subsection, we analyze the alternative assumption that there is sufficient time during the lag period required to install capacity to move the input supplier to the most advantageous location. There is evidence for the relevance of this alternative assumption, since as we described in Section 2, to sell to different markets, most vaccine sponsors chose to set up parallel supply chains, with drug substance and fill-and-finish plants rarely located in separate countries

Assume that, as part of its procurement contract, the domestic government can dictate that the input supplier locate in the domestic or foreign country. If the location of the input supplier is indeed flexible, it is not necessary that the goverment dictate the location. It could arrive at the same outcome indirectly by conditioning the capacity subsidy on the location of the input supplier, ensuring that the input supplier located as desired by eliminating profits for a firm that chose the "wrong" location.

In the absence of any modeled benefit of locating in the foreign country, it is immediate that the domestic government would require the input supplier to locate domestically as this would eliminate the distortions from foreign location without any drawbacks. To make the decision more interesting, we will introduce a tradeoff associated with input-supplier location, supposing that locating abroad reduces production cost from c to $c-\Delta$, where $\Delta \in [0,c]$. It raises no difficulties for the analysis to allow Δ to be negative, but then there would be no tradeoff in foreign location and the decision by the government to require domestic input location would be trivial. The analysis would be similar if we modeled the tradeoff involved between domestic and foreign location to reside in parameters other than c.

In the presence of international cooperation, the government would require the input supplier to locate according to where production costs are lowest, the foreign country if $\Delta > 0$. Any offsetting

distortions due to insufficient capacity subsidies or export restrictions would be eliminated by treaty.

In the absence of international cooperation, the government could evaluate its expected net surplus from two possible contracts and select the location offering the higher value. It could require domestic input supply. This would generate high production costs but the efficient subsidy conditional on domestic location and no export restriction. On the other hand, it could require (or allow) foreign input supply. This would result in lower production costs if $\Delta > 0$ but an inefficiently low capacity subsidy (due to subsidy leakage and the threat of export restriction) and possible supply curtailment if the other government imposes an export restriction. Fixing $\Delta \in (0, c)$, the government would require foreign input supply for ϕ and ξ sufficiently close to 0. Fixing $\Delta \in (0, c)$ and $\phi \in [0, 1]$, the government would require domestic input supply for ξ sufficiently close to 1.

Notice that domestic input supply can be inefficient from a global perspective even though, on the equilibrium path, there is no distortion to the capacity subsidy (it is efficient conditional on domestic input supply) and no hold up due to export restrictions. Those distortions may be great enough to deter the government from choosing the lower-cost foreign location. The distortion then would show up as a distorted extensive-margin decision to choose the high-production-cost path, with the distorted capacity subsidy and export restrictions themselves remaining off equilibrium.

One complication in the analysis of equilibrium with endogenous input-supply location is that it may lead to a complex set of location strategies. The firm may locate its at-risk investment in the foreign country but move its not-at-risk investment domestically if it learns that the foreign government is holding up investment with export restrictions. In Appendix B (in process), we work through the various location strategies and characterize fully the possible role of international cooperation when the location of the input supplier is endogenous. There we confirm that a role for international cooperation over capacity subsidies and export restrictions may arise even when governments choose local sourcing for their vaccine supply chains and the distortions in capacity subsidies and export restrictions remain off equilibrium.

6. Calibrations

The theoretical analysis has left open the question of whether any of the identified distortions are large enough to be a practical concern and if so, which are the most important, whether the social cost

of public funds, asymmetric information, or one or another distortion associated with foreign input supply. To answer these questions, in this section we turn to calibrations of the model combining empirical with plausible choices for parameter values.

We will calibrate the model to the case of a vaccine procurement program during the Covid-19 pandemic run by the US government. We will take a period to be six months, about the length of time between boosters and Covid-19 waves. Suppose the lag in capacity installation and in approval is about six months, so one period, i.e., $t_a = t_\ell = 1$. Suppose the pandemic lasts for three years, i.e., $t_e = 6$. Set vaccine efficacy to $\theta = 0.75$, the averge efficacy against symptomatic Covid-19 according to the Ssentongo et al. (2022) meta-analysis. Set the social cost of public funds to $\lambda = 0.3$, the level assumed in Laffont and Tirole (1993), based on an estimate of the deadweight loss of taxation in developed countries by Snow and Warren (1995). Set $s = 0.4$, the estimate of the probability of success of an industry-sponsored vaccine program from phase-1 trials to approval from a study of over 1,800 such programs (Lo, Siah, and Wong, 2020). Suppose firms have symmetric bargaining power, implying $\phi = 0.5$.

We have little guidance on the appropriate value of ξ, the probability that the foreign government imposes an export restriction. This will depend on the particular country, the scarcity of inputs for its own uses at the time, and its bilateral relationship with the domestic country, among other idiosyncratic factors. Following the principle of maximimum entropy (Jaynes, 1957), we set it to the entropy-maximizing value for a Bernoulli random variable, $\xi = 0.5$. We also examine $\xi = 0$ as a counterpoint.

The remaining inputs into the calibration regard the vaccine's benefits and costs. We draw on Snyder et al. (2020) for these. The authors combine economic output losses and mortality losses to arrive at an estimate of the harm from the Covid-19 pandemic. They base their estimate of economic harms from the Covid-19 pandemic on International Monetary Fund projections of growth shortfalls (International Monetary Fund, 2020). Their estimate of mortality harm starts with Walker et al. (2020) projections of deaths by country, converted into a monetary value by taking each death to result in a loss of 12 disability adjusted life years (DALYs) according to estimates from Hanlon et al. (2021), and taking each DALY to cost three times per-capita GDP in the country according to the WHO standard of cost-effective health interventions (Marseille et al., 2015). For the US, Snyder et al. (2020) calculate that the combined harm of $360 per capita per month, or $2,160 per six-month

period. We take this to be the expected harm βh relieved by a dose of the vaccine, abstracting from the offsetting factors that pandemic harms include broader losses to education and human capital (Cutler and Summers, 2020) on the one hand but that a vaccine will not fully relieve these harms if it does not eradicate the disease.[11] We assume that the US government seeks to vaccinate 250 million citizens, about 75% population coverage.

