

VACCINE SUPPLY CHAIN RESILIENCE AND INTERNATIONAL COOPERATION*

Chad P. Bown[†] Christopher M. Snyder[‡] Robert W. Staiger[§]

Preliminary
June 1 2022

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.¹ 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

*We thank Teresa Fort, Matthew Grant and seminar participants at Dartmouth for helpful feedback.

[†]Peterson Institute for International Economics; and CEPR.

[‡]Department of Economics, Dartmouth College; and NBER.

[§]Department of Economics, Dartmouth College; and NBER.

¹See Cutler and Summers (2020) and Agarwal and Gopinath (2021).

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 coordinated, expansionary subsidy policies along a supply chain - to ensure the resilience of that supply chain during a pandemic.

We begin by developing a simple theoretical model of a vaccine producer in a closed economy, motivated by the insights of Athey et al. (2022).² We characterize a pandemic as imposing a per-period cost on society 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 Snyder et al. (2020a,b), Ahuja et al. (2021), Castillo et al. (2021), 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 added speed with which it would allow a successful vaccine to be used to inoculate the population; the cost of at-risk investment is that this investment in capacity may turn out to be wasteful if the vaccine is not approved for use and the capacity cannot be repurposed.

In our benchmark model, a single vaccine dose protects an individual against the virus for only a single period, and the pandemic is assumed to end exogenously at some fixed future date, independent of vaccination status in the population; the implication is that the total demand for vaccine doses going forward declines through time as the end of the pandemic draws near, providing the firm with one possible reason for urgency when considering the timing of capacity investments. The question is whether the degree of urgency demonstrated by the firm matches that of the government/social planner. Our benchmark model yields an intuitively appealing answer to this question: since a profit-maximizing firm does not consider the social costs of the pandemic in its decision-making, it invests too little in at-risk production capacity relative to the social optimum if and only if the price at which it expects to be able to sell the vaccine is below the vaccine's social value. In this light, we show that the first-best policy to align private and social incentives for at-risk investment involves the government

²Athey et al (2022) provides the most complete summary of a series of research put forward over the course of the pandemic through the Accelerating Health Technologies (AHT) project.

offering the firm a price per vaccine dose equal to the social value of the dose. However, we also show that if efficient contracting directly over price is not an available policy option, then the government can replicate the first-best policy with an optimal subsidy for at-risk capacity investment. In the appendix, we consider a number of alternative modeling frameworks to confirm the robustness of this baseline result.

Estimates from Ahuja et al. (2021) put the social value for an individual course of COVID-19 vaccines at \$5,800, whereas the prices paid by governments were between \$4 and \$60 through 2021. Evidently, real-world governments found that a commitment to pricing vaccine doses anywhere near their social value was not feasible, suggesting according to our baseline result that subsidies to at-risk capacity investment could be an important tool. We therefore assume that the price of a vaccine dose is set (for exogenous reasons) below its social value, allowing us to focus on the role of subsidies to at-risk investment in aligning social and commercial incentives.³

We then turn to the core contribution of our analysis by considering the role of offshoring and imported inputs. To do so, we begin by adapting our model to a setting in which the domestic vaccine manufacturer can also potentially source inputs via imports from a foreign supplier. We draw from the framework introduced by Antràs and Staiger (2012) and model input trade as resulting in input prices which are determined by Nash bargaining between the domestic vaccine producer and its foreign input supplier after the foreign input supplier has sunk its capacity investments, on account of the lack of internationally enforceable complete contracts. We first show that the presence of offshoring induces a unilateral policy choice by the domestic government that results in less at-risk capacity investment by the firm than would result under the government’s unilateral policy choice if the input supplier were located in the domestic economy, due to a combination of the hold up problem faced by the foreign input supplier and the “leakage” of subsidy payments made by the domestic government to foreign entities. And all else equal, with less at-risk production capacity the expected health costs of the pandemic are then larger when the government makes unilateral policy choices and contracts with such an offshoring firm, relative to when it contracts with a vertically integrated firm or one that outsources domestically.

To consider the role of international cooperation over vaccine supply chain policy, we then introduce a second, symmetric country with a parallel market structure. We assume that each of the vaccine-producing firms can either source inputs locally or offshore inputs from the other country, and we focus on the at-risk capacity investment choices of input suppliers. We allow each government

³Of course in reality there also may be important frictions that our benchmark model misses and that would make a commitment to setting the price of a vaccine dose equal to its social value undesirable even if it were feasible, such as the frictions discussed in Athey et al., 2022, pp 6-8. In Appendix 9 we analyze one such friction that would arise with incomplete information in a linear-cost model, and there we allow governments to optimally choose the vaccine price in addition to the at-risk capacity subsidy. Our analysis confirms that while the details are altered, the main insights that derive from our benchmark model continue to apply in that setting.

to set its own vaccine policy, either unilaterally or possibly under a negotiated international agreement, maintaining our focus on the case where the price of a vaccine dose falls below the social value of the dose so that subsidies to at-risk capacity investments are desirable.

We show that our model's parameters can be partitioned into those under which the model exhibits offshoring of vaccine inputs as an equilibrium outcome when governments choose vaccine supply chain policies noncooperatively, and those under which the model exhibits local sourcing of inputs as an equilibrium outcome. And we show that the nature of the case for international cooperation over vaccine supply chain policy hinges critically on which of these noncooperative outcomes obtains.

If the noncooperative outcome features offshoring of vaccine inputs, the at-risk capacity installed by each input supplier in the noncooperative equilibrium will be inefficiently low, due to the international subsidy leakage that keeps unilateral subsidy choices for at-risk investment at inefficiently low levels. In this case an international agreement that implements free trade in vaccine inputs and commits each government to increase its per-unit subsidy to the at-risk capacity installed by its own input-producing firm can achieve the efficient level of at-risk capacity installed by each input supplier, and provides mutual gains for each country. For this case it is the positive externality generated by subsidies to at-risk capacity investment when offered to offshored input suppliers that drives the need for an international agreement covering vaccine supply chain policy. We argue that the absence of any such agreement, combined with the necessity of having to offshore certain inputs from abroad that was a binding constraint during the COVID-19 pandemic for a number of countries, may help to explain why some governments did not subsidize production capacity at risk with vaccine companies.

If instead the noncooperative outcome features local sourcing of vaccine inputs, then there may be no role for an international agreement on vaccine policy, but we show that it is also possible that a role for an international agreement on vaccine policy exists. Moreover, when there does exist such a role, we show that (i) the task of the agreement is to facilitate a switch to efficient offshoring of vaccine inputs, and (ii) while it is still the positive externality generated by subsidies to at-risk capacity investment when offered to offshored input suppliers that drives the need for the agreement, in this case the international externalities are off-equilibrium and the observed subsidies to at-risk capacity investment may either rise or fall under the international agreement.

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 we show that 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 setting 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 a potential carve out from existing WTO rules, which are generally focused on disciplining national use of subsidies, not cooperatively increasing them to globally efficient levels.⁴

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 policies. Grossman, Helpman and Liu (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 benchmark model of vaccine manufacturing in a closed economy that is experiencing a pandemic and examines optimal government policies for at-risk investment. Section 4 allows for the possibility of offshoring and assesses its impact on the equilibrium policy choices and outcomes relative to the social optimum. Section 5 introduces a second symmetric country, examines noncooperative policies, and describes key characteristics of the international agreement to deliver globally optimal cooperative policy in the presence of offshoring. Section 6 concludes. An appendix includes a number of model extensions that explore the robustness of our findings.

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), Novavax 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).⁵ In addition to

⁴That said, this feature of WTO subsidy rules poses something of a puzzle for the economic analysis of trade agreements. See, for example, Bagwell and Staiger (2001, 2002, 2006, 2012) and Sykes (2005).

⁵India'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

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 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; 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 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 “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 United States 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 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.⁶ 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” (with the European

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.

⁶The 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.

Union counting as one country).⁷ 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 United States and European Union. (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 plant manufacturing all of the drug substance for its vaccine at greater scale, the output of which could then be sent to multiple 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.

2.2 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 United States 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. While the public versions of the contracts with these companies contain considerable redactions, according to the Government Accountability Office, 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

⁷One 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.

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 United States on December 18, 2020.)

The US government had a completely different, procurement-only contract with Pfizer. The United States 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.⁸

Finally, the United States 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 capacity-enhancing contracts to other firms providing variable inputs essential to vaccine production (e.g., bioreactors, bioreactor bags, cellular material) as well as vaccine dose delivery (e.g., vials, needles, syringes).⁹

No other government matched the scale or scope of the US at-risk contracts to vaccine sponsors or to input providers. In Europe, the United Kingdom 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” (UK 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 European Union (European Investment Bank) provided at-risk financing to BioNTech and CureVac to allow them to expand production capacity. CEPI provided some at-risk funding to the Serum Institute of India to manufacture the AstraZeneca vaccine for allocation to COVAX (COVID-19 Vaccines Global Access), the facility aiming to procure and distribute COVID-19 vaccines to low-income countries globally.¹⁰

⁸According to GAO (2021, p. 18) “To minimize financial risk to the government, the parties agreed that the government would pay Pfizer only after its vaccine received authorization”.

⁹See Bown and Bollyky (2022, Table 4).

¹⁰See Bown and Bollyky (2022, Table 6 and Table 7). CEPI, the Coalition for Epidemic

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 European Union 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.¹¹ 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, 8-10 months later than many of the contracts were agreed in the United States, for example.

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.¹²

3 Vaccine Production with Local Inputs

We consider the problem of a government/benevolent social planner who seeks to minimize the overall expected cost to its country of a viral outbreak, and who deals with a single, domestic, profit-maximizing (risk neutral) firm that is in possession of a promising vaccine technology. We assume in this section that, if clinical trials show that the vaccine technology is “safe and effective” and the vaccine is approved for use, the firm can produce the vaccine with locally sourced inputs: the vaccine-producing firm either sources inputs completely within its firm boundaries (i.e., it is a vertically integrated firm), or it can outsource domestically all needed inputs for vaccine production in an environment of perfect contract enforcement. Since these two possibilities for firm structure will lead to identical outcomes for the results we emphasize here, in the formal analysis of this section we will treat the vaccine-producing firm as vertically integrated. We abstract from the impacts that a government’s policy toward vaccine production might have on the distribution of income across its citizens, and implicitly assume that the government has lump-sum instruments that it can use to redistribute to its citizens any undesired rents that its vaccine

Preparedness Innovations, is a foundation that received donations from public, private, philanthropic, and civil society organizations to finance independent research projects to develop vaccines against emerging infectious diseases.

¹¹This may have been motivated in part by EU member states imposing export controls on PPE trade with each other in March 2020 (Bown, 2022a).

¹²For US regulators and the vaccine candidates that received at-risk funding under Operation Warp Speed, this included AstraZeneca, Novavax and Sanofi-GSK.

policies create for its vaccine producers. This will allow us to make our main points with a minimum of formal complexity.¹³

We focus on the divergence between private and social incentives to undertake “at-risk” investment in capacity for vaccine production, as emphasized by Snyder et al. (2020a,b), Ahuja et al. (2021), 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 sequential investment in capacity only once the vaccine has been approved for use, which we refer to as “not-at-risk” investment – is the added speed with which at-risk investment would allow a successful vaccine to be used to inoculate the population; the cost of at-risk investment is that this investment in capacity may turn out to be wasteful if the vaccine is not approved for use and the capacity cannot be repurposed.

In this section we seek conditions under which the (vertically integrated) firm has insufficient incentive to make at-risk investments relative to the social optimum, and therefore conditions under which a government program of subsidies to at-risk investment would be optimal. Our perspective is that of a closed economy. In the next section we allow the possibility of international trade, and consider how such a program of subsidies would be complicated if the vaccine-producing firm offshores its inputs from foreign sources. In this and the next section we assume that the government designs its vaccine policy unilaterally. We consider the possible role that an international agreement on vaccine policy could play in section 5. Throughout we abstract from the R&D phase of vaccine development and focus solely on the question of optimal at-risk investment.

We undertake our analysis in several related but distinct models, featuring in this section and throughout the main body of the paper a single, simple modeling approach that we will refer to as the Benchmark Model, and relegating to the appendix our analysis under alternative modeling approaches. In all of our models the outbreak is assumed to exogenously end after a fixed length of time. The end of the outbreak could reflect the emergence of a highly contagious but benign variant of the virus which dominates other variants, precluding serious harm from the disease after that point. Alternatively, the end of the outbreak 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,

¹³Many interesting and important issues of vaccine policy design arise as a result of the desire of policy makers to achieve public health goals with a minimum of rents transferred to supplying firms (see, for example, Snyder et al., 2020a,b). By abstracting from these issues in the main body of the paper we are able to simplify the analysis and better highlight our main points, which we view as complementary to this wider set of design issues. In Appendix 9 we consider the specific possibility that incomplete information would limit the ability of the government to eliminate rents associated with its policy intervention at will.

setting the discount rate for all players to 0.

Critically, in each model that we consider the social value of installing an additional unit of vaccine production capacity declines through time. Thus, each model features a hallmark of viral outbreak situations: for purposes of public health and when it comes to the installation of vaccine production capacity during the outbreak, speed is of the essence. This feature sets the stage for an analysis of the commercial versus social incentives for at-risk (i.e., early) investments in capacity installation. But across the models that we consider, speed matters for different reasons.

In the Benchmark Model that we feature throughout the main body of the paper, the social value of an additional vaccine *dose* is constant through time as long as the outbreak continues, because the health benefit provided by a single vaccine dose is non-durable. But the social value of installing an additional unit of vaccine production capacity nevertheless declines through time, because production capacity is durable and the *number* of doses that society can make use of – and that installed, durable production capacity can make possible – declines through time as the end of the outbreak draws near.

In the model that we consider in Appendix 7, by contrast, the social value of installing an additional unit of vaccine production capacity declines through time, because the health benefit provided by a vaccine dose is now durable, and so while the number of doses that society can make use of before the outbreak ends is fixed by the size of its population, it is now the social value of an additional vaccine dose itself that declines through time, a reflection of the fact that the time-span that the dose serves to prevent illness becomes truncated as the end of the outbreak approaches. In Appendix 8 we present an extension of these two modeling approaches in which the social value of installing another unit of vaccine capacity declines through time for an additional reason, namely, because the susceptibility in the population is now assumed to be heterogeneous, and the first vaccine doses are given to those who are most in danger of severe illness from the virus. And finally, in Appendix 9 we introduce the possibility that the government operates in a setting of incomplete information when designing its vaccine policy interventions. Our goal in presenting the analysis in these alternative settings is to extract the general policy insights that apply across the disparate model environments.

3.1 Potential vaccine demand that declines through time

We consider a country that is populated by a continuum of citizens with mass L . At time $t = 1$ a new virus begins circulating in the country. We assume that, absent an effective vaccine, the virus circulates in the population for T periods and then disappears; hence period T marks the last period of the outbreak. We assume as well that in every period when the virus is circulating each unvaccinated individual will be infected by the virus with probability $\delta \in (0, 1]$ and, if infected, will be sick for the period and then recover (infected individuals do not die, and hence the population is constant through time); we therefore abstract from the dynamics of contagion and assume that an unvaccinated individual's

chances of infection are independent of how many others in the population are infected or vaccinated, and that an unvaccinated individual who has been infected in a previous period has the same chance of being infected again in a future period as long as the individual remains unvaccinated and the virus is still circulating. An individual who is infected by the virus becomes ill and imposes a one-period health cost on society which we denote by $h > 0$. This cost reflects both the private cost (economic, physical, psychological) borne by the individual and also the wider social costs associated with becoming ill (overcrowding in hospitals, disrupted supply chains). Finally, we abstract from discounting (by either the firm or society) over the T periods of the outbreak. With these assumptions, in the absence of a safe and effective vaccine the expected cost imposed on the country by the virus amounts to $LT\delta h$.

We assume that clinical trials for the vaccine begin at time $t = 1$, the period that the new virus begins circulating in the population, and that these trials take a total of M periods to complete.¹⁴ Hence, at the beginning of period $M + 1$ it becomes known whether the vaccine is safe and effective, and we assume that the vaccine will be deemed safe and effective with probability $s \in (0, 1]$. If the clinical trials fail to show that the vaccine is safe and effective, then the vaccine will not be approved for use and no sales of the vaccine can be made in any period. On the other hand, if the vaccine is found to be safe and effective, then it is immediately approved for general use in the population beginning in period $M + 1$ and until the virus stops circulating at the end of period T . We assume that one dose of the vaccine is 100% effective in preventing an individual from contracting the virus during the period that the dose is administered, and provides no protection thereafter. Hence, for a vaccine that is found to be safe and effective at the beginning of period $M + 1$, the number of vaccine doses that an individual requires to remain protected from the virus through to the end of the outbreak is $T - M$; therefore to inoculate the entire population beginning in period $M + 1$ and until the virus stops circulating at the end of period T would require $L(T - M)$ doses.

The upshot is that in period $t = M + 1$ the potential total demand for safe and effective vaccine doses through the end of the contagion period is equal to $L(T - M)$, but this potential demand declines with each passing period: in period t for $t \in [M + 1, T]$ it is given by $L(T - (t - 1))$. As we will see, for this reason the potential total demand faced by the vaccine-producing firm depends on the capacity investment choices that the firm makes – and in particular on whether at-risk capacity investment is undertaken and, if so, at what level.

We model vaccine production as a two-step process. An input (say the vac-

¹⁴In the context of the COVID-19 vaccines, phase-3 clinical trials took roughly 5 months to complete. We assume that the trials themselves are sufficiently small relative to the population that they are costless to run (reflecting the reality that clinical-trial costs are eclipsed by capacity and production costs for vaccines supplied at pandemic scales; see the cost estimates in Snyder et al., 2020), and that they have no measurable impact on population-wide health outcomes. Also, we take the completion time for clinical trials to be exogenous, but this may also be an important margin for policy intervention (indeed, the US and UK provided some funding for clinical trials, as described in section 2), and possibly for international policy cooperation.

accine’s active ingredient) produced in an upstream facility is shipped to a downstream facility that fills a vial with the input and other additives constituting a finished dose. As mentioned, in this section we assume that the two facilities are located domestically and (without loss of generality) integrated in the same firm. In any period t , input production cannot exceed the input supplier’s installed capacity. We assume that after it is installed, one unit of capacity can produce the input requirement for one vaccine dose each period thereafter, and that capacity is sunk and cannot be repurposed to manufacture other pharmaceuticals or products. Besides the capacity cost, each unit of input involves a marginal production cost of κ , with $\kappa < \delta h$ so that the social benefit of another dose of the (safe and effective) vaccine for use during the outbreak exceeds the marginal cost of producing it. 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.¹⁵

In principle, the firm can invest in any amount of input capacity $v_t \geq 0$ at the beginning of any period t in as many tranches as it wishes, and as we have noted, we assume that, once installed, this capacity is durable. But installation takes time. In particular we assume that capacity investments made at the beginning of period t will not be fully installed and ready for deployment in production until the beginning of period $t + m$. We further assume that the cost of investing in capacity at the level v , which is incurred in the period when the investment in that capacity is made, is given by the increasing and convex cost function $c(v)$, with $c(0) = 0$ and $c'(0) = 0$. The convexity of the capacity cost function provides a technological reason for spreading the investments in capacity across periods, and in particular across the at-risk ($t \in [1, M]$) and not-at-risk ($t \in [M + 1, T]$) phases of the outbreak; and if speed means concentrating these investments into a short, early time span, then the convex cost assumption ensures that speed cannot be had for free. But the question still remains whether the profit motives of the firm will lead to the socially optimal choices in this regard. Finally, for simplicity we assume that the shelf life of the input expires at the end of the period in which it is produced, and we thereby abstract from the possibility of using at-risk capacity to build a stockpile of the input for later use should the clinical trials be successful.¹⁶ To avoid uninteresting taxonomies, we also assume that $T > M + m$ and that $m < M$.¹⁷

In order to focus on the at-risk capacity margin, we introduce the following assumption that will simplify the analysis of the government’s preferred at-risk

¹⁵Of course in reality there are bound to be important at-risk capacity investment choices that have to be made for both inputs and final production of vaccines, and a more realistic model would consider the problem of at-risk capacity at each stage along the production process and the complications for policy that this might introduce. We leave this extension to future work.

¹⁶Stockpiling vaccine doses at risk is a theoretically interesting possibility, but it does not appear to have been particularly relevant in the context of COVID-19 vaccines. In any case, allowing for this possibility is beyond the scope of the current paper.