Based on data from a CEPI survey of nearly 100 potential manufacturers of a Covid-19 vaccine, Snyder et al. (2020) estimate lognormal distributions for Covid-19 vaccine capacity and production costs. Their estimates require two adjustments for our purposes. First, their estimates are for a larger international program than the country program considered in this calibration; second, they allow for randomness in capacity and production costs, whereas our model concentrates all the uncertainty solely in the capacity cost. Omitting the randomness in the full range of costs would understate the influence of asymmetric information. We proceed by using the component estimates provided in Appendix Exhibit A4 in Snyder et al. (2020) to estimate an all-in cost of installing 250 million doses of capacity and running this for one production period and allocate the result solely to capacity costs. The estimated lognormal distribution of k that results has a mean of \$13.4 per dose and standard deviation of \$13.6 per dose. Since we have already allocated some production cost to k, we set $c = 10$, slightly lower than the sample mean of around \$13.

Table 1 reports the calibration results. Expected net benefits from the US program equal \$467 billion in the first best, falling to \$384 billion in the second best when the $\lambda = 0.3$ social cost of public funds is added. The probability that k falls in the range in which at-risk investment is induced is only 32.4% in the first best, falling to 22.0% in the second best. The probability of any investment is 60.4%. Since the vaccine is so valuable relative to its expected cost in the first best, any shortfall of this probability from 1 is almost completely driven by approval failure. Vaccine failure precludes any not-at-risk investment. Adding in a social cost of public funds in the second best reduces the probability of any investment slighty to 51.5%.

Moving from the second best to the third best adds the asymmetric-information friction, resulting in about an equal decline in expected net program benefits as the distortion due to the social cost of public funds. The probability of at-risk investment falls to a paltry 17.9%.

The third best with foreign input supply is reflected in two scenarios. The first foreign sce-

[11] The \$4,320 benefit per capita annually that we use is conservative compared to the \$5,800 benefit of an annual course of capacity estimated by Castillo et al. (2021).

nario includes only the subsidy-leakage effect via a positive value of the foreign firm's profit share, $\phi = 0.5$. The second foreign scenario is the comprehensive one, also including the distortion due to possible export restrictions ($\xi = 0.5$). Relative to the third best with domestic input supply, expected net program benefits fall by about 19% due to the subsidy-leakage effect and another 42% due to the export-restriction effect. Thus, of the combined distortions due to foreign input supply, about a third can be credited to the subsidy-leakage effect and two thirds to the export-restriction effect. Both are important but export restrictions can be particularly damaging to surplus. By the time the comprehensive scenario of foreign input supply is reached, at-risk investment has all but disappeared, occuring with 1.1% probability.

The benefit of international cooperation can be measured in the calibration as the move from the third best with foreign input supply, generating net benefits of $125 billion, to the third best with domestic input supply, generating net benefits of $320 billion, a gain of $195 billion, over 1.5 times the equilibrium surplus without cooperation. This is just the benefits enjoyed by the US and the firm in the foreign country. The level of global benefits would be orders of magnitude higher.

7. Conclusion

Potentially millions of lives and trillions of dollars of economic activity could have been saved with accelerated production of Covid-19 vaccines. Why governments did not intervene with the sorts of policies needed to align private and social incentives for at risk investment in vaccine manufacturing during the pandemic remains a puzzle.

To help investigate this puzzle, we have constructed a theoretical model in which a country's government procures a vaccine from a firm to protect its citizens against pandemic harm. The firm can accelerate the vaccine's availability by investing before the vaccine is approved, but this risks stranding investment if approval fails. The government can encourage such at-risk investment by adding capacity subsidies on top of its procurement price. In this setting, we have analyzed the optimal procurement mechanism under asymmetric information about the firm's cost and where the government faces a social cost of public funds, and we have studied how this mechanism changes when the supplier of essential vaccine inputs is located in a foreign country. When the input is supplied by a foreign firm rather than locally sourced and in the absence of international cooperation

on vaccine supply chain policies, we have shown that the government cuts back on the subsidy to avoid the leakage of information rents to a firm whose profits it does not internalize. A second distortion that also arises when vaccine inputs are offshored and policy is set noncooperatively is that the foreign government may restrict exports in a crisis, holding up the capacity subsidy. These two distortions will be present on the equilibrium path if countries choose to rely on offshoring in their presence, but we argue that these distortions may continue to induce costly inefficiencies even if countries have chosen to locally source vaccine inputs and these distortions are moved off the equilibrium path, provided that offshoring of inputs would be an attractive way to reduce costs but for these distortions. Our calibrations using estimates of the costs and benefits of a Covid-19 vaccine suggest that overall the distortions identified by our model can be enormous, and that international cooperation aimed at reducing them can result in a many-fold improvement in program net benefits.

Our results therefore suggest that a role for an enforceable international agreement on vaccine supply chain policies may arise, especially when the offshoring of inputs is a potentially attractive way to reduce costs. And while this role arises as a result of international externalities associated with vaccine supply chain policy, it is important to note that the international externalities featured in our model are all pecuniary; non-pecuniary externalities, such as cross-border transmission of a disease, are ruled out in our model by construction. This suggests in turn that the World Trade Organization (WTO), rather than the World Health Organization (WHO), is the appropriate forum for the negotiation of such an agreement. At the same time, since our results indicate that increases in subsidies could be a prominent feature of such an agreement, the agreement might be best interpreted as introducing for vaccine-related subsidies a potential carve out from the existing WTO subsidy rules, which are generally focused on disciplining national use of subsidies, not cooperatively increasing them to globally efficient levels.[12]

We are not the first to suggest a possible role for international cooperation on vaccine supply chains. For example, in a September 21 2021 Opinion Piece for the *New York Times*, Jeneen Interlandi, a member of the editorial board, put the point this way:

> Pharmaceutical companies generally know how to coordinate their global supply chains. They also know how to work together to secure the resources they need to make their products. But when the situation requires changes to national and global policy,

[12]That said, this feature of WTO subsidy rules poses something of a puzzle for the economic analysis of trade agreements more generally. See, for example, Bagwell and Staiger (2001, 2002, 2006, 2012) and Sykes (2005).