¹⁷If alternatively we were to assume that $m > M$, then it is easy to show that investments in at-risk capacity would speed up the availability of vaccine doses by M periods (rather than m periods) should the vaccine be deemed safe and effective.

investment level:

$$\max\left[\frac{c(L)}{(T - (M + m))L}, \frac{c'(L)}{s}\right] < (\delta h - \kappa). \quad (\text{Assumption I})$$

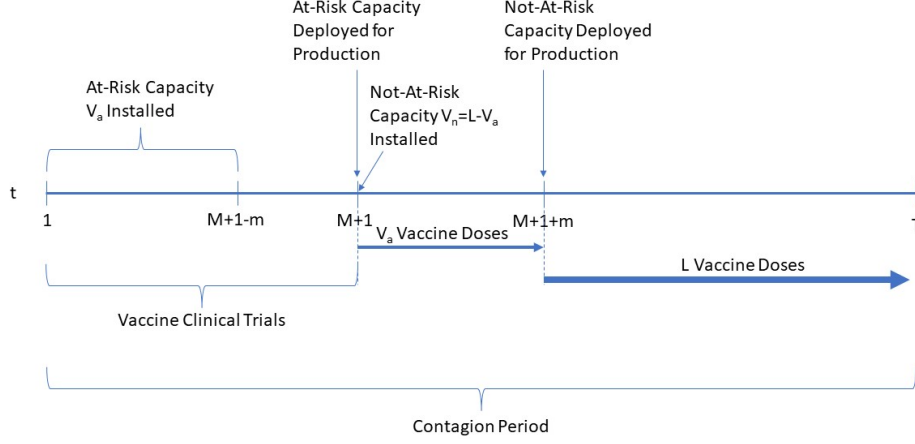
The restriction that $\frac{c(L)}{(T - (M + m))L} < (\delta h - \kappa)$ indicated by Assumption I is easy to interpret. This restriction ensures that even if the capacity necessary to inoculate the entire population (L) had to be installed in the single period $t = M + 1$, the social benefits of installing this capacity (assuming the vaccine has been found to be safe and effective) in terms of the inoculations it would make possible ($(T - (M + m))L \times (\delta h - \kappa)$) would still exceed the cost ($c(L)$).

The restriction that $c'(L) < s(\delta h - \kappa)$ implied by Assumption I is more subtle, but can be understood as follows. The term $c'(L)$ can be interpreted as the maximum capacity cost savings (defined positively) that could ever be had by delaying a small increment of capacity investment by one period and thereby achieving some investment smoothing. This maximum would occur if the initial level of capacity investment in some period were L , and if there were initially no capacity investments planned for the following period. From that (maximum-bunching) starting point, the removal of a small increment of capacity investment from the period where it was initially planned would save an amount equal to $c'(L)$ in capacity costs for that period, while adding the small increment of capacity investment to the following period where there was initially zero investment would add the amount $c'(0) = 0$ in capacity costs. That this is the maximum capacity cost savings that could be had by such investment smoothing follows from the convexity of the capacity cost function $c(\cdot)$. The term $s(\delta h - \kappa)$ is the (minimum expected) social cost of the one-period delay in vaccine delivery that would be associated with this incremental reallocation of capacity investments across the two periods. Assumption I states that the social cost of the delay in vaccine delivery created by this reallocation exceeds the capacity-cost savings from the achieved investment smoothing, and so ensures that such a reallocation would never be socially desirable.

Assumption I affords two simplifications to our characterization of the government's preferred timing of capacity investments. First, for any level of at-risk investment below L , Assumption I implies that if at the beginning of period $M + 1$ it is announced that the vaccine has passed the clinical trials and is deemed to be safe and effective, then the government will choose to "top off" the at-risk investment with not-at-risk investment in period $M + 1$ to bring operational capacity all the way up to the level L as quickly as possible (i.e., beginning in period $M + 1 + m$), in order to be able to vaccinate every citizen in every remaining period $t \in [M + 1 + m, T]$. Denoting by $V_a \equiv \sum_{\tau=1}^M v_\tau$ the stock of at-risk capacity accumulated by the end of period M , it follows from Assumption I that the government's investment choices for the not-at-risk periods $t \in [M + 1, T]$ will correspond to $v_{M+1} = L - V_a$ and $v_t = 0$ for $t \in [M + 2, T]$; and not-at-risk capacity, which we denote by $V_n \equiv \sum_{\tau=M+1}^T v_\tau$, is then

$$V_n = v_{M+1} = L - V_a. \quad (1)$$

Figure 1: Effective Time Line



And second, Assumption I implies that under the government's choices there will be no at-risk capacity investments made in periods $t \in [M + 2 - m, M]$, because relative to making these investments instead in at-risk periods $t \in [1, M + 1 - m]$ such investments would be trading off investment smoothing for delayed health benefits which we know under Assumption I is never desirable. Moreover, it is easy to see that it will be optimal to spread any at-risk investments made in periods $t \in [1, M + 1 - m]$ evenly across those periods so that $v_t = v_{t'} \equiv v_a$ for $\{t, t'\} \in [1, M + 1 - m]$. Together these last two features imply that

$$V_a = \sum_{\tau=1}^{M+1-m} v_{\tau} = (M + 1 - m) \times v_a. \quad (2)$$

We record the implications of Assumption I in

Lemma I *Under Assumption I, the socially optimal level of at-risk investment involves investing $v_a \geq 0$ in periods $t \in [1, M + 1 - m]$ leading to the stock of at-risk capacity $V_a = (M + 1 - m) \times v_a$ at the end of period $M + 1 - m$, and no investment for $t \in [M + 2 - m, M]$; the socially optimal level of not-at-risk investment involves the investment $v_{M+1} = L - V_a$ in period $M + 1$, and no investment for $t \in [M + 2, T]$.*

According to Lemma I, we can cast the government's optimal choice of capacity investment as its choice of the constant flow level of at-risk investment v_a that will be maintained over the periods $t \in [1, M + 1 - m]$, or equivalently as simply its choice of the stock of at-risk capacity V_a at the end of period $M + 1 - m$. Schematically, the effective time line of the government's choices and their implications for delivery of vaccine doses is illustrated in Figure 1.

We can now write down the objective of the government (who is taken to be a benevolent social planner). The government chooses v_t for $t \in [1, T]$ to

minimize the expected overall cost to its country of the viral outbreak. However, as noted, by Lemma I we can think of the government as simply choosing the stock of at-risk capacity V_a that it wants to accumulate by the end of period $M + 1 - m$. To express the government's choice in this way, we first define $C(V_a) \equiv (M + 1 - m) \times c(v_a(V_a))$, where $v_a(V_a) \equiv \frac{V_a}{(M+1-m)}$. The variable $C(V_a)$ represents the capacity costs associated with V_a , the stock of at-risk capacity accumulated by the end of period $M + 1 - m$, where it is acknowledged that this stock of capacity is achieved through a flow of capacity investments at a constant level over the periods $t \in [1, M + 1 - m]$. The expected cost to the country of the viral outbreak may then be written as

$$E[Cost(V_a)] = C(V_a) + sc(L - V_a) + s[mV_a + (T - (M + m))L]\kappa \\ + (1 - s)LT\delta h + s[LM\delta h + (L - V_a)m\delta h]. \quad (3)$$

The right-hand side of (3) is easy to interpret. The first two terms in the first line are respectively the cost of investing in at-risk capacity at the level V_a , before it is known whether the vaccine will be approved for use, and the expected cost of the additional capacity investment V_n in period $t = M + 1$ in the (probability s -) event that the vaccine is found to be safe and effective, with $V_n = L - V_a$ ensured under Assumption I as noted in Lemma I. The third term in the first line reflects the expected marginal cost of producing the (safe and effective) vaccine up to capacity given the at-risk capacity level V_a . Together these terms give the expected economic cost of the vaccine program itself (recall that we are abstracting from the R&D phase of vaccine development).

The expected benefits of the vaccine program accrue from the possibility that the program will reduce the health costs that the country incurs as a result of the virus from what these costs would have been in the absence of the program, which as noted above would amount to $LT\delta h$. These expected benefits are reflected in the two terms in the second line of (3). If the vaccine is found not to be safe and effective, which occurs with probability $(1 - s)$, then the vaccine program has no impact on the virus and the country will suffer the health costs that arise when all individuals are unvaccinated over all T periods, namely, the cost $LT\delta h$. On the other hand, if the vaccine is found to be safe and effective, which occurs with probability s , then for the first M periods prior to vaccine approval the country will still suffer the health costs that arise when all individuals are unvaccinated, namely, the cost $LM\delta h$, but the health costs of the virus during the remaining $T - M$ periods will be different: a portion $(L - V_a)$ of the population will remain unvaccinated for another m periods after the clinical trials are over while they wait for the additional capacity $(L - V_a)$ to be installed, and during this time the country will suffer the additional health costs $(L - V_a)m\delta h$; beyond this point, everyone is vaccinated in each period, and there is no further health cost as a result of the virus.

Using (3), the first-order condition that determines the socially optimal choice of at-risk capacity investment – that is, the at-risk capacity investment level \hat{V}_a that minimizes the expected overall cost to the country of the viral

outbreak – is given by

$$\frac{\partial E[Cost(V_a)]}{\partial V_a} = C'(\hat{V}_a) - s[c'(L - \hat{V}_a) + m(\delta h - \kappa)] = 0, \quad (4)$$

with the second-order condition guaranteed by the convexity of the cost-of-capacity $c(v)$ function. According to (4), the socially optimal level of at-risk capacity investment \hat{V}_a is determined where the marginal cost of the last unit of at-risk investment, $C'(\hat{V}_a)$, is equal to the expected marginal benefit, where the marginal benefit of the last unit of at-risk capacity is non-zero only when the vaccine is shown to be safe and effective (with probability s) and is then the sum of two terms, the cost savings $c'(L - \hat{V}_a)$ associated with the smaller level of V_n needed to top off the shortfall between \hat{V}_a and L , and the net savings on health costs $m(\delta h - \kappa)$ that the last unit of at-risk capacity generates by speeding up the delivery of a vaccine dose by m periods (and hence facilitating the provision of m additional doses, each with a social value of δh and a marginal production cost κ).

Notice that according to (4) and with $C'(0) = 0$, \hat{V}_a approaches zero as s approaches zero, because as the possibility of vaccine success becomes vanishingly small, investing in at-risk capacity is purely wasteful. On the other hand, as s approaches one and vaccine success becomes virtually guaranteed, the benefits of spreading vaccine capacity investments over time to minimize capacity costs (which follows from the convexity of $c(v)$) becomes important; indeed, in this case and as (4) confirms, but for the net savings in health costs $m(\delta h - \kappa)$ associated with at-risk capacity, it would be optimal to spread capacity investments evenly across at-risk and not-at-risk periods, and the savings in health costs associated with at-risk capacity then pushes in the direction of greater at-risk capacity investments. These are the tradeoffs that determine the socially optimal level of at-risk capacity \hat{V}_a . Finally, for future reference we record the level of expected welfare achieved by the government under the optimal at-risk investment:

$$\begin{aligned} E[Cost(\hat{V}_a)] &= C(\hat{V}_a) + sc(L - \hat{V}_a) + s[m\hat{V}_a + (T - (M + m))L]\kappa \\ &\quad + (1 - s)LT\delta h + s[LM\delta h + (L - \hat{V}_a)m\delta h]. \end{aligned} \quad (5)$$

We next consider the capacity investment choices of the vaccine-producing firm. To this end, we take as exogenous the price P that the firm receives for each dose of a vaccine that is deemed to be safe and effective, and in analogy with Assumption I we assume

$$\max\left[\frac{c(L)}{(T - (M + m))L}, \frac{c'(L)}{s}\right] < (P - \kappa). \quad (\text{Assumption II})$$

Assumption II has exactly the same interpretation as Assumption I, but with the expression on the right-hand side of the inequality in Assumption II now reflecting the operating profits rather than the social benefits generated by one more vaccine dose. We record the implications of Assumption II in

Lemma 2 *Under Assumption II, the profit-maximizing level of at-risk investment involves the investment $v_a \geq 0$ in periods $t \in [1, M+1-m]$ leading to the stock of at-risk capacity $V_a = (M+1-m) \times v_a$ at the end of period $M+1-m$, and no investment for $t \in [M+2-m, M]$; the profit-maximizing level of not-at-risk investment involves the investment $v_{M+1} = L - V_a$ in period $M+1$, and no investment for $t \in [M+2, T]$.*

According to Lemma 2, under Assumption II we can cast the firm's optimal choice of capacity investment as its choice of the stock of at-risk capacity V_a at the end of period $M+1-m$, just as we did for the government problem analyzed above, and focus on any differences across the firm and the government in the choice of V_a , assured that the firm and government incentives for not-at-risk investment V_n are aligned. The firm's expected profits can therefore be written as

$$E[\pi(V_a)] = s[V_a m + (T - (M + m))L] \times (P - \kappa) - [C(V_a) + sc(L - V_a)]. \quad (6)$$

The first line of (6) gives the expected operating profits for the firm when it invests V_a in at-risk capacity. With probability $(1 - s)$ the vaccine will fail the clinical trials and the firm will earn zero revenue from vaccine sales. But with probability s the vaccine will be found safe and effective and will be approved for sale at the beginning of period $M+1$, and in that case the firm will earn vaccine sales operating profits in the amount of $V_a m \times (P - \kappa)$ off of its at-risk capacity, and in period $M+1$ it will invest in the additional capacity $V_n = (L - V_a)$ and then, once this additional capacity comes on line, earn vaccine sales operating profits in the amount of $(T - (M + m))L \times (P - \kappa)$. The second line of (6) gives the expected costs of the capacity investment levels V_a and $V_n = (L - V_a)$.

Using (6), the first-order condition that determines the profit-maximizing choice of at-risk capacity investment for the firm, \tilde{V}_1 , is given by

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = sm(P - \kappa) - [C'(\tilde{V}_a) - sc'(L - \tilde{V}_a)] = 0, \quad (7)$$

with the second-order condition again guaranteed by the convexity of the cost-of-capacity function $c(v)$. According to (7), the profit-maximizing level of at-risk capacity investment \tilde{V}_a is determined where the expected marginal operating profits generated from the vaccine sales made possible by the last unit of at-risk investment, $sm(P - \kappa)$, is equal to the expected marginal cost of the last unit of at-risk capacity investment, $[C'(\tilde{V}_a) - sc'(L - \tilde{V}_a)]$, which in addition to $C'(\tilde{V}_a)$ includes as well the possibility (with probability s) of reduced cost in topping off capacity in period $M+1$ that the last unit of at-risk capacity affords ($-c'(L - \tilde{V}_a)$). Finally, for purposes of comparison with (4), it is helpful to rewrite (7) in the equivalent form

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = -\{C'(\tilde{V}_a) - s[c'(L - \tilde{V}_a) + m(P - \kappa)]\} = 0. \quad (8)$$

Armed with (4) and (8), we are now in a position to determine the conditions under which the at-risk capacity investment chosen by the firm will be lower than that chosen by the government/social planner. As a comparison of (4) and (8) confirms, if the price P at which each vaccine dose can be sold is equal to δh , the expected per-period health cost associated with an unvaccinated individual, then $\tilde{V}_a = \hat{V}_a$ and the choice of the firm will coincide with the choice of the government. On the other hand, if P is less than δh , then (4) and (8) imply $\tilde{V}_a < \hat{V}_a$. We may now state:

Proposition 1 *Absent additional policy intervention and relative to the social optimum, in the Benchmark Model at-risk capacity investment will be under-supplied by a vaccine-producing firm that locally sources inputs and can sell vaccines at a price P per dose if and only if $P < \delta h$.*

Proposition 1 points to a simple solution to the problem of under-investment in at-risk capacity: Why not maintain a price per dose of safe and effective vaccines that is equal to the social value of that dose, δh ? That is, by setting

$$\hat{P} \equiv \delta h, \quad (9)$$

the government would ensure that the firm's incentives for investing in at-risk capacity would align with the social incentive, delivering $\tilde{V}_a = \hat{V}_a$. An alternative and equivalent intervention suggested by Proposition 1 when $P < \delta s$ would be to offer a subsidy σ per unit of at-risk capacity investment defined so as to insure that¹⁸

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = \sigma + sm(P - \kappa) - [C'(\hat{V}_a) - sc'(L - \hat{V}_a)] = 0. \quad (10)$$

As a comparison of (4) and (10) confirms, when $P < \delta h$ this implies a subsidy in the amount

$$\hat{\sigma}(P) \equiv sm \times (\delta h - P). \quad (11)$$

That is, by adding the term $sm \times (\delta h - P)$ to the right-hand side of (8), such a subsidy to the investment in at-risk capacity would bring social and commercial incentives in line.

It is interesting to compare these policy prescriptions to the discussion in Athey et al. (2022):

Given the large expenditures involved in the recommended policy, measures should be taken to ensure the expenditures are spent efficiently, not wasted. Section 6 discusses some of these design features, including combining some of both “push” funding (contributing toward capacity costs) and “pull” funding (paying a bonus price for approved vaccines produced on an expedited schedule) and contracting on capacity rather than doses. (Athey et al., 2022 p 2)

¹⁸It can be confirmed that introducing such a subsidy cannot cause Assumption II to be violated and so we may continue to focus on the expression for $E[\pi(V_a)]$ in (6) when calculating the first-order condition that defines \tilde{V}_a in the presence of the subsidy.

According to Athey et al., the natural policy options for addressing insufficient private incentives to invest in at-risk capacity are to either contribute toward capacity costs or pay a bonus price for approved vaccines produced on an expedited schedule (see also their discussion of commercial versus social incentives on pp 6-8, and their discussion of push versus pull funding on pp 14-15). In the model above, the former would correspond to offering the subsidy $\hat{\sigma}$ defined in (11) to investment in at-risk capacity, while the latter would correspond to offering a higher price for vaccines delivered in period $M + 1$ than for vaccines delivered in period $M + m + 1$. But Proposition 1 suggests that if the vaccine price is fixed at the social value of a dose \hat{P} defined in (9), neither of these policy alternatives would be needed.

How might our analysis make more contact with the discussion in Athey et al. (2022) and the related literature? Recall that we have assumed that the outbreak will end in T periods, regardless of vaccine status, and that to prevent illness effective vaccines must be renewed with additional doses every period. These two assumptions are important for the specific result reported in Proposition 1, because as we have observed they imply that the social value of a vaccine dose is constant through time but the potential demand for vaccine doses declines with each passing period, and these assumptions therefore ensure that the faster the firm can deliver vaccines, the more doses it can sell. It is this feature of our model that, when combined with a fixed price per vaccine dose set at the social value of a dose, aligns commercial with social incentives to invest in at-risk capacity as Proposition 1 reports.

In Appendix 7 we develop an alternative model that reflects instead the feature that the number of vaccine doses needed to end the contagion is fixed, so that the number of vaccine doses the firm can sell is independent of its speed of vaccine delivery. As we establish and as might be expected, this alternative model does generate the need for a declining price of vaccine doses over time if the firm is going to deliver the socially optimal level of at-risk capacity investment, much as Athey et al. (2022) indicate. However, it turns out that in this alternative model the social value of a vaccine dose is now declining over time, and as a result the more general insight from Proposition 1 – that the firm will make the socially optimal investments in at-risk capacity if it faces a price per vaccine dose that reflects the social value of the dose – is preserved in this alternative setting as well. In the appendix we also present variants of these two models in which the social value of investing in an additional unit of vaccine capacity declines through time because the susceptibility in the population is now assumed to be heterogeneous, and the first vaccine doses are given to those who are most in danger of severe illness from the virus. We show that the insights from Proposition 1 apply equally well in those settings.

Together, these results suggest that, if a government can commit to a price schedule for vaccine doses, then committing to set the price per dose at the social value of each dose would be a natural policy intervention for ensuring that socially optimal at-risk capacity investment is made by a vaccine-producing firm that locally sources its inputs. However, if commitment to such a price schedule is not feasible, then a subsidy to investment in at-risk capacity is an

appropriate alternative that may be able to achieve the same result. We will therefore henceforth assume for our Benchmark Model that commitment to pricing doses at their social value is not feasible (or, for unmodeled reasons, not desirable) for the government, so that the implied policy from Proposition 1 is taken to be a subsidy to at-risk capacity investment. In the next section, we consider how such a policy would work (or would not work) if the vaccine producing firm is offshoring some of its inputs abroad. Of course in reality there also may be important frictions that our model has missed and that would make commitment to such a price schedule undesirable even if it were feasible, such as the frictions discussed in Athey et al., 2022, pp 6-8. In Appendix 9 we analyze one such friction that would arise with incomplete information in a linear-cost model, and there we allow governments to optimally choose the vaccine price in addition to the at-risk capacity subsidy. Our analysis confirms that while the details are altered, the main insights that derive from our Benchmark Model continue to apply in that setting.

4 Vaccine Production with Offshored Inputs

In this section we consider how the presence of offshored inputs complicates the optimal design of vaccine policy. We do so in the simplest way, by continuing with the Benchmark Model but now assuming that the domestic vaccine producer sources its inputs from a single *foreign* input supplier.¹⁹ We focus on the at-risk capacity investment chosen by the foreign input supplier, who we assume faces a cost of capacity investment $c^*(v^*)$ with the same properties as $c(v)$, and who can produce inputs at a marginal cost of κ^* up to installed capacity. The domestic vaccine producer converts each unit of input acquired from the foreign input supplier into a vaccine dose, a conversion process which we continue to assume is done costlessly and is not capacity constrained. For now we also continue to assume that the government designs its vaccine policy unilaterally, postponing consideration of the possible role that an international agreement on vaccine policy could play until section 5.

The capacity choice of the foreign input supplier v_t^* in each period t proceeds in an identical way to that described for the domestic firm's choice of v_t in section 3.1, except that in each period where there exists the capacity to produce inputs, the price and quantity of inputs exchanged between the foreign input supplier and the domestic vaccine producer are now determined by bilateral Nash bargaining similar to Antras and Staiger (2012), with capacity costs incurred up to that point in time taken as bygones and each party's outside option being zero and with the bargaining power of the domestic vaccine producer denoted by $\alpha \in [0, 1]$. And with the input price and quantity determined for the period, the input supplier then produces the contracted level of inputs at the marginal cost κ^* up to its installed capacity.

¹⁹Our analysis assumes that the foreign input supplier is not owned by the domestic vaccine producer, which as we describe in section 2 seems empirically to be the most relevant case.

It is convenient in this section to begin with the analysis of the capacity investment choices of the foreign input supplier, and then turn to the choices of the domestic government. As in section 3.1, we continue to take as exogenous the price P that the domestic vaccine producing firm receives for each delivered dose of the vaccine that is deemed to be safe and effective, and in analogy with Assumption II we now impose

$$\max\left[\frac{c^*(L)}{(T - (M + m))L}, \frac{c^{*'}(L)}{s}\right] < (1 - \alpha)(P - \kappa^*). \quad (\text{Assumption II'})$$

The difference between Assumption II' and Assumption II is the presence of the term $(1 - \alpha)$ on the right-hand side of Assumption II', reflecting the fact that in its Nash bargain with the domestic vaccine producer the foreign input supplier receives the bargaining share $(1 - \alpha)$ of the operating profits $(P - \kappa^*)$ for each dose of the vaccine sold in the domestic country.

Under Assumption II' we can cast the foreign input supplier's optimal choice of capacity investment as its choice of the stock of at-risk capacity V_a^* at the end of period $M + 1 - m$, just as we did for the local-sourcing analysis of section 3.1. The foreign input supplier's expected profits can therefore be written as

$$E[\pi^*(V_a^*)] = s(1 - \alpha)[V_a^*m + (T - (M + m))L] \times (P - \kappa^*) - [C^*(V_a^*) + sc^*(L - V_a^*)], \quad (12)$$

where $C^*(V_a^*) \equiv (M + 1 - m) \times c^*(v_a^*(V_a^*))$, where $v_a^*(V_a^*) \equiv \frac{V_a^*}{(M+1-m)}$. Using (12), the first-order condition that determines the profit-maximizing choice of at-risk capacity investment for the foreign input supplier, \tilde{V}_a^* , is given by

$$\frac{\partial E[\pi^*(V_a^*)]}{\partial V_a^*} = sm \times (1 - \alpha)(P - \kappa^*) - [C^{*'}(\tilde{V}_a^*) - sc^{*'}(L - \tilde{V}_a^*)] = 0, \quad (13)$$

with the second-order condition again guaranteed by the convexity of the cost-of-capacity function $c^*(v^*)$. As a comparison of (13) with (7) confirms, if $\kappa^* = \kappa$ and if the foreign input supplier has all the bargaining power (i.e., if $\alpha = 0$), then its at-risk capacity choice \tilde{V}_a^* is identical to \tilde{V}_a , the at-risk capacity choice of the vaccine-producing firm of section 3.1 that locally sources its inputs; but if the domestic vaccine producer has any bargaining power (i.e., if $\alpha > 0$), then $\tilde{V}_a^* < \tilde{V}_a$ reflecting the hold-up problem that the foreign input supplier faces. This marks one new dimension of the policy environment faced by the domestic government when it confronts an offshoring firm rather than the vaccine-producing firm analyzed in section 3.1 that was assumed to locally source its inputs.

However, a more consequential difference in the policy environments arises from the difference in the government's own objectives across the two settings. Recall that in our analysis of section 3.1 we abstracted from the distributional impacts that government policy toward vaccine production might have, and implicitly assumed that the government has lump-sum instruments that it can use to redistribute to its citizens any undesired rents that its vaccine policies

create for its domestic vaccine producers. This allowed us to adopt the view that the objective of the government/benevolent social planner is to minimize the overall expected cost to its country of a viral outbreak, without regard to the size of the rents transferred to vaccine producers.

In the present setting where these rents will accrue to *foreign* input suppliers, this objective would still be reasonable for a *global* social planner who takes the interests of both domestic and foreign countries into consideration, and in the next section we will consider the role that an international agreement on vaccine supply chains might play in helping to implement the outcome preferred by a global social planner. But here we maintain our focus on what unilateral interventions a domestic government would pursue, and for the domestic government any rents paid to foreign input suppliers cannot be recouped through lump-sum taxation. Hence, the objective of the domestic government/domestic social planner will be to minimize the overall expected cost to its country of a viral outbreak, where any rents transferred to foreign input suppliers are counted as part of that cost.

To write down the domestic government objective function, we first use (13) to define $\tilde{\sigma}$ as the subsidy per unit of foreign at-risk capacity investment that would induce the foreign input supplier to invest in at-risk capacity at the level V_a^* .²⁰

$$\tilde{\sigma} = [C^{*'}(V_a^*) - sc^{*'}(L - V_a^*)] - sm(1 - \alpha)(P - \kappa^*) \equiv \tilde{\sigma}(V_a^*, P - \kappa^*, \alpha), \quad (14)$$

where the subsidy $\tilde{\sigma}$ is increasing in V_a^* and decreasing in $P - \kappa^*$ and also increasing in the severity of the hold-up problem as captured by the bargaining power parameter α . We will adopt the convention that the domestic government uses a subsidy policy defined by $\tilde{\sigma}(V_a^*, P - \kappa^*, \alpha)$ to achieve any desired level of at-risk capacity investment by the foreign input supplier V_a^* given the (exogenously fixed) price of vaccines in the domestic market, though it is clear that the government could in principle achieve the same ends by setting the vaccine price appropriately.

Finally, we impose the analog of Assumption I for the setting where the vaccine input is sourced offshore, namely

$$(1 - \alpha)(P - \kappa^*) \leq (\delta h - \kappa^*). \quad (\text{Assumption I}')$$

Under Assumption I' and Assumption II', the foreign input-producing firm and domestic government incentives for not-at-risk investment V_n are aligned, because if the clinical trials of the vaccine are declared a success at the beginning of period $M + 1$ then under Assumption II' the firm will always top off its at-risk capacity investments to achieve capacity level L in period $M + 1$, and the government will approve of this choice because the payment per dose that the domestic vaccine-producing firm makes to the foreign input supplier is

²⁰As before, it can be confirmed that introducing such a subsidy cannot cause Assumption II' to be violated and so we may continue to focus on the expression for $E[\pi^*(V_a^*)]$ in (12) when calculating the first-order condition that defines the foreign input supplier's choice of V_a in the presence of the subsidy.

$(1 - \alpha)(P - \kappa^*) + \kappa^*$, which is less than the expected health benefit per dose δh according to Assumption I'. Hence, as in the analysis of section 3.1, the only remaining question is whether the incentives of the foreign input supplier and the domestic government are also aligned with regard to the level of at-risk investment, or whether instead the domestic government will want to intervene with an investment subsidy to alter the foreign input supplier's choice of at-risk investment according to the subsidy function $\tilde{\sigma}(V_a^*, P - \kappa^*, \alpha)$.²¹

To explore this possibility, we proceed as before and cast the government's problem as simply its choice of the stock of at-risk capacity V_a in period M that will minimize the expected overall cost to the country of the viral outbreak, where the government will orchestrate this choice with an appropriate subsidy to the foreign input supplier's at-risk capacity investment according to the function $\tilde{\sigma}(V_a^*, P - \kappa^*, \alpha)$. The expected overall cost to the domestic country of the viral outbreak can be written as

$$E[Cost(V_a^*)] = \tilde{\sigma}(V_a^*, P - \kappa^*, \alpha) \times V_a^* + s(1 - \alpha)[V_a^* m + (T - (M + m))L] \times (P - \kappa^*) + s[V_a^* m + L \times (T - (M + m))] \times \kappa^* + (1 - s)LT\delta h + s[LM\delta h + (L - V_a^*)m\delta h]. \quad (15)$$

The difference between the government's objective in the presence of foreign offshoring (contained in (15)) and in its absence (contained in (3)) appears in the terms that capture the economic cost to the domestic country of acquiring the inputs. In (3), where the inputs are domestically produced, the expected cost of capacity investment $c(V_a) + sc(L - V_a)$ captures the economic cost to the country of the vaccine program. In (15) by contrast, where the inputs are offshored, the economic cost to the country of the vaccine program now comes in two parts: the subsidy payments made by the domestic government to the foreign input supplier for at-risk capacity investment, $\tilde{\sigma}(V_a^*, P - \kappa^*, \alpha) \times V_a^*$, and the expected payments made by the domestic vaccine producer to the foreign input supplier for the inputs exchanged between the two, $s(1 - \alpha)[V_a^* m + (T - (M + m))L] \times (P - \kappa^*) + s[V_a^* m + (T - (M + m))L] \times \kappa^*$. And according to (15) and also using (14), the first-order condition that defines the domestic government's preferred level of at-risk capacity investment in the presence of foreign offshoring, which we denote by \bar{V}_a^* , is given by

$$\frac{\partial E[Cost(V_a^*)]}{\partial V_a^*} = C^{*'}(\bar{V}_a^*) - s[c^{*'}(L - \bar{V}_a^*) + m(\delta h - \kappa^*)] + \bar{V}_a^* \times \frac{\partial \tilde{\sigma}(V_a^*, P - \kappa^*, \alpha)}{\partial V_a^*} = 0, \quad (16)$$

with the second-order condition guaranteed by the convexity of the cost-of-capacity function $c^*(v^*)$. The subsidy offered by the government to the at-risk capacity investments of the offshore input supplier is then given by $\tilde{\sigma}(\bar{V}_a^*, P - \kappa^*, \alpha)$.

²¹With Assumption II' guaranteeing that the foreign input supplier will choose V_a^* with $V_n^* = L - V_a^*$, and with the government's choice of capacity subsidy not altering this property (see note 20), Assumption I' does not require the additional condition that is included in Assumption I where, in effect, with local sourcing of inputs the government is able to choose the levels of capacity investment directly.

Notice that if the domestic government had the ability to tax as well as subsidize the at-risk capacity investment of the foreign input supplier – that is, if $\tilde{\sigma}$ could be negative – then it might be possible that the tax revenue collected from the foreign input supplier would dominate the choice of vaccine policy, and a negative $\tilde{\sigma}$ might be chosen. We will restrict our attention to non-negative choices of $\tilde{\sigma}$, reflecting the fact that while the domestic government could reasonably offer a subsidy to the at-risk capacity investments of foreign firms operating abroad, it would not have the authority to unilaterally tax these investments. Formally, we will assume that if the (unconstrained) first-order condition (16) implies a negative value of $\tilde{\sigma}$, then this simply means that the domestic government's choice of $\tilde{\sigma}$ is driven to a corner solution with $\tilde{\sigma} = 0$.

Comparing (16) with (4) and noting that the last term in (16) is strictly positive according to (14) and is absent from (4), it is clear that if $\kappa^* = \kappa$ and when $\alpha = 0$ so that the foreign input supplier has all the bargaining power and hence faces no hold-up problem, \bar{V}_a^* is smaller than \hat{V}_a provided that the vaccine price P is fixed at a level which lies strictly below the social value of a vaccine dose (i.e., provided that $P < \delta h$). The reason is that the domestic government has to concern itself with lost subsidy payments to foreign interests when it is attempting to use subsidies to achieve the right level of at-risk capacity investment for a domestic vaccine-producing firm that offshores its inputs. And when $\alpha > 0$ and the foreign input supplier is subject to hold up from the domestic vaccine producer, the vaccine price P need only be lower than $\frac{\delta h}{1-\alpha}$ to ensure that $\bar{V}_a^* < \hat{V}_a$, because hold up provides its own reason for underinvestment.

We summarize with:

Proposition 2 *When the domestic government designs its vaccine policy unilaterally and faces a domestic vaccine-producing firm that offshores its inputs, the government's preferred level of at-risk capacity investment \bar{V}_a^* is smaller than the level \hat{V}_a that it prefers when facing a domestic vaccine-producing firm that locally sources its inputs from a supplier with the same marginal cost, provided that $P < \frac{\delta h}{1-\alpha}$ so that the vaccine price is fixed at a level strictly below the social value of a vaccine dose adjusted for the degree of hold-up that the foreign input supplier experiences in its relationship with the domestic vaccine producer.*

The fact that $\bar{V}_a^* < \hat{V}_a$ also means that, all else equal and when the government designs its vaccine policy unilaterally, the expected health cost of the pandemic, which for any level of investment in at-risk capacity V_a is given by

$$E[H(V_a)] \equiv (1-s)LT\delta h + s[LM\delta h + (L - V_a)m\delta h],$$

will be lower if the government can work with a vaccine-producing firm that locally sources its inputs than if the government must work with a domestic vaccine producer that offshores its inputs: that is, $\bar{V}_a^* < \hat{V}_a$ implies

$$E[H(V_a = \hat{V}_a)] < E[H(V_a = \bar{V}_a^*)].$$

This follows because when the government faces a domestic vaccine producer that offshores its inputs, it has to weigh the expected health benefits of inducing additional at-risk capacity investment against the higher subsidy payments that are needed for this inducement and which are accruing to foreign firms. We record this in the following:

Corollary *All else equal and when the government designs its vaccine policy unilaterally, the expected health cost of the pandemic will be lower if a government can work with a vaccine-producing firm that locally sources its inputs than if the government must work with a domestic vaccine producer that offshores its inputs.*

Our discussion above has focused on the case where $\kappa^* = \kappa$ and so has ignored the possible cost savings that might come from offshoring rather than locally sourcing vaccine inputs (i.e., the case where $\kappa^* < \kappa$). But at this point it should be clear that if the cost advantage achieved by the domestic vaccine-producing firm in offshoring its inputs is not too large, the domestic government can achieve a higher level of welfare by prohibiting imports of vaccine inputs and forcing its vaccine-producing firm to “reshore” and either domestically out-source the inputs or produce the inputs in house (vertically integrate), because this would allow the domestic government to then subsidize at-risk capacity investment to achieve the level \tilde{V}_a without losing the subsidy payments to foreign firms. However, as we detail in the next section, a policy to promote reshoring would be (at best) a second-best response. The real problem that the domestic government must confront when it faces a domestic vaccine-producing firm that offshores inputs is the positive international externality created by the government subsidy program, which can be addressed in a first-best way by an international agreement to encourage subsidies to offshored inputs for vaccine production.

5 Vaccine Policy and International Cooperation

Bown (2022c) observes that there may be a new role for the WTO in helping its member governments to cooperate on policies that can encourage vaccine supply chain resilience during a pandemic. In this section we consider the role for international cooperation over vaccine supply-chain policy that is suggested by our results of the previous sections.

To this end, we add to the offshoring model of section 4 a second vaccine-producing firm, where this second vaccine-producing firm is located in the foreign country. We assume that each of the vaccine-producing firms can either source inputs locally or offshore inputs from the other country, with Assumption I-Assumption II and Assumption I'-Assumption II' applying. And we allow each government to set its own vaccine policy, either unilaterally or possibly under a negotiated international agreement, maintaining our focus on the case where $P < \frac{\delta h}{1-\alpha}$ as emphasized in Proposition 2.

To keep focused on the main points, we assume that vaccines themselves

cannot be internationally traded, and that vaccine inputs can be freely traded. Moreover, to introduce a meaningful choice between offshoring and the local sourcing of inputs, we assume that the vaccine inputs produced in the foreign country may (or may not) be ideally suited for use in vaccine production in the domestic country, and that the vaccine inputs produced in the domestic country may (or may not) be ideally suited for use in vaccine production in the foreign country. We capture this by assuming that, for each of the vaccine-producing firms, the marginal cost of an input produced offshore is given by κ_{Off} while the marginal cost of a locally sourced input is given by κ_{Loc} . When $\kappa_{Loc} > \kappa_{Off}$, the marginal cost of vaccine inputs can be lowered by offshoring; when $\kappa_{Loc} \leq \kappa_{Off}$, the marginal cost of vaccine inputs is lowest under local sourcing. The difference $[\kappa_{Loc} - \kappa_{Off}]$ (which we allow to be positive or negative) therefore reflects the degree of cost-saving that can be achieved through the offshoring of vaccine inputs.

Finally, we assume that if an input-producing firm were to produce inputs for both local sales and export, it would install separate capacity to produce the inputs destined for each market and would therefore face separate convex capacity costs for each market. This assumption ensures that the capacity decisions of each input-producing firm are separable across the two markets. We have in mind, for example, an environment where in each country, exports must be shipped from a port that is situated far way from the geographic center of local demand, and where internal transport costs in each country are high. The high internal transport costs would then create a reason for the input-producing firm to locate production capacity for exports at the port and to place production capacity for serving local demand at an interior location that coincides with the geographic center of local demand. As before, we abstract from any international health ramifications of a country's vaccine policies and assume that the health outcomes in each country are determined entirely by the country's own vaccine policy choices.

Notice that with these assumptions the foreign country is now a mirror image of the domestic country, and that the mirror-image vaccine problem for each country is separable. With each country's input producer facing a separable capacity installation problem across its local and export market, the policy choices of the two governments are therefore strategically independent. This is not necessary for any of the results we emphasize below, but it allows us to make direct use of the results from sections 3 and 4 in deriving these results.

5.1 Noncooperative outcome

Consider first the situation faced by each government if there is no possibility of negotiating an agreement over vaccine supply chain policies. In keeping with our assumption that governments have lump-sum instruments that they can use to redistribute to their respective citizens any undesired rents that their vaccine policies create for their vaccine producers, in what follows we assume that each government can use these instruments to ensure that the organization of its firm's supply chain – vertical integration/domestic outsourcing versus

offshoring – is chosen by its vaccine-producing firm in accordance with the government’s preferences. So in effect, each government chooses (noncooperatively) whether its vaccine-producing firm will offshore inputs or rather source the inputs domestically, and then selects (noncooperatively) the subsidy to at-risk capacity investment that it will offer. Given the strategic independence of the policy decisions made by each government and the symmetric features of the two countries, we can without loss of generality focus on the policy decision of the domestic government.

If the government induces its vaccine-producing firm to offshore, then it faces the situation modeled in section 4 (with the marginal cost parameter κ_{Off} now replacing κ^*), and as we have shown there it would offer to the foreign input supplier a subsidy to investment in at-risk capacity in the amount $\tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha)$, inducing the at-risk capacity level \bar{V}_a^* and achieving a welfare level implied by its own noncooperative policy choices of

$$E[Cost_{Off}^N(\bar{V}_a^*)] = \tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha) \times \bar{V}_a^* + s(1 - \alpha)[\bar{V}_a^* m + (T - (M + m))L] \times (P - \kappa_{Off}) + s[\bar{V}_a^* m + (T - (M + m))L] \times \kappa_{Off} + (1 - s)LT\delta h + s[LM\delta h + (L - \bar{V}_a^*)m\delta h]. \quad (17)$$

where we now use the notation $E[Cost_{Off}^N(\bar{V}_a^*)]$ to denote the expected overall cost to the country of the viral outbreak when the government induces offshoring and selects noncooperatively the subsidy that it offers to at-risk capacity investment.