> world leaders need to step in. So far, they have not. For all its successes, the race to vaccinate the world against Covid-19 has unfolded like a symphony without a conductor. The corralling of manufacturing sites has been haphazard. The channeling of equipment and ingredients has been messy and at times wasteful. And the flow of vaccines has been recklessly uneven: More than 80 percent of the four billion vaccine doses that had been distributed as of early August went to high- and upper-middle-income countries. . . .
>
> Boosters for the wealthy and scraps for everyone else will neither get us out of this pandemic nor prepare us for the next one. But nearly a year since the first shots were administered, world leaders have yet to put forth a bolder or more comprehensive plan. "Nobody is saying unequivocally, 'Here is what we need, and here is how we are going to get it,' " said Zain Rizvi, a health law expert at the consumer advocacy nonprofit Public Citizen. "We were promised a war effort, and instead we got a pillow fight." (Interlandi, 2021)

Our findings provide some guidance for this general call to action, by suggesting that a focus specifically on international efforts to cooperate over subsidies to at-risk capacity investments and avoid export restrictions on vaccine inputs could be especially fruitful.

Given the public health and economic costs of the pandemic, our findings also raise a number of questions for future research. The first and most important is clearly empirical. While our model is consistent with a number of stylized facts on vaccine supply chains and the policy environment to emerge from Covid-19, and while our calibration exercise suggests that the distortions featured in our theoretical analysis and the benefits from international cooperation that we have identified in the presence of these distortions could be quantitatively important, the question remains: How important are the channels identified here relative to other economic, political, and public health factors impacting at-risk investment and vaccine production decisions? Second, while the results here suggest a novel form of international policy cooperation—multilateral commitments not to impose export restrictions, and also commitments to coordinated expansionary subsidy policies along a supply chain—are there other WTO principles that might inform policy makers on how to negotiate such an agreement in practice? Third, given the prisoner's dilemma nature of the problem, how might the agreement be enforced to prevent unilateral defections and increase the chance of compliance in its time of need during the next public health emergency?[13] We leave answers to these and other related questions to future research.

[13]The difficulty of enforcing international commitments during a pandemic is emphasized, for example, in Staiger (2022, chapter 12).

Appendix A. Proofs of Propositions

This appendix provides proofs omitted from the text for space considerations.

Proof of Lemma 1

The maximization problem (8)–(11) is a linear program for which the Kuhn-Tucker conditions are necessary and sufficient for an optimum. Letting γ_1, γ_2, μ_1, and μ_2 be Lagrange multipliers, the associated Lagrangian is

$$\begin{aligned}\mathscr{L} = sQ_1(\theta\beta h - c)t_\ell + sQ_2(\theta\beta h - c)(t_e - t_a - t_\ell) - kX_r - skX_n \\ + \gamma_1(X_r - Q_1) + \gamma_2(X_r + X_n - Q_2) + \mu_1(1 - Q_1) + \mu_2(1 - Q_2). \quad \text{(A1)}\end{aligned}$$

The Kuhn-Tucker conditions include inequalities bounding partials of the Lagrangian,

$$\frac{\partial \mathscr{L}}{\partial X_r} = -k + \gamma_1^* + \gamma_2^* \leq 0 \quad \text{(A2)}$$

$$\frac{\partial \mathscr{L}}{\partial X_n} = -sk + \gamma_2^* \leq 0 \quad \text{(A3)}$$

$$\frac{\partial \mathscr{L}}{\partial Q_1} = s(\theta\beta h - c)t_\ell - \gamma_1^* - \mu_1^* \leq 0 \quad \text{(A4)}$$

$$\frac{\partial \mathscr{L}}{\partial Q_2} = s(\theta\beta h - c)(t_e - t_a - t_\ell) - \gamma_2^* - \mu_2^* \leq 0, \quad \text{(A5)}$$

complementary slackness conditions associated with nonnegativity constraints on the choice variables,

$$X_r^*(-k + \gamma_1^* + \gamma_2^*) = 0 \quad \text{(A6)}$$

$$X_n^*(-sk + \gamma_2^*) = 0 \quad \text{(A7)}$$

$$Q_1^*[s(\theta\beta h - c)t_\ell - \gamma_1^* - \mu_1^*] = 0 \quad \text{(A8)}$$

$$Q_2^*[s(\theta\beta h - c)(t_\ell - t_a - t_\ell) - \gamma_2^* - \mu_2^*] = 0, \quad \text{(A9)}$$

complementary slackness conditions associated with constraints (9)–(11),

$$\gamma_1^*(X_r^* - Q_1^*) \quad \text{(A10)}$$

$$\gamma_2^*(X_r^* + X_n^* - Q_2^*) \quad \text{(A11)}$$

$$\mu_1^*(1 - Q_1^*) \quad \text{(A12)}$$

$$\mu_2^*(1 - Q_2^*), \quad \text{(A13)}$$

as well as nonnegativity of Lagrange multipliers, nonnegativity of the choice variables, and constraints (9)–(11) themselves.

The proof proceeds by analyzing an series of exhaustive and mutually exclusive cases.

Case (i) Suppose $X_r^* = X_n^* = 0$. Then $0 \leq Q_2^* \leq X_r^* + X_n^* = 0$ implies $Q_2^* = 0$, in turn implying $\mu_2^* = 0$ by (A13). Then

$$\begin{aligned} s\bar{k}^* &= s(\theta\beta h - c)(t_e - t_a - t_\ell) && \text{(A14)} \\ &\leq \gamma_2^* + \mu_2^* && \text{(A15)} \\ &= \gamma_2^* && \text{(A16)} \\ &\leq sk, && \text{(A17)} \end{aligned}$$

where the first step follows from the definition of $\bar{k}^*$ in (12), the second step from (A5), the third step from $\mu_2^* = 0$, and the last step from (A3). Hence, $k \geq \bar{k}^*$. The reader can verify that setting choice variables $X_r^* = X_n^* = Q_1^* = Q_2^* = 0$ and Lagrange multipliers $\gamma_1^* = (1-s)k$, $\gamma_2^* = sk$, and $\mu_1^* = \mu_2^* = 0$ satisfy all Kuhn-Tucker conditions and thus provide a solution.