On the other hand, if the government induces its vaccine-producing firm to “re-shore” inputs and either produce them itself or outsource them domestically, then it faces the situation modeled in section 3 (with the marginal cost parameter κ_{Loc} now replacing κ), and as we have shown there it would offer to the local input supplier a subsidy to investment in at-risk capacity in the amount $\hat{\sigma}(P)$, inducing the at-risk capacity level \hat{V}_a^* and achieving a welfare level implied by its own noncooperative policy choices of

$$E[Cost_{Loc}^N(\hat{V}_a)] = C(\hat{V}_a) + sc(L - \hat{V}_a) + s[m\hat{V}_a + (T - (M + m))L]\kappa_{Loc} + (1 - s)LT\delta h + s[LM\delta h + (L - \hat{V}_a)m\delta h]. \quad (18)$$

Note that when $\kappa_{Loc} = \kappa_{Off} \equiv \kappa$, Proposition 2 and its Corollary apply and so $\hat{V}_a > \bar{V}_a^*$. And with \hat{V}_a declining in κ_{Loc} and \bar{V}_a^* declining in κ_{Off} , it follows that for $\kappa_{Loc} - \kappa_{Off}$ positive but below a threshold level Proposition 2 and its Corollary continue to apply as well. On the other hand, if the marginal cost reductions associated with offshoring are sufficiently large in the sense that $\kappa_{Loc} \gg \kappa_{Off}$ (with both κ_{Loc} and κ_{Off} still satisfying Assumption I-Assumption II and Assumption I'-Assumption II'), then the results of Proposition 2 and its Corollary will be reversed: the implied efficiency loss from reshoring ($\kappa_{Loc} - \kappa_{Off}$) would be so high that this alone would lead to a smaller level of at-risk capacity preferred by the government under reshoring, despite the avoidance of subsidy leakage that reshoring affords.

In any event, armed with (17) and (18), we can now partition the noncooperative outcome into two cases. If $E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] < 0$, then with noncooperative policy choices the expected overall cost to a country of the viral outbreak will be lower under offshoring than under local sourcing, and each government will induce its vaccine-producing firm to offshore inputs, and it will offer the subsidy $\tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha)$ to the at-risk capacity investment of offshore input producers in order to achieve for its vaccine producers the at-risk capacity level \bar{V}_a^* for their offshored inputs. And if $E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] \geq 0$, then the opposite choice will be made, and in the noncooperative equilibrium each government will induce its vaccine-producing firm to reshore inputs, and it will offer the subsidy $\hat{\sigma}(P)$ to the at-risk capacity investment of domestic input producers in order to achieve for its vaccine producers the at-risk capacity level \hat{V}_a for their locally sourced inputs.

To gain insight into the conditions under which either offshoring or reshoring occurs in the noncooperative equilibrium, it is useful to write $E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)]$ in the more explicit form

$$E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] = \left(\left(E[\pi_{Loc}(\hat{V}_a)] - \hat{\sigma}(P) \times \hat{V}_a \right) - E[\pi_{Off}(\bar{V}_a^*)] \right) + \bar{V}_a^* \times \tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha) + [\hat{V}_a - \bar{V}_a^*] \times \hat{\sigma}(P), \quad (19)$$

where $E[\pi_{Loc}(\hat{V}_a)]$ is the expected profit earned by the domestic vaccine-producing firm under local sourcing, which is given by

$$E[\pi_{Loc}(\hat{V}_a)] = s[\hat{V}_a m + (T - (M + m))L] \times (P - \kappa_{Loc}) - [C(\hat{V}_a) - \hat{\sigma}(P) \times \hat{V}_a + sc(L - \hat{V}_a)],$$

and where $E[\pi_{Off}(\bar{V}_a^*)]$ is the expected profit earned by the domestic vaccine-producing firm under offshoring, which is given by

$$E[\pi_{Off}(\bar{V}_a^*)] = s\alpha[\bar{V}_a^* m + (T - (M + m))L] \times (P - \kappa_{Off}),$$

its share of the bargaining surplus with the foreign input supplier. According to (19), for any level of the difference in expected net-of-subsidy profits earned by the domestic firm under local sourcing relative to its profits under offshoring $\left(\left(E[\pi_{Loc}(\hat{V}_a)] - \hat{\sigma}(P) \times \hat{V}_a \right) - E[\pi_{Off}(\bar{V}_a^*)] \right)$, it is more likely that $E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] \geq 0$, and therefore more likely that in the noncooperative equilibrium each government will induce its vaccine-producing firm to reshore inputs, the larger is the equilibrium subsidy leakage under offshoring $\bar{V}_a^* \times \tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha)$, and the larger is the discrepancy between the social and commercial value of any increased at-risk capacity that comes with reshoring $[\hat{V}_a - \bar{V}_a^*] \times \hat{\sigma}(P)$.²² Moreover, note that

²²Recall that $\hat{V}_a > \bar{V}_a^*$ provided that $(\kappa_{Loc} - \kappa_{Off})$ is below a critical threshold. If $(\kappa_{Loc} - \kappa_{Off})$ is above this threshold, the statement in the text still applies, with the phrase “the smaller in magnitude is the discrepancy between the social and commercial value of any decreased at-risk capacity that comes with reshoring” replacing the phrase “the larger is the discrepancy between the social and commercial value of any increased at-risk capacity that comes with reshoring” that appears in the text.

$\left(\left(E[\pi_{Loc}(\hat{V}_a)] - \hat{\sigma}(P) \times \hat{V}_a \right) - E[\pi_{Off}(\bar{V}_a^*)] \right)$ can be written as

$$\begin{aligned} & \left(\left(E[\pi_{Loc}(\hat{V}_a)] - \hat{\sigma}(P) \times \hat{V}_a \right) - E[\pi_{Off}(\bar{V}_a^*)] \right) = \\ & s[\hat{V}_a m + (T - (M + m))L] \times (P - \kappa_{Loc}) - s\alpha[\bar{V}_a^* m + (T - (M + m))L] \times (P - \kappa_{Off}) \\ & \quad - [C(\hat{V}_a) + sc(L - \hat{V}_a)], \end{aligned}$$

implying that, through its effects on the expected profits of the domestic firm and all else equal, a larger $[\kappa_{Loc} - \kappa_{Off}]$ (higher marginal cost under local sourcing as compared to offshoring) and/or a larger α (higher domestic firm bargaining power under offshoring) tends to tilt government preferences toward offshoring.

We summarize in:

Proposition 3 *In the noncooperative equilibrium and for any level of the difference in expected net-of-subsidy profits earned by the domestic firm under local sourcing relative to its profits under offshoring, it is more likely that each government will induce its vaccine-producing firm to reshore inputs: (i) the larger is the equilibrium subsidy leakage under offshoring; and (ii) the larger is the discrepancy between the social and commercial value of any increased at-risk capacity that comes with reshoring. Moreover, through its effects on the expected profits of the domestic firm and all else equal, a higher marginal cost under local sourcing as compared to offshoring and/or a higher domestic firm bargaining power under offshoring tends to tilt the noncooperative equilibrium toward offshoring.*

Proposition 3 partitions the noncooperative equilibrium outcomes into those that exhibit offshoring of vaccine inputs and those that exhibit local sourcing of inputs. In the next section we allow for the possibility of international cooperation over vaccine policy, and we establish that the nature of the case for international cooperation hinges critically on which of these noncooperative outcomes obtains.

5.2 Cooperation over vaccine supply chain policies

In section 5.1 we considered the situation faced by each government when there is no possibility of negotiating an agreement over vaccine supply chain policies. We now ask whether governments could find an international agreement that could offer joint benefits to the governments over the noncooperative outcomes described in section 5.1. Again given the symmetry of the model, we can focus on the payoffs of the domestic government. We organize our analysis in this section around the partition of noncooperative outcomes characterized in section 5.1.

When the noncooperative equilibrium features offshoring Consider first the case where $E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] < 0$, and therefore where the noncooperative outcome features offshoring of vaccine inputs and a subsidy

to investment in at-risk capacity offered by the government to its offshore input supplier in the amount $\tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha)$, inducing the at-risk capacity level \bar{V}_a^* and achieving for the government the welfare level implied by its own noncooperative policy choices of

$$E[Cost_{Off}^N(\bar{V}_a^*)] = \tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha) \times \bar{V}_a^* + s(1 - \alpha)[\bar{V}_a^* m + (T - (M + m))L] \times (P - \kappa_{Off}) + s[\bar{V}_a^* m + (T - (M + m))L] \times \kappa_{Off} + (1 - s)LT\delta h + s[LM\delta h + (L - \bar{V}_a^*)m\delta h].$$

In section 5.1 it was sufficient to focus on the welfare level implied by a government's own noncooperative policy choices, given separability and the strategic independence of the policy choices of the two governments and the fact that we were simply interested in characterizing noncooperative policy choices. But to assess the possibility of government welfare gains from international cooperation over vaccine policy, we now need an expression for the overall welfare that a government experiences in the noncooperative equilibrium, including the welfare implications of the other government's policy intervention.

The overall welfare of the domestic government in the noncooperative equilibrium when offshoring prevails, as measured by the expected overall cost to the country of the viral outbreak at home and abroad, and which we denote by $E[Cost_{Off, Off}^N(\bar{V}_a^*, \bar{V}_a)]$, is given by

$$\begin{aligned} E[Cost_{Off, Off}^N(\bar{V}_a^*, \bar{V}_a)] &= \\ &\tilde{\sigma}(\bar{V}_a^*, P - \kappa_{Off}, \alpha) \times \bar{V}_a^* + s(1 - \alpha)[\bar{V}_a^* m + (T - (M + m))L] \times (P - \kappa_{Off}) \\ &+ s[\bar{V}_a^* m + (T - (M + m))L] \times \kappa_{Off} + (1 - s)LT\delta h + s[LM\delta h + (L - \bar{V}_a^*)m\delta h] \\ &- (s(1 - \alpha)[\bar{V}_a^* m + (T - (M + m))L] \times (P - \kappa_{Off}) - [C(\bar{V}_a) - \tilde{\sigma}(\bar{V}_a, P - \kappa_{Off}, \alpha) \times \bar{V}_a + sc(L - \bar{V}_a)]) \\ &= C(\bar{V}_a) + sc(L - \bar{V}_a) + s[\bar{V}_a^* m + (T - (M + m))L] \times \kappa_{Off} \\ &\quad + (1 - s)LT\delta h + s[LM\delta h + (L - \bar{V}_a^*)m\delta h]. \end{aligned}$$

The terms in the first two lines are simply $E[Cost_{Off}^N(\bar{V}_a^*)]$, the welfare level implied by the domestic government's own noncooperative policy choices, while the terms in the third line represent the economic surplus captured by the domestic input producer. The expression after the second equality follows from observing that $\bar{V}_a^* = \bar{V}_a$. Notably, with the noncooperative subsidy payments made by both governments to their offshore input suppliers offsetting each other, and with direct consequences of the division of bargaining surplus over the exchanged inputs cancelling out as well given that each country has a firm that plays each role in the bilateral bargain, this expression is identical to $E[Cost(V_a = \bar{V}_a)]$ for $\kappa = \kappa_{Off}$, where $E[Cost(V_a)]$ is defined in (3) and gives the expression for the expected overall cost to a country of the viral outbreak in the local sourcing model of section 3.1.

But it is now immediate that the two governments can gain from an international agreement over vaccine supply chain policies that implements free trade in inputs and has each government agree to offer to its own input-producing firm

a subsidy to at-risk capacity investment at the level $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$, thereby achieving for the input producer in each country the level of at-risk capacity \hat{V}_a . This follows because, as established in section 3.1, the level of at-risk capacity \hat{V}_a minimizes $E[Cost(V_a)]$ by (4), and by Proposition 2 (and with marginal costs held fixed at $\kappa = \kappa_{Off}$) \hat{V}_a is strictly greater than \bar{V}_a^* , therefore ensuring that

$$E[Cost(\hat{V}_a)] < E[Cost(\bar{V}_a)] = E[Cost_{Off, Off}^N(\bar{V}_a^*, \bar{V}_a)],$$

with symmetric arguments holding for the foreign country.

Intuitively, the subsidy $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$ induces the level of at-risk capacity \hat{V}_a defined by (4) in section 3.1 that a government would choose for its locally sourcing vaccine-producing firm – where there is no hold-up problem and where the government has no concern over subsidy leakage – and with the marginal cost of producing vaccine inputs κ set equal to κ_{Off} . The subsidy $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$ is designed to solve the hold-up problem that arises for $\alpha > 0$ (as in Antras and Staiger, 2012), and given free trade the vaccine-producing firms in each country will be able to achieve the full efficiency gains from offshoring and face marginal costs κ_{Off} of vaccine input production. And finally, given the symmetry between countries there will in fact be no net subsidy leakage under this arrangement, since each government is offering the same subsidy payments to at-risk capacity investments for inputs used by the other country’s vaccine producers.²³

With this agreement the two governments would therefore have found a way to enjoy the full efficiency benefits from offshoring while neutralizing the subsidy leakage problem that in the absence of an agreement would have led them to choose inefficiently low subsidies to at-risk capacity in the presence of offshoring. We summarize with

Proposition 4 *If the noncooperative outcome features offshoring of vaccine inputs and provided that $P < \frac{\delta h}{1-\alpha}$, the at-risk capacity chosen by each input supplier in the noncooperative equilibrium, \bar{V}_a^* , will be inefficiently low. An international agreement that implements free trade in vaccine inputs and commits each government to offer a per-unit subsidy in the amount of $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$ to the at-risk capacity investments of its own input-producing firm can achieve the efficient level of at-risk capacity for each input supplier, \hat{V}_a , and provides mutual gains for each country.*

We have characterized the international agreement on vaccine supply chain policy described in Proposition 4 as an agreement under which each government commits to offering the subsidy $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$ to the at-risk capacity investment of its *local* input-supplying firm, even though under the free-trade policies dictated by the agreement each country’s local input-supplying firm produces inputs only for export. This particular design of the efficient subsidy commitments seems the most natural in the context of an international

²³Symmetry is not required for this result, as long as the international agreement allows countries to exchange transfers, but symmetry obviates the need for international transfers in our interpretation of the agreement described above.

agreement because under this design each government provides direct subsidy payments only to local firms. But we could have alternatively and equivalently characterized the agreement as involving a commitment by each government to subsidize the at-risk capacity of its vaccine producer's *offshore* input supplier at the level $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$, more closely mirroring the geographical pattern of subsidy payments that would be suggested by the results of section 4 if governments were committed to free trade in vaccine inputs but did not have an agreement on subsidies to investment in at-risk capacity and therefore chose these subsidies unilaterally.

With this equivalence clarified, a second point is now also apparent: the subsidy commitment described by Proposition 4 represents a commitment to a *higher* subsidy than each government would choose unilaterally. This follows from the fact that the subsidy function $\tilde{\sigma}$ is increasing in the level of at-risk capacity to be induced by the subsidy, and that the unilateral subsidy choice would induce \hat{V}_a^* which is smaller than \hat{V}_a as confirmed in section 4. The implied subsidy $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$ called for in the agreement is therefore larger than $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha)$, the noncooperative subsidy when the noncooperative outcome features offshoring of vaccine inputs. Intuitively, the underlying feature that creates the inefficiency of unilateral at-risk capacity subsidies is the positive international externality created by each government's subsidy program and associated with the subsidy leakage that occurs when it is efficient for vaccine producers to offshore their inputs and when those input suppliers must make important at-risk capacity decisions. An international agreement can internalize these positive externalities, and thereby raise to the first-best level the subsidies that governments offer to at-risk capacity investment for the production of offshored vaccine inputs.

We summarize this point with the following

Corollary *When the noncooperative outcome features offshoring of vaccine inputs and provided that $P < \frac{\delta h}{1-\alpha}$, it is the positive externality generated by subsidies to at-risk capacity investment when offered to offshored input suppliers that drives the need for the agreement characterized in Proposition 4.*

When the noncooperative equilibrium features local sourcing Consider next the case where $E[Cost_{Off}^N(\hat{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] \geq 0$, and therefore where the noncooperative outcome features local sourcing of vaccine inputs and a subsidy to investment in at-risk capacity offered by the government to the local input supplier in the amount $\hat{\sigma}(P)$, inducing the at-risk capacity level \hat{V}_a and achieving for the government the welfare level implied by its own noncooperative policy choices of

$$E[Cost_{Loc}^N(\hat{V}_a)] = C(\hat{V}_a) + sc(L - \hat{V}_a) + s[m\hat{V}_a + (T - (M + m))L]\kappa_{Loc} \\ + (1 - s)LT\delta h + s[LM\delta h + (L - \hat{V}_a)m\delta h].$$

In contrast to the case of offshoring, the magnitude of $E[Cost_{Loc}^N(\hat{V}_a)]$ *does* reflect the full measure of welfare for the domestic government when the nonco-

operative outcome features local sourcing, since in this case there is no international trade related to vaccine supply chains in either direction. And conditional on local sourcing, the level of at-risk capacity achieved under noncooperative policies \hat{V}_a is in fact efficient, because it is the minimizer of $E[Cost_{Loc}^N(V_a)]$ as confirmed by (4).

Hence, to determine whether there is a role for an international agreement on vaccine supply chains in this case, the only remaining question is whether local sourcing is in fact the efficient arrangement of the supply chain when $E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] \geq 0$: if it is, then there is no role for an international agreement, because conditional on local sourcing the unilateral choices that each government makes related to vaccine policy will be efficient; but if local sourcing is not the efficient arrangement, then an international agreement that facilitates a switch to efficient offshoring can provide mutual gains to the two governments. Denoting by $E[Cost_{Off,Off}^C(\hat{V}_a, \hat{V}_a)]$ the overall welfare of the domestic government in the cooperative equilibrium under offshoring, as measured by the expected overall cost to the country of the viral outbreak at home and abroad and defined as

$$E[Cost_{Off,Off}^C(\hat{V}_a, \hat{V}_a)] = C(\hat{V}_a) + sc(L - \hat{V}_a) + s[\hat{V}_a m + (T - (M + m))L] \times \kappa_{Off} + (1 - s)LT\delta h + s[LM\delta h + (L - \hat{V}_a)m\delta h], \quad (20)$$

we may therefore state:

Lemma 3 *A role for an international agreement on vaccine supply chains will arise when $E[Cost_{Off}^N(\bar{V}_a^*)] - E[Cost_{Loc}^N(\hat{V}_a)] \geq 0$ if and only if it is also true that*

$$E[Cost_{Off,Off}^C(\hat{V}_a, \hat{V}_a)] - E[Cost_{Loc}^N(\hat{V}_a)] < 0. \quad (\text{Condition I})$$

According to Condition I, if each country's offshore input suppliers were to invest in the efficient level of at-risk capacity \hat{V}_a , then each country would experience a lower expected overall cost over the period of contagion than it would under local sourcing with the at-risk investment level \hat{V}_a . Lemma 3 therefore simply points out that even if governments would choose local sourcing of vaccine inputs over offshoring when they expect *noncooperative* vaccine policies, it is still possible that under *cooperative* vaccine policies they would prefer offshoring.

An interesting feature of the case for an international agreement on vaccine policies identified in Lemma 3 arises when $\alpha = 0$ and there is no holdup problem associated with offshoring. Using (4) and (14), both evaluated at κ_{Off} , it follows that $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha) = sm \times (\delta h - \kappa_{Off} - (1 - \alpha)(P - \kappa_{Off}))$, so when $\alpha = 0$ we have $\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha = 0) = sm \times (\delta h - P) = \hat{\sigma}(P)$. That is, in this case the level of the subsidy to at-risk capacity offered to offshore input suppliers under the international agreement ($\tilde{\sigma}(\hat{V}_a, P - \kappa_{Off}, \alpha = 0)$) *remains unchanged* from the level of the subsidy to at-risk capacity offered to locally sourced input suppliers in the noncooperative equilibrium ($\hat{\sigma}(P)$). Nevertheless, while the international agreement does not alter the observed subsidy level to

at-risk capacity in this case, it is still the positive externality generated by such subsidies when offered to offshored input suppliers that drives the need for an international agreement to internalize these externalities: it is just that in the case identified in Lemma 3, these international externalities are off-equilibrium.

Finally, to confirm that there are model parameters under which the conditions of Lemma 3 hold, we fix κ_{Loc} and consider different values of κ_{Off} , and we define $\omega \equiv [\kappa_{Loc} - \kappa_{Off}]$ as the marginal cost savings that is associated with offshoring (for $\omega > 0$) or local sourcing (for $\omega < 0$) of inputs. It is easy to see that for ω negative and sufficiently large in magnitude, (i) $E[Cost_{Off}^N(\bar{V}_a^*(\omega))] - E[Cost_{Loc}^N(\hat{V}_a)] > 0$ is assured so that the noncooperative equilibrium features local sourcing, but (ii) Condition I will be violated so that no mutually beneficial international agreement on vaccine supply chains is possible. To see that it is also possible to find a range of ω for which Condition I will be satisfied, note that for $\omega = 0$ and fixing all other model parameters, we must have $E[Cost_{Off}^N(\bar{V}_a^*(\omega = 0))] - E[Cost_{Loc}^N(\hat{V}_a)] > 0$ while for $\bar{\omega}$ positive and sufficiently large we must have $E[Cost_{Off}^N(\bar{V}_a^*(\omega = \bar{\omega}))] - E[Cost_{Loc}^N(\hat{V}_a)] < 0$, so that, conditional on these other model parameters, there exists $\tilde{\omega} \in (0, \bar{\omega})$ such that $E[Cost_{Off}^N(\bar{V}_a^*(\omega = \tilde{\omega}))] - E[Cost_{Loc}^N(\hat{V}_a)] = 0$. But for $\tilde{\omega}$ (and a left neighborhood of $\tilde{\omega}$) Condition I is surely satisfied, because then

$$\begin{aligned} & E[Cost_{Off,Off}^C(\hat{V}_a(\tilde{\omega}), \hat{V}_a(\tilde{\omega}))] - E[Cost_{Loc}^N(\hat{V}_a)] \\ &= E[Cost_{Off,Off}^C(\hat{V}_a(\tilde{\omega}), \hat{V}_a(\tilde{\omega}))] - E[Cost_{Off}^N(\bar{V}_a^*(\tilde{\omega}))] \\ &< E[Cost_{Off,Off}^C(\hat{V}_a(\tilde{\omega}), \hat{V}_a(\tilde{\omega}))] - E[Cost_{Off,Off}^N(\bar{V}_a^*(\tilde{\omega}))] < 0 \end{aligned}$$

where the last inequality is an implication of Proposition 4.

We may therefore state:

Proposition 5 *If the noncooperative outcome features local sourcing of vaccine inputs and provided that $P < \frac{\delta h}{1-\alpha}$, then there may be no role for an international agreement on vaccine policy, but it is also possible that a role for an international agreement on vaccine policy exists. Moreover, when there does exist such a role, (i) the task of the agreement is to facilitate a switch to efficient offshoring of vaccine inputs, and (ii) while it is still the positive externality generated by subsidies to at-risk capacity investment when offered to offshored input suppliers that drives the need for the agreement, in this case the international externalities are off-equilibrium and the observed subsidies to at-risk capacity investment need not rise under the international agreement.*

6 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.

This paper presents a theoretical framework to help investigate that puzzle. It begins by providing a closed economy model to illustrate a first-best role for government policy to induce the socially optimal level of at-risk vaccine production capacity. Motivated by the industry’s production structure - which is characterized by considerable outsourcing and offshoring - it then investigates the role of trade. In the presence of production offshoring or imported inputs, a benevolent government may induce its firms to under-invest in at-risk capacity, leading to larger pandemic-related health costs for its economy than if its vaccine manufacturer had set up a purely domestic supply chain. Nevertheless, even in the presence of production offshoring, governments can induce the socially optimal level of capacity investment, but doing so requires international cooperation. The form of policy cooperation involves commitments to greater at-risk-capacity-enhancing subsidies for input-providers along the supply chain. And the role for international cooperation may extend as well to situations where absent that cooperation vaccine-producing firms would eschew offshoring and source their inputs locally, pointing to the possibility that off-equilibrium international externalities associated with subsidies can also underpin the need for international cooperation. Since the international externalities in our setup are all pecuniary by construction, the WTO is the natural institution to facilitate this novel form of international policy cooperation.

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, world leaders need to step in. So far, they have not. For all its successes, the race to vaccinate the world against Covid 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 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, 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 - 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?²⁴ We leave answers to these and other related questions to future research.

²⁴The difficulty of enforcing international commitments during a pandemic is emphasized, for example, in Staiger (forthcoming, chapter 12).

References

- [1] Agarwal, Ruchir and Gita Gopinath. 2021. “A Proposal to End the COVID-19 Pandemic,” *IMF Staff Discussion Note* 2021/04, May.
- [2] Ahuja, Amrita, Susan Athey, Arthur Baker, Eric Budish, Juan Camilo Castillo, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Jean Lee, Canice Prendergast, Christopher Snyder, Alex Tabarrok, Brandon Joel Tan, and Witold Wiecek. 2021a. “Preparing for a Pandemic: Accelerating Vaccine Availability,” *American Economic Association Papers and Proceedings* 111: 331-335.
- [3] Ahuja, Amrita, Susan Athey, Arthur Baker, Eric Budish, Juan Camilo Castillo, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Jean Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon Joel Tan and Witold Wiecek. 2021b. “Preparing for a Pandemic: Accelerating Vaccine Availability,” *American Economic Association Papers and Proceedings*: Online Appendix.
- [4] Antras, Pol and Robert W. Staiger. 2012. “Offshoring and the Role of Trade Agreements,” *American Economic Review* 102: 3140-83.
- [5] Athey, Susan, Juan Camilo Castillo, Esha Chaudhuri, Michael Kremer, Alexandre Simoes Gomes, Christopher M. Snyder. 2022. “Expanding Capacity for Vaccines Against Covid-19 and Future Pandemics: A Review of Economic Issues,” mimeo, January.
- [6] Bagnoli, Mark and Ted Bergstrom. 2005. “Log-concave Probability and its Applications,” *Economic Theory* 26: 445–469.
- [7] Bagwell, Kyle and Robert W. Staiger. 2001. “Strategic Trade, Competitive Industries and Agricultural Trade Disputes,” *Economics & Politics* 13: 113-128.
- [8] Bagwell, Kyle, and Robert W. Staiger. 2002. *The Economics of the World Trading System*. Cambridge, MA and London: MIT Press.
- [9] Bagwell, Kyle and Robert W. Staiger. 2006. “Will International Rules on Subsidies Disrupt the World Trading System?,” *American Economic Review* 96: 877-895.
- [10] Bagwell, Kyle, and Robert W. Staiger. 2012. “The Economics of Trade Agreements in the Linear Cournot Delocation Model.” *Journal of International Economics* 88 (1): 32–46.
- [11] Bown, Chad P. 2022a. “How COVID-19 medical supply shortages led to extraordinary trade and industrial policy,” *Asian Economic Policy Review* 17: 114-135.

- [12] Bown, Chad P. 2022b. “COVID-19 vaccine supply chains and the Defense Production Act,” mimeo, May.
- [13] Bown, Chad P. 2022c. “The WTO and vaccine supply chain resilience during a pandemic,” mimeo, May.
- [14] 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.
- [15] Bown, Chad P., and Thomas J. Bollyky. 2022. “How COVID-19 vaccine supply chains emerged in the midst of a pandemic,” *The World Economy* 45: 468-522.
- [16] Castillo, Juan Camilo, Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Greg Larson, Jean Lee, Canice Prendergast, Christopher Snyder, Alex Tabarrok, Brandon Joel Tan, and Witold Wiecek. 2021. “Market Design to Accelerate COVID-19 Vaccine Supply,” *Science* 371 (6534): 1107-1109.
- [17] 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.
- [18] Grossman, Gene M., Elhanan Helpman and Hugo Lhuillier. 2021. “Supply Chain Resilience: Should Policy Promote Diversification or Reshoring?,” mimeo, Princeton University, September 27.
- [19] Interlandi, Jeneen. 2021. “The World Is at War With Covid: Covid Is Winning,” *Opinion Piece, The New York Times*, September 21.
- [20] Snyder, Christopher M., Kendall Hoyt, Demitrios Gouglas, Thomas Johnston, and James Robinson. 2020a. “Designing Pull Funding for a COVID-19 Vaccine,” *Health Affairs* 39 (9): 1633-1642.
- [21] Snyder, Christopher M., Kendall Hoyt, Demitrios Gouglas, Thomas Johnston, and James Robinson. 2020. “Supplementary Appendix for ‘Designing Pull Funding for a COVID-19 Vaccine’,” mimeo.
- [22] Staiger, Robert W. Forthcoming. *A World Trading System for the Twenty-First Century*, MIT Press.
- [23] 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.
- [24] Sykes, Alan O. 2005. “Subsidies and Countervailing Measures,” *The World Trade Organization: Legal, Economic and Political Analysis*, vol. 2, ed. Arthur E. Appleton, Patrick F. J. Macrory, and Michael G. Plummer, 1682-706. Boston: Springer.

- [25] 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.
- [26] US Government Accountability Office (GAO). 2021. “COVID-19: Efforts to Increase Vaccine Availability and Perspectives on Initial Implementation,” Washington, DC: GAO-21-443, April.

Appendices

A series of appendices analyze a variety of alternative models to assess the robustness of the results derived in our Benchmark Model and emphasized in the main body of the paper. In Appendices 7 and 8, we focus on the case of local sourcing and show that Proposition 1 extends to a setting where demand for vaccine doses is not time-sensitive and where susceptibility to the virus is heterogeneous in the population. In Appendix 9 we confirm that all of the results of the Benchmark Model extend to a setting with a linear cost-of-capacity function, and there we also extend the analysis to the presence of incomplete information.

7 Potential Vaccine Demand Not Time-Sensitive

In the Benchmark Model developed in the main body of the paper, we assume that absent an effective vaccine the virus circulates in the population for T periods and then disappears. As we observed there, this assumption might be rationalized on the grounds that a highly contagious but benign variant of the virus emerges at the end of period T and becomes dominant, inoculating the entire population in period $T + 1$ without causing serious illness. Together with our assumption that one vaccine dose can protect one individual for one period against the virus, this meant that the faster the firm could deliver vaccine doses to the population, the more vaccine doses it could sell.

We now consider an alternative to these assumptions, which lies at the other extreme. We assume that one dose of the vaccine is sufficient to permanently inoculate an individual, and that the benign variant arrival date period T is sufficiently far out in the future that in effect, as long as the vaccine turns out to be safe and effective, the contagion ends only when everyone in the population has been vaccinated.²⁵ This implies that the firm will be able to sell L doses of the vaccine, regardless of the speed with which it delivers doses to the population. Reality probably lies somewhere in between these two extremes but, as we now demonstrate, there is a sense in which the design of optimal vaccine policy does not hinge on which extreme is closer to the truth. In particular, we show that the insights from Proposition 1 apply equally well in this setting.

Relative to the Benchmark Model, this change in assumptions alters the optimal at-risk capacity choice both for the government and for the firm. For the government, this is so because a single dose of the vaccine is sufficient to permanently inoculate an individual, and so a vaccine dose has greater social value in terms of preventing multiple periods of illness when it is administered earlier in the contagion period: in other words, the social value of a vaccine dose now declines through time as the world approaches period T . In particular, the social value of a vaccine dose administered in period $t \leq T$, in terms of the health costs avoided as a result of the administration of this dose net of the

²⁵The role of the benign variant arrival date is then simply to ensure that the cost to society of the virus remains finite even if the vaccine turns out not to be safe and effective.

marginal cost of producing the dose, is equal to $[1 + (T - t)] \times \delta h - \kappa$. For the firm, the optimal at-risk capacity choice is altered because at-risk capacity will no longer provide the firm with any advantage in terms of greater overall vaccine dose sales; what investment in at-risk capacity could do instead is provide the firm with greater revenues due to early sales of vaccine doses at higher prices, but only if the price per vaccine dose declines through time.

Notice that the social value of a dose delivered in period $t = M + 1 + m$, the period that capacity investments made in $t = M + 1$ could first be deployed in production, is equal to

$$[1 + (T - t)] \times \delta h - \kappa = (T - (M + m)) \times \delta h - \kappa \text{ for } t = M + 1 + m.$$

This means that the analog of Assumption I that we imposed in the Benchmark Model can be stated for the present model as

$$\max\left[\frac{c(L) + L\kappa}{(T - (M + m))L}, \frac{c'(L)}{s}\right] < \delta h. \quad (\text{Assumption AI})$$

Assumption AI delivers the same properties in the model we develop here as does Assumption I in the Benchmark Model: namely, we can cast the government's optimal choice of capacity investment as its choice of the constant flow level of at-risk investment v_a that will be maintained over the periods $t \in [1, M + 1 - m]$, or equivalently as simply its choice of the stock of at-risk capacity V_a at the end of period $M + 1 - m$.

Similarly, the analog of Assumption II that we imposed in the Benchmark Model can be stated for the present model as

$$\frac{c(L) + L\kappa}{(T - (M + m))L} < P(t = M + 1 + m), \quad \text{and} \quad \frac{c'(L)}{s} < -[P'(t)] \text{ for } t \in [1, T]. \quad (\text{Assumption AII})$$

Assumption AII delivers the same properties in the model we develop here as does Assumption II in the Benchmark Model: namely, we can cast the firm's optimal choice of capacity investment as its choice of the constant flow level of at-risk investment v_a that will be maintained over the periods $t \in [1, M + 1 - m]$, or equivalently as simply its choice of the stock of at-risk capacity V_a at the end of period $M + 1 - m$. Notice that the second condition in Assumption AII, which ensures that the maximum possible capacity-cost savings associated with a one-period delay in a small amount of capacity investment would not be worth it because this delay would cause the price at which the vaccine can be sold to fall by a sufficiently great amount, is guaranteed to be violated if the vaccine price is constant (or rising) through time. On the other hand, if the price is set equal to the social value of a dose so that $P(t) = [1 + (T - t)] \times \delta h - \kappa$, then $P'(t) = -\delta h$ and this second condition would amount to $\frac{c'(L)}{s} < \delta h$, the same as the second condition in Assumption AI.

With these assumptions, we can proceed to characterize the capacity investments chosen by the government/social planner and by the firm as we did in

the Benchmark Model. The expected overall cost to the country of the viral outbreak can be written as

$$E[Cost(V_a)] = C(V_a) + sc(L - (m + 1)V_a) + sL\kappa + (1 - s)LT\delta h + s[LM\delta h + \sum_{t=1}^m (L - tV_a) \times \delta h], \quad (A1)$$

where in writing (A1) we have implicitly assumed that the parameters s and m are sufficiently small so that the government's choice of \hat{V}_a satisfies $L > (m + 1)\hat{V}_a$.²⁶ The interpretation of (A1) is analogous to that of (3) for the Benchmark Model. Using (A1), the first-order condition that determines the socially optimal choice of at-risk capacity investment – that is, the at-risk capacity investment level \hat{V}_a that minimizes the expected overall cost to the country of the viral outbreak – is given by

$$\frac{\partial E[Cost(V_a)]}{\partial V_a} = C'(\hat{V}_a) - s[(m + 1) \times c'(L - (m + 1)\hat{V}_a) + \delta h \sum_{t=1}^m t] = 0, \quad (A2)$$

with the second-order condition again guaranteed by the convexity of the cost-of-capacity-installation $c(v)$ function.

We next turn to the capacity installation choices of the vaccine-producing firm. Again we take as exogenous the price that the firm receives for each dose of a vaccine that is deemed to be safe and effective, but we now allow this price to be a function of time, $P(t)$. The firm's expected profits can be written as

$$E[\pi(V_a)] = s \left[\left(\sum_{t=M+1}^{t=M+m} [V_a \times P(t)] + [L - mV_a] \times P(M + 1 + m) \right) - L\kappa \right] - [C(V_a) + sc(L - (m + 1)V_a)]. \quad (A3)$$

The interpretation of (A3) is analogous to that of (6) for the Benchmark Model. Using (A3), the first-order condition that determines the profit-maximizing choice of at-risk capacity investment for the firm, \tilde{V}_a , is given by

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = s \left[\sum_{t=M+1}^{t=M+m} P(t) - mP(M + m + 1) \right] - [C'(\tilde{V}_a) - s(m + 1) \times c'(L - (m + 1)\tilde{V}_a)] = 0, \quad (A4)$$

with the second-order condition again guaranteed by the convexity of the cost-of-capacity-installation function $c(v)$. Finally, for purposes of comparison with (A2), it is helpful to rewrite (A4) in the equivalent form

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = - \left\{ C'(\tilde{V}_a) - s \left[(m + 1) \times c'(L - (m + 1)\tilde{V}_a) + \sum_{t=M+1}^{t=M+m} P(t) - mP(M + m + 1) \right] \right\} = 0. \quad (A5)$$

²⁶As will become clear below, for any value of m , \hat{V}_a will approach zero as s approaches zero and this assumption must be satisfied.

We are now ready to compare the social and commercial incentives to install at-risk capacity in this modified model. We confirm for this setting the central message of Proposition 1 for the Benchmark Model, namely, that what is needed to align commercial and social incentives for investment in at-risk capacity is simply that the price per dose be set at the social value of the dose.

To establish this, we set the vaccine price equal to the social value of a vaccine dose. That is, suppose the government could commit to purchase vaccine doses from the firm at the price

$$\hat{P}(t) \equiv [1 + (T - t)] \times \delta h - \kappa \text{ for } t \leq T. \quad (\text{A6})$$

Facing the price schedule in (A6), the firm's first-order condition in (A5) becomes

$$\begin{aligned} \frac{\partial E[\pi(V_a)]}{\partial V_a} = & -\{C'(\tilde{V}_a) - s[(m + 1) \times c'(L - (m + 1)\tilde{V}_a) + \\ & \sum_{t=M+1}^{t=M+m} ([1 + (T - t)] \times \delta h - \kappa) - m([1 + T - M - m - 1] \times \delta h) - \kappa)\} = 0, \end{aligned}$$

which simplifies to

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = -\{C'(\tilde{V}_a) - s[(m + 1) \times c'(L - (m + 1)\tilde{V}_a) + \delta h \sum_{t=1}^m t]\} = 0. \quad (\text{A7})$$

A comparison of (A7) with (A2) confirms that the government can induce the firm to invest in the socially optimal level of at-risk capacity – that is the government can ensure that $\tilde{V}_a = \hat{V}_a$ – if it can commit to allow the firm to earn the social value of every vaccine dose that it sells. We summarize with:

Proposition A1 *In the model of this appendix, the timing of the firm's capacity investments do not impact the total number of vaccine doses it can sell. Nevertheless, if the price per vaccine dose is set at the social value of each dose, $\hat{P}(t) = [1 + (T - t)] \times \delta h - \kappa$, then commercial and social incentives for investment in at-risk capacity will be aligned, and $\tilde{V}_a = \hat{V}_a$.*

Together, Propositions 1 and A1 suggest that, if a government can commit to a price schedule for vaccine doses, then committing to set the price per dose at the social value of each dose would be a natural policy intervention for ensuring that socially optimal at-risk capacity investment is made by a vaccine-producing firm that locally sources its inputs. Notice that this policy intervention could be designed to take the form of an up-front commitment to pay the amount $\hat{P}(M + 1)L$ for delivery of L doses of a safe and effective vaccine at the beginning of period $t = M + 1$, combined with an appropriate “penalty” for late delivery (i.e., if fewer than L doses are delivered at the beginning of period $t = M + 1$, then the lower price $\hat{P}(M + 2)$ will be paid for each dose delivered at the beginning of period $t = M + 2$ up to the maximum total delivery of L doses, and so on

applying the price schedule $\hat{P}(t)$ for any deliveries that occur at even later dates $t = M + 3$ all the way until the final possible delivery at date T). By design, in equilibrium a series of late deliveries would occur and penalties would be paid, reflecting the vaccine delivery and declining price schedule described by Proposition A1. This is reminiscent of the bonus payments for early delivery of vaccine doses observed in some of the US contracts as described in section 2.

8 Heterogeneous Susceptibility

In this appendix we show that the main results from the Benchmark Model of section 3.1 and the model of appendix 7, as summarized in Propositions 1 and A1 respectively, are preserved when heterogeneity in susceptibility to the virus is introduced into the population.