Cases (ii)–(iv) Having provided a solution for $k \geq \bar{k}^*$, for the remainder of the proof suppose $k < \bar{k}^*$. We then have

$$\begin{aligned} \mu_2^* &\geq s(\theta\beta h - c)((t_e - t_a - t_\ell) - \gamma_2^* && \text{(A18)} \\ &\geq s(\theta\beta h - c)((t_e - t_a - t_\ell) - sk && \text{(A19)} \\ &\geq s(\theta\beta h - c)((t_e - t_a - t_\ell) - s\bar{k}^* && \text{(A20)} \\ &= 0. && \text{(A21)} \end{aligned}$$

where the first step follows from (A5), the second step from (A7), and the third step from $k < \bar{k}^*$, and the last step from (12). But $\mu_2^* > 0$ together with (A13) implies

$$Q_2^* = 1. \tag{A22}$$

We will use this fact in the analysis of the cases below.

Case (ii) Suppose $X_r^* > 0$ and $X_n^* = 0$. Since $X_r^* > 0$, $\gamma_1^* + \gamma_2^* = k > 0$ by (A6). Hence, either $\gamma_1^* > 0$ or $\gamma_2^* > 0$. If $\gamma_1^* > 0$, then $X_r^* = Q_1^* \leq 1$, where the last inequality follows from constraint (9). If $\gamma_2^* > 0$, then $X_r^* = X_r^* + X_n^* = Q_2^* = 1$, where the first step follows from $X_n^* = 0$, the second from (A11), and the third from (A22). We have shown $X_r^* \leq 1$ whether $\gamma_1^* > 0$ or $\gamma_2^* > 0$. We can thus sandwich X_r^* as $1 = Q_2^* \leq X_r^* \leq 1$, implying $X_r^* = 1$.

Case (iii) Suppose $X_r^* = 0$ and $X_n^* > 0$. Since $X_r^* = 0$, $Q_1^* = 0$ by (10). Since $X_r^* > 0$, $\gamma_2^* = sk > 0$ by (A7), implying $Q_2^* = X_r^*$ by (A11), in turn implying $X_r^* = 1$ by (A22).

Case (iv) Suppose $X_r^* > 0$ and $X_n^* > 0$. Since $X_n^* > 0$, $\gamma_2^* = sk > 0$ by (A6), implying $1 = Q_2^* = X_r^* + X_n^*$ by (A11), in turn implying $X_r^* < 1$ since $X_n^* > 0$. But then $Q_1^* \leq X_r^* < 1$ implies $\mu_1^* = 0$ by (A12). Then

$$\begin{aligned} \gamma_1^* &\geq s(\theta\beta h - c)t_\ell - \mu_1^* && \text{(A23)} \\ &= s(\theta\beta h - c)t_\ell && \text{(A24)} \\ &> 0, && \text{(A25)} \end{aligned}$$

where the first step follows from (A4), the second step from $\mu_1^* = 0$, and the third step from equation (23). Now $\gamma_1^* > 0$ implies $Q_1^* = X_r^*$ by (A10), implying $Q_1^* > 0$ since $X_r^* > 0$, in turn implying

$$s(\theta\beta h - c)t_\ell = \gamma_1^* \tag{A26}$$

from (A8) and $\mu_1^* = 0$. Since $X_r^* > 0$, $\gamma_1^* + \gamma_2^* = k$, implying $\gamma_1^* = (1-s)k$ since $\gamma_2^* = sk$ as previously shown. Substituting this value of γ_1^* in (A26) yields

$$s(\theta\beta h - c)t_\ell = (1-s)k. \tag{A27}$$

Substituting (A27), $Q_1^* = X_r^*$, $Q_2^* = X_r^* + X_n^*$ into equation (8) and rearranging yields a new expression for the planner's objective function,

$$-\beta h t_e + s(X_r^* + X_n^*)(\theta\beta h - c)(t_e - t_a - t_\ell) - sk(X_r^* + X_n^*). \tag{A28}$$

Only the sum $X_r^* + X_n^*$ is pinned down in an optimum, not the two capacity tranches separately. We showed in the previous paragraph that $X_r^* + X_n^* = 1$. The optimum can be achieved by a linear combination of capacity tranches, including $X_r^* = 0$ and $X_n^* = 1$, which does not involve at-risk investment.

Summary Combining cases (i)–(iv) together, we have shown that, across an exhaustive set of cases, the optimum can be attained by setting $Q_1^* = X_r^* = 0$ and $Q_2^* = X_n^* = 1$ or, if not, by setting $Q_1^* = Q_2^* = X_r^* = 1$ and $X_n^* = 0$. *Q.E.D.*

Proof of Proposition 3

We will refer to a single integrated firm for brevity but the analysis also nests unintegrated firms and both domestic and foreign input supply. Consider a general contract specifying that the firm announces its capacity-cost type $\tilde{k}$. The contract requires the firm to install capacity $X_r(\tilde{k})$ at risk and $X_n(\tilde{k})$ not at risk, where

$$X_r(\tilde{k}) \geq 0 \tag{A29}$$

$$X_n(\tilde{k}) \geq 0 \tag{A30}$$

$$X_r(\tilde{k}) + X_n(\tilde{k}) \leq 1. \tag{A31}$$

Conditional on approval, the firm is required to produce $Q_1(\tilde{k})$ within t_ℓ periods after approval and $Q_2(\tilde{k})$ in the remaining periods, where

$$Q_1(\tilde{k}) \geq 0 \tag{A32}$$

$$Q_2(\tilde{k}) \geq 0 \tag{A33}$$

$$Q_1(\tilde{k}) \leq X_r(\tilde{k}) \tag{A34}$$

$$Q_2(\tilde{k}) \leq X_r(\tilde{k}) + X_n(\tilde{k}). \tag{A35}$$

The firm is paid $p_1(\tilde{k})$ within t_ℓ periods after approval and sustained price $p_2(\tilde{k})$ afterwards. The firm is paid per-unit subsidies for at-risk and not-at-risk capacity $\sigma_r(\tilde{k})$ and $\sigma_n(\tilde{k})$, respectively.

Let $\pi(k,\tilde{k})$ denote the firm's expected profit from the contract when its type is k but it announces $\tilde{k}$. We have

$$\pi(k,\tilde{k}) = R(\tilde{k}) - X(\tilde{k})k, \tag{A36}$$

defining expected revenue paid to the firm

$$R(\tilde{k}) = \sigma_r(\tilde{k})X_r(\tilde{k}) + s\sigma_n(\tilde{k})X_n(\tilde{k}) + s[p_1(\tilde{k}) - c]Q_1(\tilde{k})t_\ell + s[p_2(\tilde{k}) - c]Q_2(\tilde{k})(t_e - t_a - t_\ell) \tag{A37}$$

and expected total capacity installed

$$X(\tilde{k}) = X_r(\tilde{k}) + sX_n(\tilde{k}). \tag{A38}$$

If investment is nontrivial, $X(\tilde{k}) > 0$, then either $X_r(\tilde{k}) > 0$ or $X_n(\tilde{k}) > 0$. In that case, the government can deliver arbitrary revenue $R(\tilde{k})$ to the firm via one or the other of the capacity subsidies $\sigma_r(\tilde{k})$ or $\sigma_n(\tilde{k})$. Without loss of generality, prices can be set to production cost:

$$p_1(\tilde{k}) = p_2(\tilde{k}) = c. \tag{A39}$$

The government designs the contract to maximize expected net program benefits, which upon substituting from (A39) for the prices can be written

$$\int_0^K \Big\{ s[\theta\beta h - (1+\lambda)c][Q_1(k)t_\ell + Q_2(k)(t_e - t_a - t_\ell)] \\ - (1+\lambda)[\sigma_r(k)X_r(k) + s\sigma_n(k)X_n(k)] + (1-\phi)\pi(k,k) \Big\} f(k)dk, \tag{A40}$$

subject to incentive compatibility

$$\pi(k,k) \geq \pi(k,\tilde{k}) \ \ \forall k \geq 0, \tilde{k} \geq 0 \tag{A41}$$

and individual rationality

$$\pi(k,k) \geq 0 \ \ \forall k \geq 0. \tag{A42}$$

An increase in each of $Q_1(k)$ and $Q_2(k)$ increases the objective function while leaving constraints (A41)–(A42) unchanged. To see this, note that $Q_1(\tilde{k})$ and $Q_2(\tilde{k})$ enter $\pi(k,\tilde{k})$ only through $R(\tilde{k})$, which is independent of $Q_1(\tilde{k})$ and $Q_2(\tilde{k})$ when (A39) holds. Hence, we can take (A34) and (A35) to bind without loss of generality. Subsituting for $Q_1(k)$ and $Q_2(k)$ treating (A34)–(A35) as equalities as well as from (A36)–(A39) into (A40) yields, after rearranging,

$$\int_0^K \Big\{ \big\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a) - (1-\phi)k \big\} X_r(k) \\ + s\big\{ [\theta\beta h - (1+\lambda)c](t_e - t_a - t_\ell) - (1-\phi)k \big\} X_n(k) - (\phi+\lambda)R(k) \Big\} f(k)dk. \tag{A43}$$

Using optimal-control techniques that are standard in mechanism design, we can express $R(k)$ in terms of $X_r(k)$ and $X_n(k)$ to reduce the number of controlled functions. Define $\Pi(k) = \pi(k,k)$. Then (A41) implies

$$\Pi(k) = R(k) - X(k)k = \max_{\tilde{k} \geq 0} [R(\tilde{k}) - X(\tilde{k})k]. \tag{A44}$$

By the Envelope Theorem, (A44) implies $\Pi'(k) = -X(k)$. The Fundamental Theorem of Calculus then gives

$$\Pi(k) = \int_k^K X(z)dz - \Pi(K) = \int_k^K X(z)dz, \tag{A45}$$

since $\Pi(K) = 0$ because it is optimal not to leave rents to the least productive type. Then

$$R(k) = \Pi(k) + X(k)k = \int_k^K X(z)dz + X(k)k, \tag{A46}$$

where the first equality follows from (A44) and the second from (A45). Integrating,

$$\int_0^K R(k)f(k)dk = \int_0^K \left[\int_k^K X(z)dz\right] f(k)dk + \int_0^K X(k)kf(k)dk. \tag{A47}$$

Using integration by parts to evaluate the first integral on the right-hand side of (A47),

$$\int_0^K \left[\int_k^K X(z)dz\right] f(k)dk = \left[F(k)\int_k^K X(z)dz\right]_{k=0}^{k=K} + \int_0^K F(k)X(k)dk \tag{A48}$$

$$= \int_0^K F(k)X(k)dk. \tag{A49}$$

Substituting from (A49) into (A47) and then substituting for $X(k)$ from (A38) yields

$$\int_0^K R(k)f(k)dk = \int_0^K F(k)[X_r(k) + sX_n(k)]dk + \int_0^K [X_r(k) + sX_n(k)]kf(k)dk. \tag{A50}$$

We can then use (A50) to substitute out for the term in $R(k)$ in objective function (A43), obtaining

$$\int_0^K \Big\{ \big\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a) - (1+\lambda)k \big\} X_r(k) + s\big\{ [\theta\beta h - (1+\lambda)c](t_e - t_a - t_\ell) - (1+\lambda)k \big\} X_n(k) \Big\} f(k)dk - (\phi+\lambda)\int_0^K F(k)[X_r(k) + sX_n(k)]dk. \tag{A51}$$

We claim that this objective function incorporates the information from constraints (A41)–(A42). We will demonstrate this by maximizing (A51) ignoring those constraints and verifying at the end that the solution satisfies them. Constraints (A41)–(A42) are the only place where different realizations of the random variable k are linked. Ignoring (A41)–(A42) severes this linkage, allowing the problem to be maximized pointwise.

The problem becomes one of finding $X_r(k)$ and $X_n(k)$ maximizing the integrand in (A51) for each k subject to (A29)–(A31). This is a linear program in the choice variables, implying that the Kuhn-Tucker conditions are necessary and sufficient for a maximum.