To proceed, we introduce one new assumption: rather than assuming that in every period when the virus is circulating each unvaccinated individual will be infected by the virus with probability $\delta \in (0, 1]$, we now assume that the susceptibility to contracting the virus differs across the population. We introduce this feature by assuming that the country under consideration is populated with a measure L of citizens indexed by $z \in [0, L]$, and we assume that the probability that individual z will contract the virus in any period when individual z is unvaccinated is $\delta(z)$, where the function δ is strictly decreasing (by the convention of indexing citizens in order of decreasing susceptibility to the virus) and continuous (by assumption).²⁷ We can then define $\Delta(V)$ for $V \in [0, L]$ as the measure of citizens $z \in [0, V]$ that, if unvaccinated, would contract the virus in any period, where $\Delta(V)$ is given by

$$\Delta(V) \equiv \int_0^V \delta(z) dz \quad \text{for } V \in [0, L]. \quad (\text{A8})$$

This implies that, in the absence of a safe and effective vaccine the expected cost imposed on the country by the virus amounts to $\Delta(L)Th$ (rather than $LT\delta h$ as in the Benchmark Model and the model of appendix 7).

Observe that (A8) implies that $\Delta'(V) = \delta(V) > 0$ and $\Delta''(V) = \delta'(V) < 0$, so that $\Delta(V)$ is an increasing and concave function. This property ensures that the optimal allocation of V vaccine doses for $V < L$ in terms of maximizing the country's health benefits of the V doses is to allocate doses to the most susceptible citizens, i.e., to citizens $z \in [0, V]$. The introduction of heterogeneous susceptibility therefore provides an additional reason, absent from the Benchmark Model and the model of appendix 7, that capacity installed at an earlier date generates higher social value than additional capacity that is installed at a later date: the first installment of capacity will be used to produce vaccine doses for those who are most susceptible to the virus and who therefore provide society with the greatest health benefits from inoculation. Nevertheless, as we

²⁷We could alternatively assume that the severity of illness h associated with contracting the virus varies across the population, and derive identical results.

next establish, it is still the case in each model that setting the price per vaccine dose equal to the social value, appropriately defined, is sufficient to provide the firm with the incentive to choose the socially optimal level of at-risk investment.

8.1 Heterogeneous susceptibility in the Benchmark Model

We first consider the Benchmark Model, augmented with the assumption of heterogeneous susceptibility to the virus across the population as defined above. Our Assumption I and Assumption II are sufficient to deliver the same implications in the present model as they do in the Benchmark Model, provided that $\delta(L) \geq \bar{\delta}$ for some $\bar{\delta} \in (0, 1)$ and with $\bar{\delta}$ then playing the role of δ in the statement of Assumption I, which we assume is the case henceforth. With these assumptions imposed, and with all other assumptions of the Benchmark Model in place, it is straightforward following the steps in section 3.1 to derive the objectives of the government and the vaccine-producing firm.

For the government, under Assumption I and using

$$\int_{V_a}^L \delta(z) dz = \int_0^L \delta(z) dz - \int_0^{V_a} \delta(z) dz \equiv \Delta(L) - \Delta(V_a) \text{ for } V_a \leq L,$$

the expected overall cost to the country of the viral outbreak is now

$$E[\text{Cost}(V_a)] = C(V_a) + sc(L - V_a) + (1 - s)\Delta(L)Th + s[\Delta(L)Mh + (\Delta(L) - \Delta(V_a))mh],$$

and the associated first-order condition that defines \hat{V}_a is given by

$$\frac{\partial E[\text{Cost}(V_a)]}{\partial V_a} = C'(\hat{V}_a) - s[c'(L - \hat{V}_a) + m\delta(\hat{V}_a)h] = 0, \quad (\text{A9})$$

with the second-order condition guaranteed by the convexity of the cost-of-capacity-installation $c(v)$ function and the concavity of $\Delta(V)$. And for any fixed price per dose \hat{P} that satisfies Assumption II, the expected profits of the vaccine-producing firm are now

$$E[\pi(V_a)] = s[V_a(T - M) + (L - V_a) \times (T - (M + m))] \times \hat{P} - [C(V_a) + sc(L - V_a)],$$

and the associated first-order condition that defines \tilde{V}_a is given by

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = sm\hat{P} - [C'(\tilde{V}_a) - sc'(L - \tilde{V}_a)] = 0,$$

or, equivalently,

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = -\{C'(\tilde{V}_a) - s[c'(L - \tilde{V}_a) + m\hat{P}]\} = 0, \quad (\text{A10})$$

with the second-order condition guaranteed by the convexity of the cost-of-capacity-installation function $c(v)$.

It is now immediate from (A9) and (A10) that, if the price of a vaccine dose is set equal to the social value of the last dose produced under the optimal level of at-risk capacity, that is, if

$$\hat{P} \equiv \delta(\hat{V}_a)h, \quad (\text{A11})$$

then $\tilde{V}_a = \hat{V}_a$. Hence, relative to the result reported in Proposition 1, the only difference here is that the social value of the marginal vaccine dose is no longer a constant δh but is rather a decreasing function $\delta(V)h$ of the number of doses that can be produced, V , and the price per dose must be set at the social value of the last dose that would be produced under the optimal level of at-risk capacity \hat{V}_a .

8.2 Heterogeneous susceptibility in the model of appendix 7

We next consider the model of appendix 7, augmented with the assumption of heterogeneous susceptibility to the virus across the population as defined above. Here our Assumption AI and Assumption AII are sufficient to deliver the same implications in the present model as they do in the model of appendix 7, provided that $\delta(L) \geq \bar{\delta}$ for some $\bar{\delta} \in (0, 1)$ and with $\bar{\delta}$ then playing the role of δ in the statement of Assumption AI, which we assume is the case henceforth. With these assumptions imposed, and with all other assumptions of the model of appendix 7 in place, it is straightforward to derive the objectives of the government and the vaccine-producing firm.

For the government, and using

$$\int_{tV_a}^L \delta(z)dz = \int_0^L \delta(z)dz - \int_0^{tV_a} \delta(z)dz \equiv \Delta(L) - \Delta(tV_a) \quad \text{for } tV_a \leq L,$$

the expected overall cost to the country of the viral outbreak is now

$$\begin{aligned} E[\text{Cost}(V_a)] &= C(V_a) + sc(L - (m+1)V_a) \\ &\quad + (1-s)\Delta(L)Th + s[\Delta(L)Mh + \sum_{t=1}^m (\Delta(L) - \Delta(tV_a))h], \end{aligned}$$

and the associated first-order condition that defines \hat{V}_a is given by

$$\frac{\partial E[\text{Cost}(V_a)]}{\partial V_a} = C'(\hat{V}_a) - s[(m+1) \times c'(L - (m+1)\hat{V}_a) + \sum_{t=1}^m t\delta(t\hat{V}_a)s] = 0, \quad (\text{A12})$$

with the second-order condition guaranteed by the convexity of the cost-of-capacity-installation $c(v)$ function and the concavity of $\Delta(V)$. Similarly, the expected profits of the vaccine-producing firm facing a per-dose price of $P(t)$

are now

$$E[\pi(V_a)] = s \left[\left(\sum_{t=M+1}^{t=M+m} [V_a \times P(t)] + [L - mV_a] \times P(M + m + 1) \right) - L\kappa \right] - [C'(V_a) + sc(L - (m + 1)V_a)].$$

and the associated first-order condition that defines \tilde{V}_a is given by

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = s \left[\sum_{t=M+1}^{t=M+m} P(t) - mP(M + m + 1) \right] - [C'(\tilde{V}_a) - s(m+1) \times c'(L - (m+1)\tilde{V}_a)] = 0,$$

or equivalently

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = - \left\{ C'(\tilde{V}_a) - s \left[(m + 1) \times c'(L - (m + 1)\tilde{V}_a) + \sum_{t=M+1}^{t=M+m} P(t) - mP(M + m + 1) \right] \right\} = 0, \quad (\text{A13})$$

with the second-order condition guaranteed by the convexity of the cost-of-capacity-installation function $c(v)$.

As in the model of appendix 7, it is immediate from a comparison of (A12) and (A13) that if the price per vaccine dose is fixed at its social value, where the social value of a vaccine dose in period t is evaluated at the socially optimal level of at-risk capacity and is defined for period t as the social value of the last dose produced in period t given the socially optimal level of at-risk capacity, then we would have

$$\hat{P}(t) \equiv [1 + (T - t)] \times \delta((t - M)\hat{V}_a)h - \kappa \quad \text{for } t \in [M + 1, M + m], \quad (\text{A14})$$

with $\hat{P}(M + m + 1)$ then set at

$$\hat{P}(M + m + 1) \equiv \frac{1}{m} \sum_{\iota=1}^m [\hat{P}(M + \iota) - \iota\delta(\iota\hat{V}_a)h] - \kappa, \quad (\text{A15})$$

and facing the price schedule defined by (A14)-(A15), the firm's first-order condition in (A13) simplifies to

$$\frac{\partial E[\pi(V_a)]}{\partial V_a} = - \{ C'(\tilde{V}_a) - s[(m + 1) \times c'(L - (m + 1)\tilde{V}_a) + \sum_{t=1}^m t\delta(t\hat{V}_a)h] \} = 0, \quad (\text{A16})$$

which together with (A12) implies $\tilde{V}_a = \hat{V}_a$.

9 Linear Costs and Incomplete Information

This appendix provides a fresh analysis of many of the major results of the paper but in a model with a simpler form for capacity costs, assuming linear

rather than convex. This simpler functional form allows us to extend the analysis in several directions. We endogenize the prices in procurement contracts offered by the government. We extend the model to allow the firm to have private information about its productivity. This extension has a major advantage. Without private information, it turns out that the government can obtain the second best if allowed sufficiently rich procurement contracts. With private firm information, we show here that the government cannot obtain the second best under private information. We show that the best the government can do with optimal contracts involves more distortion dealing with a foreign input supplier than domestic. This provides a role for international cooperation that is robust regardless of how rich is the contracting space.

Some of the longer proofs are relegated to the end of the appendix.

9.1 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 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. This 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 one panel of Figure A1. In this case, the firm produces the vaccine with domestically sourced inputs produced completely within the firm bound-

aries. Later, we will analyze cases, shown in 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 in 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 after the completion of clinical trials, 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, reflective of this 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.


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.²⁸ 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. 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 the t_a , while clinical trials are still underway, referred to as “at-risk” investment. At-risk investment trades off of wasted capacity

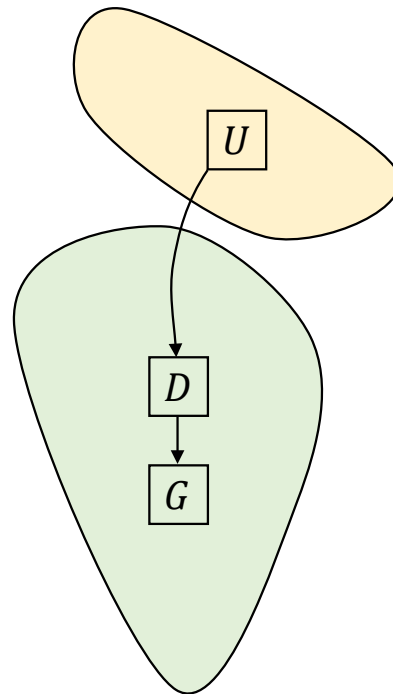
²⁸An equivalent alternative assumption is that capacity is fungible but only imperfectly so, with k capturing the value lost in repurposing.

Figure A1: Organization and Trade Structures Analyzed


Unintegrated domestic supply



Unintegrated foreign supply



Integrated domestic supply



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 installation process is underway, we abstract from any option value from abandoning it; the k capacity cost per unit of the capacity 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

$$t_\ell \leq t_a < t_a + t_\ell < t_e. \quad (\text{A17})$$

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 has 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:


$$1 - s > \frac{t_\ell}{t_e - t_a}. \quad (\text{A18})$$

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, the capacity strategy that is both weakly profitable and weakly socially efficient is to undertake each investment all together as early as possible. Installation of X_r should start in period 1, and installation of X_n should start 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 both a theoretically interesting strategy and practically

Figure A2: Timeline for Model



relevant in the Covid-19 pandemic; but to simplify the analysis, 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 A2. 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 A2). 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.

In the subsequent analysis, we will see that various organization and trade structures generate the same equilibria for some parameter vectors, only revealing distinctions for other vectors. Similarly, equilibria “tie” the first best for some parameter vectors, only revealing inefficiencies for other vectors. To

smooth out comparisons, we will introduce the device of a distribution of parameter vectors from which the one under consideration is drawn. This device will allow us to conduct the analysis for a single, representative parameter vector while still allowing us to make coherent statements comparing structures such as at-risk investment being more likely or welfare higher on average in one structure than another. Capacity cost k turns out to span the relevant outcome space, so for concreteness we will specify just that parameter being a random variable drawn from a distribution having probability density function (pdf) $f(k)$ and cumulative distribution function (cdf) $F(k)$ on support $[0, K]$, where K can be infinite. Assume $F(k)$ is logconcave, a property exhibited by most commonly used distributions in practice (Bagnoli and Bergstrom 2005). We will initially assume the draw of k is publicly observed ex ante, so the game remains one of complete information.

9.2 Integrated Domestic Supply

In this section, we analyze the case of integrated domestic supply, adhering to the model introduced in the previous section, corresponding to the structure depicted in the lower left panel of Figure A1. We begin by characterizing the first best, maximizing domestic social welfare. We then solve for the firm's equilibrium capacity and output levels for an arbitrary per-dose price offered by the government. Finally, we solve for the optimal procurement policy for a government that faces a social cost of public funds or that discounts producer relative to consumer surplus.

9.2.1 First Best

To characterize the first best, we will solve for the optimum chosen by a social planner 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 planner controls all firm operations, including the firm's capacity decisions, which, as argued in Section 9.1, can be reduced to two choices: the amount of at-risk capacity X_r and not-at-risk capacity X_n . The planner 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 9.1, 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 $t_a + t_\ell$ to t_e using combined capacity $X_r + X_n$.

Expected consumer harm over the pandemic is given by

$$\sum_{t=1}^{t_a-1} \beta h + \sum_{t=t_a}^{t_a+t_\ell-1} \beta h(1 - s\theta Q_1) + \sum_{t=t_a+t_\ell}^{t_e-1} \beta h(1 - s\theta Q_2) \quad (\text{A19})$$

$$= \beta h t_e - s\beta h \theta Q_1 t_\ell - s\beta h \theta Q_2 (t_e - t_a - t_\ell). \quad (\text{A20})$$

The first term in equation (A20) reflects the disease harms experienced over the pandemic in the absence of a vaccine, 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.

Expected total production cost over the pandemic is given by

$$\sum_{t=t_a}^{t_a+t_\ell-1} scQ_1 + \sum_{t=t_a+t_\ell}^{t_e-1} scQ_2 = scQ_1t_\ell + scQ_2(t_e - t_a - t_\ell). \quad (\text{A21})$$

Production costs are only expended if the vaccine is approved, so all terms in (A21) are scaled by the probability of success s .

Expected total capacity cost is given by

$$kX_r + skX_n. \quad (\text{A22})$$

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 minimizes three sources of cost in (A20), (A21), and (A22). Adding the three equations, rearranging, and recasting the minimization of costs as the equivalent maximization of its negative (to facilitate later comparison to the firm's profit-maximization problem), the first-best problem can be expressed as

$$\max_{X_r, X_n, Q_1, Q_2 \geq 0} [-\beta ht_e + sQ_1(\theta\beta h - c)t_\ell + sQ_2(\theta\beta h - c)(t_e - t_a - t_\ell) - kX_r - skX_n], \quad (\text{A23})$$

subject to

$$Q_1, Q_2 \leq 1 \quad (\text{A24})$$

$$Q_1 \leq X_r \quad (\text{A25})$$

$$Q_2 \leq X_r + X_n. \quad (\text{A26})$$

Constraint (A24) ensures that there is no benefit from vaccinating more than the total population mass, normalized to 1. Constraints (A25) and (A26) ensure output cannot exceed available capacity in the relevant periods.

The maximization problem in equations (A23)–(A26) 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 solely with at-risk capacity, or if not, solely 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 two stars added as subscripts indicate first-best values of endogenous variables.

Lemma A1 *The maximization problem in (A23)–(A26) has the trivial solution $X_r^* = X_n^* = Q_1^* = Q_2^* = 0$ if and only if $k \geq \bar{k}^*$, where*

$$\bar{k}^* = (\theta\beta h - c)(t_e - t_a - t_\ell). \quad (\text{A27})$$

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 at the end of this appendix, is based on the Kuhn-Tucker conditions from the linear program (A23)–(A26).

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 (A23), the planner's expected surplus from fully investing at risk is

$$-\beta h t_e + s(\theta\beta h - c)(t_e - t_a) - k. \quad (\text{A28})$$

Substituting $X_r^* = Q_1^* = 0$ and $X_n^* = Q_2^* = 1$ into (A23), the planner's expected surplus from fully investing not at risk is

$$-\beta h t_e + s(\theta\beta h - c)(t_e - t_a - t_\ell) - s k. \quad (\text{A29})$$

Setting (A28) equal to (A29) and rearranging shows that the planner is indifferent between the two investment strategies if and only if

$$s(\theta\beta h - c)t_\ell = (1 - s)k. \quad (\text{A30})$$

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 (A30) reflects 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.

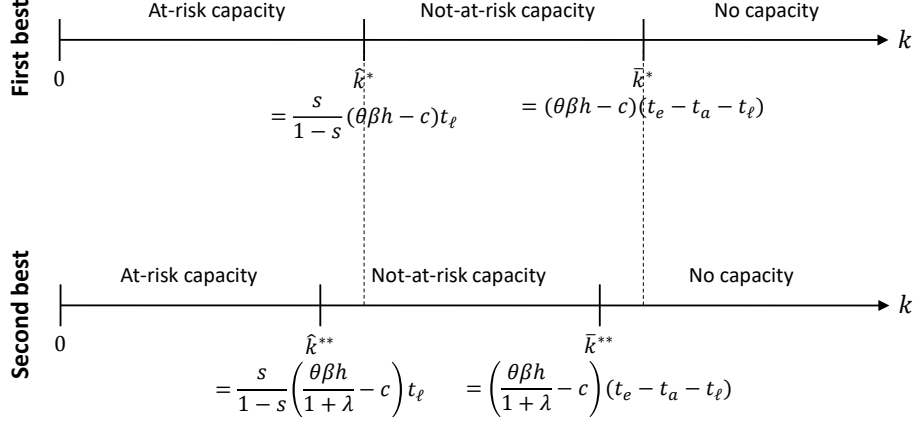
Rearranging equation (A30) gives a threshold capacity cost in the first best

$$\hat{k}^* = \frac{s}{1 - s}(\theta\beta h - c)t_\ell, \quad (\text{A31})$$

such that the planner strictly prefers at-risk investment if and only if $k < \hat{k}^*$. The planner finds at-risk investment more attractive the higher the probability of success s , the higher is efficacy θ , the more infectious and harmful the disease β and h , the longer the lag in capacity installation t_ℓ , and the lower is production cost c .

Figure A3 provides a schematic diagram of first-best capacity investment as k varies. For $k \in [0, \hat{k}^*)$, the first-best involves investing in enough at-risk capacity to serve the population. For $k \in (\hat{k}^*, \bar{k}^*)$, the first best involves

Figure A3: Optimal Capacity Investments as Functions of k



investing in enough not-at-risk capacity to serve the population. For $k > \bar{k}^*$, the first best involves no investment. The interval of at-risk investing has positive measure if and only if

$$\theta\beta h > c, \quad (\text{A32})$$

ensuring that a dose's social value exceeds its production cost, a minimal condition to justify any capacity investment.²⁹ The interval of not-at-risk investing is nonempty if and only if, in addition to condition (A32), condition (A18) holds.

9.2.2 Procurement from Nationalized Firm

We next turn to the optimal vaccine procurement policy for the government. Assume the government maximizes domestic social welfare, the sum of the surpluses of domestic consumers and producers. Government expenditures must be financed 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.