The associated Lagrangian is

$$\mathscr{L} = \Big\{ \big\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a) - (1+\lambda)k \big\} f(k) - (\phi+\lambda)F(k) \Big\} X_r(k) + s\Big\{ \big\{ [\theta\beta h - (1+\lambda)c](t_e - t_a - t_\ell) - (1+\lambda)k \big\} f(k) - (\phi+\lambda)F(k) \Big\} X_n(k) + \gamma[1 - X_r(k) - X_n(k)], \tag{A52}$$

where γ is the Lagrange multiplier on constraint (A31). The Kuhn-Tucker conditions include inequalities bounding partials of the Lagrangian,

$$\frac{\partial \mathscr{L}}{\partial X_r(k)} = \left\{s[\theta\beta h-(1+\lambda)c](t_e-t_a)-(1+\lambda)k\right\}f(k)-(\phi+\lambda)F(k)-\gamma^d \leq 0 \tag{A53}$$

$$\frac{\partial \mathscr{L}}{\partial X_n(k)} = \left\{[\theta\beta h-(1+\lambda)c](t_e-t_a-t_\ell)-(1+\lambda)k\right\}f(k)-(\phi+\lambda)F(k)-\gamma^d \leq 0, \tag{A54}$$

complementary slackness conditions associated with nonnegativity constraints on the choice variables,

$$X_r^d(k)\frac{\partial \mathscr{L}}{\partial X_r(k)} = 0 \tag{A55}$$

$$X_n^d(k)\frac{\partial \mathscr{L}}{\partial X_n(k)} = 0, \tag{A56}$$

complementary slackness conditions associated with constraint (A31),

$$\gamma^d[1-X_r^d(k)-X_n^d(k)] \tag{A57}$$

as well as nonnegativity of Lagrange multipliers, nonnegativity of the choice variables, and constraint (A31) itself. We have added superscript d to variables $X_r^d(k)$, $X_n^d(k)$, and γ^d satisfying the Kuhn-Tucker conditions to distinguish the optimum under domestic input supply.

The proof proceeds by analyzing an series of exhaustive and mutually exclusive cases.

Case (i) Suppose $X_r^d(k) = X_n^d(k) = 0$. Then $\gamma^d = 0$ by (A57). Substituting $\gamma^d = 0$ into (A53)–(A54), dividing by $(1+\lambda)f(k)$, and rearranging yields

$$s\left(\frac{\theta\beta h}{1+\lambda}-c\right)(t_e-t_a) \leq k+\left(\frac{\phi+\lambda}{1+\lambda}\right)\frac{F(k)}{f(k)} \tag{A58}$$

$$\left(\frac{\theta\beta h}{1+\lambda}-c\right)(t_e-t_a-t_\ell) \leq k+\left(\frac{\phi+\lambda}{1+\lambda}\right)\frac{F(k)}{f(k)}. \tag{A59}$$

Let $\bar{k}^d$ be the infimum of the values of k satisfying (A59) with equality. The derivative of the right-hand side of (A59) with respect to k is

$$1+\left(\frac{\phi+\lambda}{1+\lambda}\right)\left[\frac{f(k)^2-f'(k)F(k)}{f(k)^2}\right]. \tag{A60}$$

The second term is positive if and only if $F(k)$ is logconcave. Since $F(k)$ is logconcave by assumption, derivative (A60) is positive, implying that (A54) is satisfied for all $k \geq \hat{k}^d$. By (2), (A58) is a weaker condition than (A59). Hence, (A58) is satisfied for all $k \geq \hat{k}^d$. Thus $k \geq \hat{k}^d$ is a necessary and sufficient condition for solution $X_r^d(k) = X_n^d(k) = 0$. By (19), the $\hat{k}^d$ satisfying (A59) with equality here is the is same $\hat{k}^d$ given in equation (37) in the text.

Case (ii) Suppose $X_r^d(k) > 0$ and $X_n^d(k) > 0$. By (A55)–(A56) both (A53) and (A54) must be satisfied with equality. Since both equal 0, the left-hand sides must be equal. Setting the left-hand sides of (A53) equal to the left-hand side of (A54) and rearranging yields

$$s(t_e-t_a) = t_e-t_a-t_\ell, \tag{A61}$$

violating assumption (2), ruling out this case.

Case (iii) Suppose $X_r^d(k) > 0$ and $X_n^d(k) = 0$. Since $X_r^d(k) > 0$, (A53) must be satisfied with equality by (A55). Thus,

$$\big\{s[\theta\beta h-(1+\lambda)c](t_e-t_a)-(1+\lambda)k\big\}f(k)-(\phi+\lambda)F(k) \tag{A62}$$

$$=\gamma^d \tag{A63}$$

$$\geq\big\{[\theta\beta h-(1+\lambda)c](t_e-t_a-t_\ell)-(1+\lambda)k\big\}f(k)-(\phi+\lambda)F(k) \tag{A64}$$

Rearranging (A62)–(A64) gives

$$k+\left(\frac{\phi+\lambda}{1+\lambda}\right)\frac{F(k)}{f(k)} \leq \frac{s}{1-s}\left(\frac{\theta\beta h}{1+\lambda}-c\right)t_\ell. \tag{A65}$$

Let $\hat{k}^d$ be the supremum of the values of k satisfying (A65). The positive sign on the derivative (A60) implies that the left-hand side of (A65) is increasing in k. Hence, (A65) is satisfied for all $k \leq \hat{k}^d$. By (20), the $\hat{k}^d$ satisfying (A65) with equality here is the is same $\hat{k}^i$ given in equation (36) in the text.

Case (iv) Suppose $X_r^d(k) = 0$ and $X_n^d(k) > 0$. By process of elimination, this solution is optimal for all $k \in (\hat{k}^d, \bar{k}^d)$.

Ignored Constraints Satisfied Combining cases (i)–(iv) together gives the optimal contract in the statement of the proposition. As a final step, we need to verify that this solution satisfies ignored constraints (A41)–(A42). Consider a firm with type $k' \in [0, \hat{k}^d]$ that invests at risk in equilibrium according to the optimal contract. The firm's expected profit equals $\Pi(k') = (1-s)\hat{k}^d + s\bar{k}^d - k'$ in equilibrium. The firm earns the same expected profit by deviating to any announcement $\tilde{k} \in [0, \hat{k}^d]$, because the contractual provisions are the same for all types in that interval. If the firm deviates to announcement $\tilde{k} \in (\hat{k}^d, \bar{k}^d]$, its expected profit equals $\pi(k', \tilde{k}) = s(\bar{k}^d - k') \leq \Pi(k')$ for all $k' \leq \hat{k}^d$. If the firm deviates to $\tilde{k} > \bar{k}^d$, its expected profit equals 0, which is strictly less than $\Pi(k')$ for $k' \leq \hat{k}^d$. Thus, the incentive-compatibility constraint (A41) is satisfied for all $k' \leq \hat{k}^d$. We also just verified that individual-rationality constraint (A42) is satisfied for them as well. The reader can similarly verify that constraints (A41) and (A42) are satisfied for other types as well. *Q.E.D.*

Appendix B. Details on Analysis of Endogenous Location

TBA

References

Agarwal, Ruchir and Gita Gopinath. 2021. "A Proposal to End the COVID-19 Pandemic," *IMF Staff Discussion Note* 2021/04, May.

Ahuja, Amrita, Susan Athey, Arthur Baker, Eric Budish, Juan C. Castillo, Rachel Glennerster, Scott D. Kominers, Michael Kremer, Jean Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon J. Tan, and Witold Wiecek. 2021a. "Preparing for a Pandemic: Accelerating Vaccine Availability," *American Economic Association Papers and Proceedings* 111: 331–335.

Antras, Pol and Robert W. Staiger. 2012. "Offshoring and the Role of Trade Agreements," *American Economic Review* 102: 3140–3183.

Athey, Susan, Juan C. Castillo, Esha Chaudhuri, Michael Kremer, Alexandre Simoes Gomes, and Christopher M. Snyder. 2022. "Expanding Capacity for Vaccines Against Covid-19 and Future Pandemics: A Review of Economic Issues," mimeo, January.

Bagnoli, Mark and Ted Bergstrom. 2005. "Log-concave Probability and its Applications," *Economic Theory* 26: 445–469.

Bagwell, Kyle and Robert W. Staiger. 2001. "Strategic Trade, Competitive Industries and Agricultural Trade Disputes," *Economics & Politics* 13: 113–128.

Bagwell, Kyle and Robert W. Staiger. 2002. *The Economics of the World Trading System*. MIT Press.

Bagwell, Kyle and Robert W. Staiger. 2006. "Will International Rules on Subsidies Disrupt the World Trading System?" *American Economic Review* 96: 877–895.

Bagwell, Kyle and Robert W. Staiger. 2012. "The Economics of Trade Agreements in the Linear Cournot Delocation Model." *Journal of International Economics* 88: 32–46.

Bollyky, Thomas J. and Chad P. Bown. 2021. "The Real Vaccine Procurement Problem: Why America Should Make Its Supply Chain More Transparent," *Foreign Affairs*, June 24.

Bown, Chad P. 2022a. "How COVID-19 Medical Supply Shortages Led to Extraordinary Trade and Industrial Policy," *Asian Economic Policy Review* 17: 114–135.

Bown, Chad P. 2022b. "COVID-19 Vaccine Supply Chains and the Defense Production Act," mimeo, May.

Bown, Chad P. 2022c. "The WTO and Vaccine Supply Chain Resilience During a Pandemic," mimeo, May.

Bown, Chad P. 2022d. "The US Produced Millions of COVID-19 Vaccine Doses Quickly. But Then its Capacity Petered Out," *PIIE Realtime Economic Issues Watch*, June.

Bown, Chad P. and Thomas J. Bollyky. 2022. "How COVID-19 Vaccine Supply Chains Emerged in the Midst of a Pandemic," *World Economy* 45: 468–522.

Bown, Chad P. and Christopher Rogers. 2021. "The US Did Not Ban Exports of Vaccine Supplies. But More Help Is Needed," *PIIE Trade and Investment Policy Watch*, June 7.

Castillo, Juan C., Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Rachel Glennerster, Scott D. Kominers, Michael Kremer, Greg Larson, Jean Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon J. Tan, and Witold Wiecek. 2021. "Market Design to Accelerate COVID-19 Vaccine Supply," *Science* 371: 1107–1109.

Cutler, David M. and Lawrence H. Summers. 2020. "The COVID-19 Pandemic and the $16 Trillion Virus," *Journal of the American Medical Association* 324: 1495–1496.

Grossman, Gene M., Elhanan Helpman and Hugo Lhuillier. 2021. "Supply Chain Resilience: Should Policy Promote Diversification or Reshoring?," mimeo, Princeton University, September 27.

Grout, Paul A. 1984. "Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach," *Econometrica* 52: 449–460.

Hanlon, Peter, Fergus Chadwick, Anoop Shah, Rachael Wood, Jon Minton, Gerry McCartney, Colin Fischbacher, Frances S. Mair, Dirk Musmeier, Jason Mattiopoulos, et al. 2021. "COVID-19—Exploring the Implications of Long-term Condition Type and Extent of Multimorbidity on Years of Life Lost: A Modelling Study," *Wellcome Open Research* 5: 75.

Interlandi, Jeneen. 2021. "The World Is at War With Covid: Covid Is Winning," *New York Times*, op-ed, September 21.

International Monetary Fund. 2020. *World Economic Outlook, April 2020: The Great Lockdown*. April. Washington, DC.

Jaynes, E. T. 1957. "Information Theory and Statistical Mechanics," *Physical Review* 106: 620–630.

Kremer, Michael, Jonathan Levin, and Christopher M. Snyder. 2020. " Advance Market Commitments: Insights from Theory and Experience," *American Economic Association Papers and Proceedings* 110: 269–273.

Kremer, Michael, Jonathan Levin, and Christopher M. Snyder. (2022) "Designing Advance Market Commitments for New Vaccines," *Management Science* forthcoming.

Laffont, Jean-Jacques and Jean Tirole. 1993. *A Theory of Incentives in Procurement and Regulation*. MIT Press.

Lo, Andrew W., Kien W. Siah, and Chi H. Wong. 2020. "Estimating Probabilities of Success of Vaccine and Other Anti-Infective Therapeutic Development Programs," *Harvard Data Science Review*, special issue 1, doi:10.1162/99608f92.e0c150e8.

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

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

Snyder, Christopher M., Wills Begor, and Ernst R. Berndt. 2011. "Economic Perspectives on the Advance Market Commitment for Pneumococcal Vaccines," *Health Affairs* 30: 1508–1517.

Snyder, Christopher M., Kendall Hoyt, Demitrios Gouglas, Thomas Johnston, and James Robinson. 2020a. "Designing Pull Funding for a COVID-19 Vaccine," *Health Affairs* 39: 1633–1642.

Ssentongo, Paddy, Anne E. Ssentongo, Navya Votlti, Destin Groff, Ashley Sun, Djibril M. Ba, Jonthan Nunez, Leslie J. Parent, Vernon M. Chinchilli, and Catharine I. Paules. 2022. "SARS-CoV-2 Vaccine Effectiveness Against Infection, Symptomatic and Severe COVID-19: A Systematic Review and Meta-analysis," *BMC Infectious Disease* 22: article 439.

Staiger, Robert W. 2022. *A World Trading System for the Twenty-First Century*. MIT Press, forthcoming.