We start with the case of a nationalized firm, allowing the government control over all the firm's operations including capacity and output decisions. The government's objective function is the same as the planner's in the first-best problem (A23) 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 sometimes refer to this outcome as the second best: "best" in the sense that the principal has full control of the agent's operations, just as the planner in the first best did; only "second" best, however, because

²⁹We will not introduce (A32) as a maintained assumption here, waiting to introduce a slightly stronger version of the condition (integrating consideration of the social cost of public funds, to be defined) in the next subsection.

of the distortion involved in making transfers. Whereas first-best values of variables were distinguished by one star, second-best values will be distinguished 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 A1 holds, except that the cutoff capacity cost above which there is no investment is now given by

$$(1 + \lambda)\bar{k}^{**} = [\theta\beta h - (1 + \lambda)c](t_e - t_a - t_\ell). \quad (\text{A33})$$

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

$$\bar{k}^{**} = \left(\frac{\theta\beta h}{1 + \lambda} - c \right) (t_e - t_a - t_\ell). \quad (\text{A34})$$

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

$$\hat{k}^{**} = \frac{s}{1 - s} \left(\frac{\theta\beta h}{1 + \lambda} - c \right) t_\ell. \quad (\text{A35})$$

For $k < \hat{k}^{**}$, the government orders full investment at risk, providing surplus

$$-\beta ht_e + s[\theta\beta h - (1 + \lambda)c](t_e - t_a) - (1 + \lambda)k. \quad (\text{A36})$$

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

$$-\beta ht_e + s[\theta\beta h - (1 + \lambda)c](t_e - t_a - t_\ell) - s(1 + \lambda)k. \quad (\text{A37})$$

As expected, these expressions for second-best surplus are identical to the first best in (A28)–(A29) except that all cost parameters have been scaled up by $1 + \lambda$. It is immediate that the first best is recovered from the second best in the limit as $\lambda \downarrow 0$.

Figure A3 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}^{**} down 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

$$\theta\beta h > (1 + \lambda)c. \quad (\text{A38})$$

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 (A38) throughout the remainder of the paper to avoid a taxonomy of trivial cases.

9.2.3 Procuring from Privatized Firm Via Linear Price

We next study the optimal government policy toward a privatized firm. In this structure, the firm operates independently of the government, undertaking capacity and output decisions to maximize profit. Even though the government does not control the operations of the firm, we continue to assume that the government internalizes its profits. Still, the government seeks to limit transfers to the firm because the social cost of public funds generates deadweight loss λ for every dollar transferred.

In principle, the government and firm might bargain over the terms of the bilateral deal. To simplify the analysis and to give this structure the best chance of generating a socially efficient outcome, we will have the government move first, making a take-it-or-leave-it offer. For now, we restrict the government to offering a simple linear price p , as observed in many bilateral deals during the Covid-19 pandemic.

These assumptions set up a sequential game of perfect information between the government and the firm. The government moves first, offering price p . The firm moves second, basing its capacity and output decisions on p . We solve for the subgame perfect equilibrium using backward induction, starting by determining the firm's best response to any p .

Firm's Best Response As was the case with the social planner, the firm has two relevant capacity choices, X_r and X_n , and two relevant output choices, Q_1 and Q_2 . The firm chooses these variables to maximize its profit

$$\max_{X_r, X_n, Q_1, Q_2 \geq 0} [sQ_1(p-c)t_\ell + sQ_2(p-c)(t_e - t_a - t_\ell) - kX_r - skX_n], \quad (\text{A39})$$

subject to the same constraints, (A24)–(A26), faced by the planner.

One can show that the characterization of the planner's optimum in Lemma A1 as one of three corner solutions also applies to the firm's profit-maximization problem. The proof, based on Kuhn-Tucker conditions, uses nearly identical logic as the proof of Lemma A1 and so is omitted. Determining the firm's best response thus boils down to a comparison of the alternative of investing in population-level capacity at risk, yielding expected profit

$$\pi_r = s(p-c)(t_e - t_a) - k, \quad (\text{A40})$$

versus investing in population-level capacity not-at-risk, yielding expected profit

$$\pi_n = s(p-c)(t_e - t_a - t_\ell) - sk, \quad (\text{A41})$$

versus not investing at all, yielding zero profit.

The firm's profit from investing at risk (A40) weakly exceeds its profit from investing not at risk (A41) if

$$p \geq c + \frac{(1-s)k}{st_\ell}. \quad (\text{A42})$$

One can show that (A42) ensures that (A40) is positive when assumption (A18) holds. Thus, the firm is willing to invest at risk when (A42) holds. If the price is in the intermediate range

$$p \in \left[c + \frac{k}{t_e - t_a - t_\ell}, c + \frac{(1-s)k}{st_\ell} \right], \quad (\text{A43})$$

then the firm is willing to invest not at risk instead of investing at risk or not investing at all. One can show that assumption (A18) ensures that the interval in (A43) has positive measure. If the price is below the interval in (A43), then the firm prefers not to invest.

Government Price Setting Fold the game back to the government's first-stage choice of p . By varying p over the intervals defined in (A42) and (A43), the government can induce one of three firm strategies. First, suppose the government wishes to incentivize at-risk investment by setting a p satisfying (A42). The firm installs capacity $X_r^* = 1$ and $X_n^* = 0$, producing $Q_1^* = Q_2^* = 1$ each period from t_a on, providing the government with surplus

$$-\beta ht_e + s[\theta\beta h - (1 + \lambda)p](t_e - t_a) + \pi_r. \quad (\text{A44})$$

The first two terms reflect health losses and distortionary expenditures paying for the vaccine. (Both can be appropriately categorized as aspects of consumer surplus supposing that consumers pay for expenditures through taxes.) The last term reflects domestic firm's profit, which the government internalizes. Substituting for π_r from (A40) and simplifying, (A44) can be rewritten as

$$-\beta ht_e + s[\theta\beta h - c - \lambda p](t_e - t_a) - k. \quad (\text{A45})$$

This objective is decreasing in p since transfers involve a distortion due to the social cost of public funds. Thus, the optimal price inducing at-risk investing is the lowest price satisfying (A42), i.e.,

$$p = c + \frac{(1-s)k}{st_\ell}. \quad (\text{A46})$$

Substituting (A46) into (A45) and rearranging yields government surplus

$$-\beta ht_e + s[\theta\beta h - (1 + \lambda)c](t_e - t_a) - (1 + \lambda)k - M(k, \lambda), \quad (\text{A47})$$

defining

$$M(k, \lambda) = \frac{\lambda k}{t_\ell} [(1-s)(t_e - t_a) - t_\ell]. \quad (\text{A48})$$

For all $\lambda > 0$ and $k > 0$, assumption (A18) ensures $M(k, \lambda) > 0$. Since $\partial M(k, \lambda)/\partial \lambda = M(k, \lambda)/\lambda$, assumption (A18) ensures $\partial M(k, \lambda)/\partial \lambda > 0$ as well.

Comparing (A47) to (A36), we see that the government obtains less surplus when it inducing at-risk investment from a privatized firm than from a nationalized firm. The surplus difference is equal to $M(k, \lambda)$, reflecting the government's loss due to the moral-hazard problem, requiring it to pay a premium to induce at-risk investment. The price premium would net out of social welfare if $\lambda = 0$ and there were no social cost of public funds but is socially costly when $\lambda > 0$.

Next, suppose instead that the government wishes to incentivize not-at-risk investment by setting a p satisfying (A43). The firm installs capacity $X_r^* = 0$ and $X_n^* = 1$, producing $Q_1^* = 0$ in the period during which capacity is ramped up and $Q_2^* = 1$ each period from $t_a + t_\ell$ on, providing the government with surplus

$$-\beta h t_e + s[\theta \beta h - (1 + \lambda)p](t_e - t_a - t_\ell) + \pi_n. \quad (\text{A49})$$

As before, the first two terms reflect health losses and distortionary expenditures to pay for the vaccine, and the last term represents the domestic firm's profit, which the government internalizes. Substituting for π_n from (A41) and simplifying, (A49) can be written

$$-\beta h t_e + s[\theta \beta h - c - \lambda p](t_e - t_a - t_\ell) - s k. \quad (\text{A50})$$

This objective is decreasing in p , so is maximized subject to (A43) by setting a price equal to the lower bound of the interval:

$$p = c + \frac{k}{t_e - t_a - t_\ell}. \quad (\text{A51})$$

Substituting (A51) into (A50) and rearranging yields the same expression for government surplus as it obtains in the second best, given in equation (A37).

Equating the government's continuation payoff from inducing at-risk investment (A47) to its continuation payoff from inducing not-at-risk investment (A37) and solving for k yields the threshold capacity cost

$$\hat{k}^p = \hat{k}^{**} - \frac{M(\hat{k}^p, \lambda)}{(1 - s)(1 + \lambda)}, \quad (\text{A52})$$

where the superscript indicates an optimal variable in the case of a privatized firm offered a price contract. The equilibrium involves at-risk investment if $k < \hat{k}^p$ and not-at-risk investment if $k \in (\hat{k}^p, \bar{k}^{**})$.

Comparing (A35) and (A52), we see that \hat{k}^p and \hat{k}^{**} differ only by the last term in (A52), which is proportional to $M(\hat{k}^p, \lambda)$. Since $M(\hat{k}^p, \lambda) > 0$, we have $\hat{k}^p < \hat{k}^{**}$, implying that there is less at-risk investment in equilibrium than in the second best. This establishes the following proposition.

Proposition A2 *Under maintained assumptions, $\hat{k}^p < \hat{k}^{**}$, implying the government induces less at-risk investment on average across the distribution of k*

when procuring vaccine via a linear-price contract from a privatized firm compared to the second best in which the government controls the operations of a nationalized firm.

Counterfactual Control over Investment Timing As a counterfactual exercise to build intuition, we will imagine that the the government has one additional lever of control over the privatized firm, the power to control the timing of investment, perhaps choosing to ban at-risk investment or to ban not-at-risk investment. Even with this additional authority, the government still has less authority than it has over a nationalized firm. There, the government can control capacity and output levels directly at extensive and intensive margins. In the counterfactual exercise here, the government has no direct authority over intensive-margin decisions and cannot force the firm to invest. It can only incentivize those decisions indirectly via the price, with one additional lever of control over when capacity can be installed. That limited additional authority to control investment timing is all the government needs to attain the second best.

We propose that the following government policy will deliver the second best. If the conditions are such that the second best involves at-risk investment, i.e., $k < \hat{k}^{**}$, the government bans not-at-risk investment and offers the firm a price allowing it to just break even under at-risk investment,

$$p = c + \frac{k}{s(t_e - t_a)}, \quad (\text{A53})$$

found by setting at-risk profit (A40) equal to 0. Similarly, if the conditions are such that the second best involves not-at-risk investment, i.e., $k \in (\hat{k}^{**}, \bar{k}^{**})$, the government bans at-risk investment and offers the firm a price allowing it to just break even under not-at-risk investment. This price was computed above in equation (A51).

To verify that the proposed policy delivers the second best, first suppose $k < \hat{k}^{**}$. The firm breaks even at the offered price (A53) so has no strict incentive to deviate to not investing. Given it is forbidden from investing not-at-risk, it invests at risk in equilibrium. Substituting price (A53) into the relevant surplus function for the government (A45) yields the expression for second-best government surplus in (A36). Supposing, next, that $k \in (\hat{k}^{**}, \bar{k}^{**})$, similar arguments can be used to rule out deviations by the firm from the desired not-at-risk investment strategy and to show that the government earns the second-best surplus.

The proposed policy has the government banning at-risk investment $k \in (\hat{k}^{**}, \bar{k}^{**})$. In fact, this ban is unnecessary. The discussion following equation (A43) showed that the firm earns negative profit from investing at risk if paid the break-even price for investing not at risk.

Proposition A3 *Suppose that the government, in addition to offering a linear-price contract to a privatized firm, also has the authority to forbid not-at-risk*

investment. This limited amount of additional authority is sufficient to allow the government to attain the second-best surplus.

In the language of contract theory, we can think of incentive-compatibility constraints ensuring that the firm undertakes the desired form of investment (at-risk or not-at-risk) and individual-rationality constraints ensuring that the firm at least breaks even under that form of investment. Incentive compatibility only binds for at-risk investment. For not-at-risk investment, if individual rationality binds, incentive compatibility is slack.

The analysis in this subsection has pinpointed the source of moral hazard with a privatized firm. The linearity of costs eliminates the problems of undercapacity and undersupply; the government can obtain the second best without controlling those as shown in Proposition A3. The moral-hazard problem regards the timing of capacity installation: the firm has too little private incentive for at-risk investment. If the government could (counterfactually) ban later investment, it could eliminate moral hazard entirely and obtain the second best.

9.2.4 Procuring from Privatized Firm Via Richer Contracts

The government can improve upon the surplus obtained from its deal with a privatized firm by moving to richer alternatives to the linear-price contract.

Capacity Subsidy One alternative is to combine a linear price with a subsidy for at-risk capacity investment. The government can observe when the investment is made and agree only to subsidize capacity installed before period t_a . Let σ denote the per-unit subsidy. Assume σ is paid at the start of period t_a for each unit of capacity completed by then whether or not the vaccine is successful. We claim that a contract of this form can attain the second-best surplus. In particular, the per-dose price can be set to cover production cost, $p^r = c$, and the at-risk capacity subsidy can be set to cover the per-unit capacity cost, $\sigma^r = k$, where the superscript indicates the optimal provisions of a richer alternative to a linear-price contract. By construction, this contract allows the firm to just break even if it invests at risk. It would earn negative profit from investing not at risk because capacity investment undertaken later would not be covered. Thus, the contract induces the firm to invest at risk.

It is immediate that the government can attain the second-best surplus with this contract for all k such that the second best involves at-risk investment, i.e., $k < \hat{k}^{**}$. The contract induces the same investment strategy as in the second best in that case. The contract only transfers enough to the firm to allow it to break even in expectation, so economizes on the social cost of public funds to the extent possible.

For $k \in (\hat{k}^{**}, \bar{k}^{**})$, the second best involves investing not at risk. The contract with capacity subsidy $\sigma^r = k$ and $p^r = c$ works for these parameters as well. But the richer contracts are not needed for these parameters. We showed in the previous subsection that there is no moral-hazard problem in inducing not-at-risk investment with a linear price. The government can attain the second

best offering a simple linear-price contract to the privatized firm. No additional capacity subsidy is needed. In sum, for all $k > 0$, some combination of a capacity subsidy and a linear price allows the government to obtain the second-best surplus.

Bonus Price for Early Doses Another alternative contract that allows the government to attain the second-best surplus firm involves two prices, a bonus price p_1 for doses sold in periods t_a through $t_a + t_\ell - 1$, which could only have been produced with at-risk capacity, and a second, potentially lower price p_2 covering production cost for doses sold after that.

One can check that setting $p_1^r = c + k/st_\ell$ and allows the firm to just break even if it invests at risk. If the firm invests not at risk, it earns negative profit under this two-price contract because it would not earn any bonus. Only its production costs, not its capacity costs, would then be covered and no bon. The contract with bonus for early delivery thus works to induce at-risk investment, and provides the government with the second-best surplus. Again, the firm can use the bonus-price strategy if $k < \hat{k}^{**}$, and default back to a single linear-price contract that induces investment not at risk if $k \in (\hat{k}^{**}, \bar{k}^{**})$. A single-price contract can be viewed as a trivial special case of a two-price contract. In this way, the government can attain the second-best surplus for all $k > 0$.

Comparing Alternative Contracts The alternative contracts are more closely aligned than might first appear, as revealed by a comparison of the total government expenditure entailed beyond the constant linear price. The government's total expenditure on the capacity subsidy (at least in states in which it is optimal to induce at-risk investment) is $\sigma^r = k$. We envisioned this amount being paid in a lump sum in period t^a whether or not the vaccine is approved. Equivalently, the payment could be prorated over the t_ℓ periods starting from t_a during which vaccines were only available due to at-risk investment. The prorated payment, again made whether or not the vaccine is approved, would be k/t_ℓ per period. Equivalently, the prorated payments could be promised only conditional on success. To keep the expected payment constant, the conditional payment would need to be scaled up by the reciprocal of the probability of success s , in which case the per-period capacity payment would become k/st_ℓ . But note this is exactly the markup over cost paid for early doses in the bonus-price scheme.

In our simple model, therefore, the capacity subsidy is isomorphic to the bonus-price scheme. They both provide the same extra incentive payment in expectation. The capacity payment is an example of what the innovation literature calls “push funding,” covering investment cost unconditional on the project's success. While push funding is sometimes criticized for dulling incentives, here push funding is only provided for the type of investment desired; strong, direct incentives for at-risk investment are provided. The bonus-price scheme is an example of what the innovation literature calls “pull funding,” paying the bonus conditional on success. Conditioning on success is not the source of incentives for at-risk investing here. Rather, it is that the bonus is

only paid for early doses, providing incentives for the only type of investment that can produce them. Since the prorating of the payment and the conditioning on success are irrelevant to at-risk investment incentives, we can specify the formulas can convert from one to the other without changing the power of the incentive scheme.

9.3 Unintegrated Input Supply

In this section, 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.

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 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.

9.3.1 Domestic Input Supplier

It is obvious that a move from an integrated to an unintegrated supplier does not change the first-best outcome if the planner 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 planner control. Similarly, the move from integrated to unintegrated supplier does not change the outcome with nationalized firms 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.

9.3.2 Foreign Input Supplier

In this subsection, we continue to suppose that D is a domestic, downstream firm but now move U 's location to a foreign country. The structure is depicted in the right panel of Figure A1.

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 an indirect effect of the combined organization and location change. The combined change will change the government 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 provide the government with stronger incentives to economize on transfers, leading it to induce less at-risk investment, reducing social welfare.

Suppose for now that the domestic government procures vaccine from D using a linear-price contract. The analysis of a linear contract offered to an integrated, domestic firm provided in Section 9.2.3 carries over with little modification here. Firms best respond to a given price p with the same investment and output decisions as before. Efficient bargaining leads the unintegrated firm to operate as joint-profit maximizers, and their location in different countries is irrelevant to their interaction.

The expressions for government surplus change. Government surplus from prices inducing at-risk investment and not-at-risk investment, respectively, given by equations (A44) and (A49) when the government procures from an integrated domestic firm become

$$-\beta ht_e + s[\theta\beta h - (1 + \lambda)p](t_e - t_a) + (1 - \phi)\pi_r \quad (\text{A54})$$

$$-\beta ht_e + s[\theta\beta h - (1 + \lambda)p](t_e - t_a - t_\ell) + (1 - \phi)\pi_n. \quad (\text{A55})$$

These are identical to their analogues except that the profit terms are scaled by $1 - \phi$ to reflect the fact that the government internalizes domestic but not foreign firm profit.

This small change to the specification of government surplus does not affect the logic of the analysis in from Section 9.2.3. Sparing the details, one can show that the threshold capacity cost below which the government induces at-risk investment and above which it induces not-at-risk investment is given by

$$\hat{k}^f = \hat{k}^{**} - \frac{M(\hat{k}^f, \phi + \lambda)}{(1 - s)(1 + \lambda)}, \quad (\text{A56})$$

This is identical to the analogous threshold when the government procures from an integrated domestic supplier except that the second argument of the moral-hazard term $M(k, \phi + \lambda)$ reflects a greater disutility from making transfers to the firm. In addition to the social cost of public funds λ , the share of profits leaking to the foreign firm ϕ is an additional factor pushing the government further toward economizing on the price premium needed to induce at-risk investment.

We showed above that $M(k, \lambda)$ is increasing in its second argument. Since it is being subtracted off, it is immediate that $\phi > 0$ implies $\hat{k}^f < \hat{k}^p$, so equilibrium with a foreign input supplier involves less at-risk investment than with an integrated domestic supplier as long as the foreign supplier has some positive bargaining weight. The amount of at-risk investment with a foreign input higher falls the higher is ϕ .

Proposition A4 *Under maintained assumptions, $\hat{k}^f \leq \hat{k}^p$, with strict inequality when $\phi > 0$. Thus, the government induces less at-risk investment on average across the distribution of k when procuring vaccine via a linear-price contract from an unintegrated privatized firm with foreign input supplier than from an integrated domestic firm.*

If instead of just the input supplier, both firms (or the whole integrated firm) is located in a foreign country, at-risk investment would be lower still. Instead of $\phi + \lambda$, the second argument of the moral-hazard term $M(\hat{k}^f, \phi + \lambda)$ in (A56) changes from $\phi + \lambda$ to $1 + \lambda$, an increase for all $\phi < 1$, increasing the moral-hazard term, reducing the at-risk investment threshold.

If the government is able to use richer contracts like those analyzed in Section 9.2.4, then the location of the input supplier may no longer matter. For example, suppose it offers a capacity subsidy $\sigma^r = k$ and linear price $p^r = c$. As we saw, this works to induce the second-best investment strategy. But it also allows the government to attain the second-best surplus. To see, this, note that the contract with capacity subsidy just allows the firm to break even. Hence the profits π_r and π_n equal 0 in equilibrium, so the fact that they have lower coefficients in the surplus function facing a foreign firm is irrelevant. The surpluses are the same as in the second best.