Sun, Hongmei, Fuminori Toyasaki, and Ioanna Falagara Sigala. 2021. "Incentivizing At-risk Production Capacity Building for COVID-19 Vaccines," *Production and Operations Management* Special Issue Article: 1–17.

Sykes, Alan O. 2005. "Subsidies and Countervailing Measures," In Arthur E. Appleton, Patrick F. J. Macrory, and Michael G. Plummer, eds., *The World Trade Organization: Legal, Economic and Political Analysis* (vol. 2), Springer, pp. 1682-1706.

UK National Audit Office (NAO). 2020. "Investigation into Preparations for Potential COVID-19 Vaccines," Department for Business, Energy and Industrial Strategy, Department of Health & Social Care, NHS England and NHS Improvement, and Public Health England HC 1071, Session 2019–2021, December 16.

US Government Accountability Office (GAO). 2021. "COVID-19: Efforts to Increase Vaccine Availability and Perspectives on Initial Implementation," Washington, DC: GAO-21-443, April.

Walker Patrick G. T., Charles Whittaker, Oliver J. Watson, Marc Baguelin, Kylie E. C. Ainslie, Sangeeta Bhatia, Samir Bhatt, Adhiratha Bonyasiri, Olivia Boyd, Lorenzo Cattarino, et al. 2020. "The Global Impact of COVID-19 and Strategies for Mitigation and Suppression," Imperial College COVID-19 Response Team working paper, March 26. London.

Table 1: Calibration Results

Scenario	Probability of investment: At-risk capacity	Probability of investment: Any capacity	Expected net program benefit (bil. $)
First best	0.324	0.604	467
Second best	0.220	0.515	384
Third best domestic input	0.179	0.434	320
Third best foreign input			
• No export restriction ($\xi = 0$)	0.135	0.357	258
• Export restriction ($\xi = 0.5$)	0.011	0.295	125

Notes: Results for calibrated US program using parameter values described in text. Probability of at-risk capacity investment equals $F(\hat{k})$, where $\hat{k}$ is the threshold for at-risk investing in the relevant scenario. Probability of any capacity investment equals $F(\hat{k}) + [F(\bar{k}) - F(\hat{k})]s$, where $\bar{k}$ is the threshold for any investing in the relevant scenario.

Figure 1: Organization and Trade Structures Analyzed

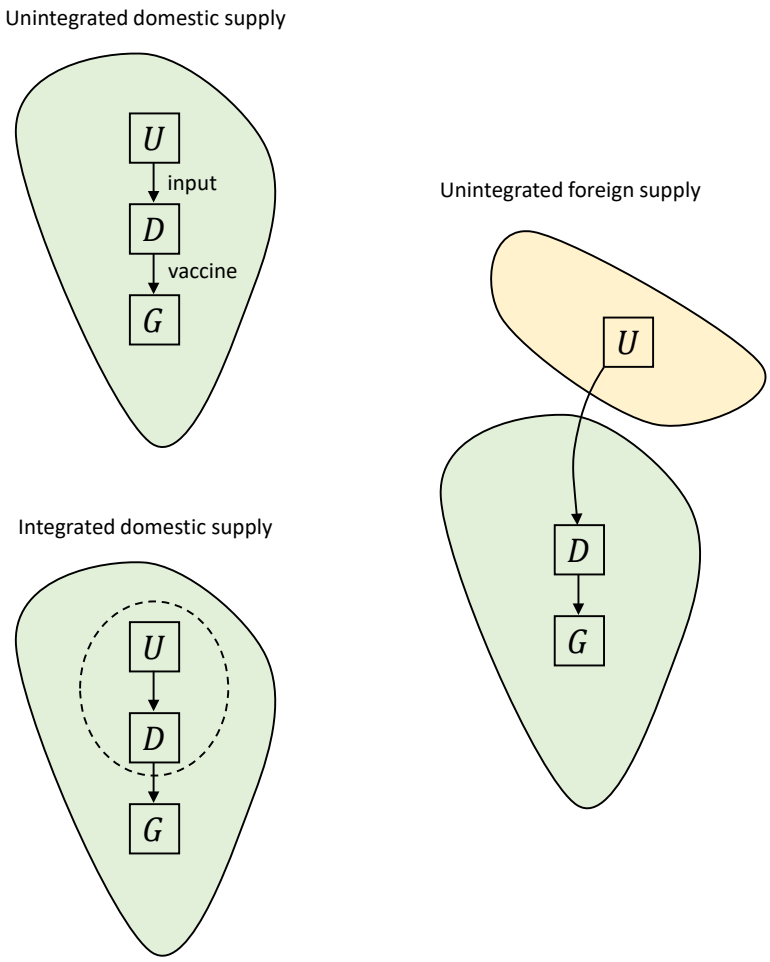


Figure 2: Timeline for Model

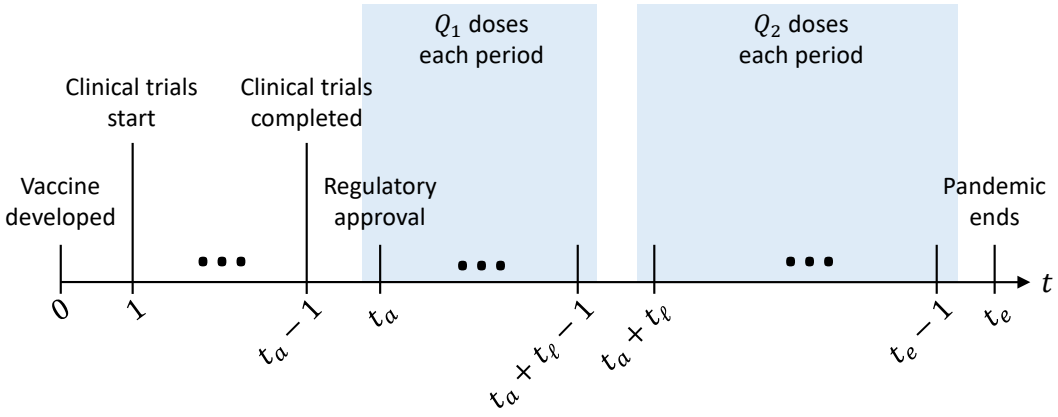


Figure 3: Optimal Capacity Investments as Functions of k