9.4 Incomplete Information

The irrelevance of the foreign firm location when the government can offer capacity subsidies is an artifact of the simplified economic environment we have studied to this point. The simplified environment involved so few frictions that the government can solve the moral-hazard problem and obtain the second best with a domestic or foreign firm. To detect nuances between the performance of domestic or foreign firms requires sufficient frictions that the second best is not so easily attained.

We have already seen this in the analysis restricting to linear-price contracts. With that friction, the domestic government was more inclined to sacrifice global efficiency to extract rent from the firm with a foreign firm whose surplus the government does not internalize. We will obtain the same flavor of a result in this section in a model introducing private information on the side of firms. The advantage of the approach here is more rigor, obtaining the looked-for results by suitably enriching the economic environment, not placing exogenous restrictions on contractual forms. The disadvantage is that the analysis is more complex, requiring the tools of mechanism-design. The disadvantage is mitigated by the

fact that the ultimate results delivered by the mechanism-design tools are not any more complex than earlier ones. Instead of having a moral-hazard term distorting the investment thresholds, we will have an adverse-selection term that looks quite similar.

As before, we will suppose unit capacity cost k is a random variable drawn from a distribution having pdf $f(k)$ and logconcave cdf $F(k)$ on support $[0, K]$. We now suppose that the draw of k is private observed by the integrated firm or by both U and D in model variants with unintegrated firms. 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, which have firm announces its type \tilde{k} and receive the provisions specified in the initial contract offer for that announcement. As an analytical device, we will allow rich contractual provisions, perhaps unrealistically rich, but end up showing that a quite simple contract not varying along many of those dimensions is optimal, proving that we could have restricted contracts to that simple form 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 $\sigma_n(\tilde{k})$. Rather than solving for the optimal contract separately first for an integrated domestic firm and then for a foreign input supplier, we will dive right in to analyzing the latter case by including the variable ϕ for the foreign firm's bargaining share and note as before that setting $\phi = 0$ nests the case of a fully domestic firm.

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 $k < \hat{k}$, 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

$$\pi_r(k) = \sigma_r - k \tag{A57}$$

$$\pi_n(k) = s(\sigma_n - k) \tag{A58}$$

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

$$\begin{aligned}
& -\beta ht_e + \int_0^{\hat{k}} \left\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a) - (1+\lambda)\sigma_r + (1-\phi)\pi_r(k) \right\} f(k) dk \\
& \int_{\hat{k}}^{\bar{k}} \left\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a - t_\ell) - s(1+\lambda)\sigma_n + (1-\phi)\pi_n(k) \right\} f(k) dk.
\end{aligned} \tag{A59}$$

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:

$$\pi_r(k) \geq \pi_n(k) \quad \forall k \in [0, \hat{k}] \tag{A60}$$

$$\pi_n(k) \geq \pi_r(k) \quad \forall k \in (\hat{k}, \bar{k}]. \tag{A61}$$

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

$$\pi_r(k) \geq 0 \quad \forall k \in [0, \hat{k}] \tag{A62}$$

$$\pi_n(k) \geq 0 \quad \forall k \in (\hat{k}, \bar{k}]. \tag{A63}$$

As usual, incentive compatibility binds only for the more not less productive types, so we can ignore constraints (A61). Also as usual, individual rationality binds only for the less not more productive types, so we can ignore constraints (A62). In view of the profit expressions in (A57)–(A58), one can show that a necessary and sufficient condition for the suite of constraints to hold is that they hold at the upper boundary:

$$\pi_r(\hat{k}) \geq \pi_n(\hat{k}) \tag{A64}$$

$$\pi_n(\bar{k}) \geq 0. \tag{A65}$$

These constraints bind at an optimum. Treating the constraints as equalities, substituting from (A57)–(A58), and solving the system of equations for the capacity subsidies yields

$$\sigma_r = (1-s)\hat{k} + s\bar{k} \tag{A66}$$

$$\sigma_n = \bar{k}. \tag{A67}$$

Substituting from (A66)–(A67) into (A59) and rearranging yields a new expression for the government's objective function

$$\begin{aligned}
& -\beta ht_e + \int_0^{\hat{k}} \left\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a) - (\phi+\lambda)[(1-s)\hat{k} + s\bar{k}] - (1-\phi)k \right\} f(k) dk \\
& \int_{\hat{k}}^{\bar{k}} \left\{ s[\theta\beta h - (1+\lambda)c](t_e - t_a - t_\ell) - (\phi+\lambda)s\bar{k} - (1-\phi)sk \right\} f(k) dk.
\end{aligned} \tag{A68}$$

Let \hat{k}^i denote the optimal threshold between the types that invest at risk and not at risk. Taking the first-order condition of (A68) with respect to \hat{k} , rearranging, and substituting the definition of \hat{k}^{**} from (A35) yields

$$\hat{k}^i = \hat{k}^{**} - \left(\frac{\phi + \lambda}{1 + \lambda} \right) \frac{F(\hat{k}^i)}{f(\hat{k}^i)}. \quad (\text{A69})$$

Similarly, let \bar{k}^i denote the optimal threshold between the types that invest not at risk and do not invest. Taking the first-order condition of (A68) with respect to \bar{k} , rearranging, and substituting the definition of \bar{k}^{**} from (A34) yields

$$\bar{k}^i = \bar{k}^{**} - \left(\frac{\phi + \lambda}{1 + \lambda} \right) \frac{F(\bar{k}^i)}{f(\bar{k}^i)}. \quad (\text{A70})$$

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

Proposition A5 *The following is an optimal procurement mechanism for the domestic government under incomplete information. For $k \in [0, \hat{k}^i]$, the firms (or firm) invest in sufficient at-risk capacity to serve the whole population each period, produce up to available capacity each period, are paid a per-dose price c and a capacity subsidy $\sigma_r^i = (1 - s)\hat{k}^i + s\bar{k}^i$. For $k \in (\hat{k}^i, \bar{k}^i]$, the firms (or firm) invest in sufficient not-at-risk capacity to serve the whole population each period, produce up to available capacity each period, are paid a per-dose price c and a capacity subsidy $\sigma_n^i = \bar{k}^i$.*

It is straightforward to compare the thresholds under incomplete information for at-risk investment \hat{k}^i and for any investment \bar{k}^i to their analogues in the second best, respectively, \hat{k}^{**} and \bar{k}^{**} , since they differ only by the last term in (A69) and (A70), 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 greater the greater the profit share ϕ that leaks to the foreign firm and the greater the social cost of public funds λ . The distortion term is inversely proportional to the ratio $f(k)/F(k)$ at the respective threshold, 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.

Incomplete information is a friction on the optimal-contracting problem, and it causes a similar distortion to the amount of at-risk investment as restriction to a linear price caused in the complete-information setting. To see this, compare equation (A69) to (A56).

9.5 International Cooperation

The formulas for the threshold \hat{k}^i for at-risk investment and the cutoff \bar{k}^i for any investment suggest a clear role for international cooperation. When the input supplier is located in a foreign country, the domestic government offering the vaccine procurement contract distorts the capacity subsidies downward to extract more rent from the firm, reducing the probability of at-risk investment relative to the second best as well as the probability of any investment relative to the second best.

The domestic and foreign governments can cooperate to increase the subsidy for at-risk investment from $\sigma_r^i = (1-s)\hat{k}^i + s\bar{k}^i$ to $(1-s)\hat{k}^{**} + s\bar{k}^{**}$. They can increase the subsidy for not-at-risk investment from $\sigma_n^i = \bar{k}^i$ to \bar{k}^{**} . This cooperative move would allow the countries to attain second best level of global surplus.

Proof of Lemma A1

The maximization problem (A23)–(A26) 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

$$\begin{aligned} \mathcal{L} = & -\beta h t_e + sQ_1(\theta\beta h - c)t_\ell + sQ_2(\theta\beta h - c)(t_e - t_a - t_\ell) - kX_r - skX_n \\ & + \gamma_1(X_r - Q_1) + \gamma_2(X_r + X_n - Q_2) + \mu_1(1 - Q_1) + \mu_2(1 - Q_2). \end{aligned} \quad (\text{A71})$$

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

$$\frac{\partial \mathcal{L}}{\partial X_r} = -k + \gamma_1^* + \gamma_2^* \leq 0 \quad (\text{A72})$$

$$\frac{\partial \mathcal{L}}{\partial X_n} = -sk + \gamma_2^* \leq 0 \quad (\text{A73})$$

$$\frac{\partial \mathcal{L}}{\partial Q_1} = s(\theta\beta h - c)t_\ell - \gamma_1^* - \mu_1^* \leq 0 \quad (\text{A74})$$

$$\frac{\partial \mathcal{L}}{\partial Q_2} = s(\theta\beta h - c)(t_e - t_a - t_\ell) - \gamma_2^* - \mu_2^* \leq 0, \quad (\text{A75})$$

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

$$X_r^*(-k + \gamma_1^* + \gamma_2^*) = 0 \quad (\text{A76})$$

$$X_n^*(-sk + \gamma_2^*) = 0 \quad (\text{A77})$$

$$Q_1^*[s(\theta\beta h - c)t_\ell - \gamma_1^* - \mu_1^*] = 0 \quad (\text{A78})$$

$$Q_2^*[s(\theta\beta h - c)(t_e - t_a - t_\ell) - \gamma_2^* - \mu_2^*] = 0, \quad (\text{A79})$$

complementary slackness conditions associated with constraints (A24)–(A26),

$$\gamma_1^*(X_r^* - Q_1^*) \quad (\text{A80})$$

$$\gamma_2^*(X_r^* + X_n^* - Q_2^*) \quad (\text{A81})$$

$$\mu_1^*(1 - Q_1^*) \quad (\text{A82})$$

$$\mu_2^*(1 - Q_2^*), \quad (\text{A83})$$

as well as nonnegativity of Lagrange multipliers, nonnegativity of the choice variables, and constraints (A24)–(A26) themselves.

The proof proceeds by analyzing a 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 (A83). Then

$$s\bar{k}^* = s(\theta\beta h - c)(t_e - t_a - t_\ell) \quad (\text{A84})$$

$$\leq \gamma_2^* + \mu_2^* \quad (\text{A85})$$

$$= \gamma_2^* \quad (\text{A86})$$

$$\leq sk, \quad (\text{A87})$$

where the first step follows from the definition of \bar{k}^* in (A27), the second step from (A75), the third step from $\mu_2^* = 0$, and the last step from (A73). 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

$$\mu_2^* \geq s(\theta\beta h - c)((t_e - t_a - t_\ell) - \gamma_2^*) \quad (\text{A88})$$

$$\geq s(\theta\beta h - c)((t_e - t_a - t_\ell) - sk) \quad (\text{A89})$$

$$\geq s(\theta\beta h - c)((t_e - t_a - t_\ell) - s\bar{k}^*) \quad (\text{A90})$$

$$= 0. \quad (\text{A91})$$

where the first step follows from (A75), the second step from (A77), and the third step from $k < \bar{k}^*$, and the last step from (A27). But $\mu_2^* > 0$ together with (A83) implies

$$Q_2^* = 1. \quad (\text{A92})$$

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 (A76). 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 (A24). 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 (A81), and the third from (A92). 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 (A25). Since $X_r^* > 0$, $\gamma_2^* = sk > 0$ by (A77), implying $Q_2^* = X_r^*$ by (A81), in turn implying $X_r^* = 1$ by (A92).

Case (iv) Suppose $X_r^* > 0$ and $X_n^* > 0$. Since $X_n^* > 0$, $\gamma_2^* = sk > 0$ by (A76), implying $1 = Q_2^* = X_r^* + X_n^*$ by (A81), in turn implying $X_r^* < 1$ since $X_n^* > 0$. But then $Q_1^* \leq X_r^* < 1$ implies $\mu_1^* = 0$ by (A82). Then

$$\gamma_1^* \geq s(\theta\beta h - c)t_\ell - \mu_1^* \quad (\text{A93})$$

$$= s(\theta\beta h - c)t_\ell \quad (\text{A94})$$

$$> 0, \quad (\text{A95})$$

where the first step follows from (A74), the second step from $\mu_1^* = 0$, and the third step from equation (A38). Now $\gamma_1^* > 0$ implies $Q_1^* = X_r^*$ by (A80), implying $Q_1^* > 0$ since $X_r^* > 0$, in turn implying

$$s(\theta\beta h - c)t_\ell = \gamma_1^* \quad (\text{A96})$$

from (A78) 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 (A96) yields

$$s(\theta\beta h - c)t_\ell = (1 - s)k. \quad (\text{A97})$$

Substituting (A97), $Q_1^* = X_r^*$, $Q_2^* = X_r^* + X_n^*$ into equation (A23) and rearranging yields a new expression for the planner's objective function,

$$-\beta ht_e + s(X_r^* + X_n^*)(\theta\beta h - c)(t_e - t_a - t_\ell) - sk(X_r^* + X_n^*). \quad (\text{A98})$$

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 A5

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

$$X_r(\tilde{k}) \geq 0 \quad (\text{A99})$$

$$X_n(\tilde{k}) \geq 0 \quad (\text{A100})$$

$$X_r(\tilde{k}) + X_n(\tilde{k}) \leq 1. \quad (\text{A101})$$

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

$$Q_1(\tilde{k}) \geq 0 \quad (\text{A102})$$

$$Q_2(\tilde{k}) \geq 0 \quad (\text{A103})$$

$$Q_1(\tilde{k}) \leq X_r(\tilde{k}) \quad (\text{A104})$$

$$Q_2(\tilde{k}) \leq X_r(\tilde{k}) + X_n(\tilde{k}). \quad (\text{A105})$$

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

$$\pi(k, \tilde{k}) = R(\tilde{k}) - X(\tilde{k})k, \quad (\text{A106})$$

defining expected revenue paid to the firm

$$\begin{aligned} R(\tilde{k}) &= \sigma_r(\tilde{k})X_r(\tilde{k}) + s\sigma_n(\tilde{k})X_n(\tilde{k}) \\ &\quad + s[p_1(\tilde{k}) - c]Q_1(\tilde{k})t_\ell + s[p_2(\tilde{k}) - c]Q_2(\tilde{k})(t_e - t_a - t_\ell) \end{aligned} \quad (\text{A107})$$

and expected total capacity installed

$$X(\tilde{k}) = X_r(\tilde{k}) + sX_n(\tilde{k}). \quad (\text{A108})$$

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:

$$p_1(\tilde{k}) = p_2(\tilde{k}) = c. \quad (\text{A109})$$

The government designs the contract to maximize its expected surplus, which upon substituting from (A109) for the prices can be written

$$\begin{aligned} & -\beta ht_e + \int_0^K \left\{ s[\theta\beta h - (1 + \lambda)c][Q_1(k)t_\ell + Q_2(k)(t_e - t_a - t_\ell)] \right. \\ & \quad \left. - (1 + \lambda)[\sigma_r(k)X_r(k) + s\sigma_n(k)X_n(k)] + (1 - \phi)\pi(k, k) \right\} f(k)dk, \end{aligned} \quad (\text{A110})$$

subject to incentive compatibility

$$\pi(k, k) \geq \pi(k, \tilde{k}) \quad \forall k \geq 0, \tilde{k} \geq 0 \quad (\text{A111})$$

and individual rationality

$$\pi(k, k) \geq 0 \quad \forall k \geq 0. \quad (\text{A112})$$

An increase in each of $Q_1(k)$ and $Q_2(k)$ increases the objective function while leaving constraints (A111)–(A112) 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 (A109) holds. Hence, we can take (A104) and (A105) to bind without loss of generality. Substituting for $Q_1(k)$ and $Q_2(k)$ treating (A104)–(A105) as equalities as well as from (A106)–(A109) into (A110) yields, after rearranging,

$$\begin{aligned} -\beta ht_e + \int_0^K \left\{ \{s[\theta\beta h - (1 + \lambda)c](t_e - t_a) - (1 - \phi)k\} X_r(k) \right. \\ \left. + s\{\theta\beta h - (1 + \lambda)c](t_e - t_a - t_\ell) - (1 - \phi)k\} X_n(k) \right. \\ \left. - (\phi + \lambda)R(k) \right\} f(k) dk. \quad (\text{A113}) \end{aligned}$$

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 (A111) implies

$$\Pi(k) = R(k) - X(k)k \quad (\text{A114})$$

$$= \max_{\tilde{k} \geq 0} [R(\tilde{k}) - X(\tilde{k})k]. \quad (\text{A115})$$

By the Envelope Theorem, (A115) implies $\Pi'(k) = -X(k)$. The Fundamental Theorem of Calculus then gives

$$\Pi(k) = \int_k^K X(z) dz - \pi(\infty) \quad (\text{A116})$$

$$= \int_k^K X(z) dz, \quad (\text{A117})$$

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

$$R(k) = \Pi(k) + X(k)k \quad (\text{A118})$$

$$= \int_k^K X(z) dz + X(k)k, \quad (\text{A119})$$

where (A118) follows from (A114) and (A119) from (A117). Integrating,

$$\int_0^K R(k) f(k) dk = \int_0^K \left[\int_k^K X(z) dz \right] f(k) dk + \int_0^K X(k) k f(k) dk. \quad (\text{A120})$$

Using integration by parts to evaluate the first integral on the right-hand side,

$$\int_0^K \left[\int_k^K X(z) dz \right] f(k) dk = \left[F(k) \int_k^K X(z) dz \right]_{k=0}^{k=K} + \int_0^K F(k) X(k) dk \quad (\text{A121})$$

$$= \int_0^K F(k) X(k) dk. \quad (\text{A122})$$

Substituting from (A122) into (A120) and then substituting for $X(k)$ from (A108) yields

$$\int_0^K R(k)f(k)dk = \int_0^K F(k)[X_r(k) + sX_n(k)]dk + \int_0^K [X_r(k) + sX_n(k)]kf(k)dk. \quad (\text{A123})$$

We can then use (A123) to substitute out for the term in $R(k)$ in objective function (A113), obtaining

$$\begin{aligned} & -\beta ht_e + \int_0^K \left\{ s[\theta\beta h - (1 + \lambda)c](t_e - t_a) - (1 + \lambda)k \right\} X_r(k) \\ & + s \left\{ [\theta\beta h - (1 + \lambda)c](t_e - t_a - t_\ell) - (1 + \lambda)k \right\} X_n(k) \right\} f(k)dk \\ & - (\phi + \lambda) \int_0^K F(k)[X_r(k) + sX_n(k)]dk. \quad (\text{A124}) \end{aligned}$$

We claim that this objective function incorporates the information from constraints (A111)–(A112). We will demonstrate this by maximizing (A124) ignoring those constraints and verifying at the end that the solution satisfies them. Constraints (A111)–(A112) are the only place where different realizations of the random variable k are linked. Ignoring (A111)–(A112) severs 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 (A124) for each k subject to (A99)–(A101). 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

$$\begin{aligned} \mathcal{L} = & \left\{ s[\theta\beta h - (1 + \lambda)c](t_e - t_a) - (1 + \lambda)k \right\} f(k) - (\phi + \lambda)F(k) \Big\} X_r(k) \\ & + s \left\{ [\theta\beta h - (1 + \lambda)c](t_e - t_a - t_\ell) - (1 + \lambda)k \right\} f(k) - (\phi + \lambda)F(k) \Big\} X_n(k) \\ & + \gamma[1 - X_r(k) - X_n(k)], \quad (\text{A125}) \end{aligned}$$

where γ is the Lagrange multiplier on constraint (A101). The Kuhn-Tucker conditions include inequalities bounding partials of the Lagrangian,

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial X_r(k)} = & \left\{ s[\theta\beta h - (1 + \lambda)c](t_e - t_a) - (1 + \lambda)k \right\} f(k) \\ & - (\phi + \lambda)F(k) - \gamma^i \leq 0 \quad (\text{A126}) \end{aligned}$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial X_n(k)} = & \left\{ [\theta\beta h - (1 + \lambda)c](t_e - t_a - t_\ell) - (1 + \lambda)k \right\} f(k) \\ & - (\phi + \lambda)F(k) - \gamma^i \leq 0, \quad (\text{A127}) \end{aligned}$$

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

$$X_r^i(k) \frac{\partial \mathcal{L}}{\partial X_r(k)} = 0 \quad (\text{A128})$$

$$X_n^i(k) \frac{\partial \mathcal{L}}{\partial X_n(k)} = 0, \quad (\text{A129})$$

complementary slackness conditions associated with constraint (A101),

$$\gamma^i [1 - X_r^i(k) - X_n^i(k)] \quad (\text{A130})$$

as well as nonnegativity of Lagrange multipliers, nonnegativity of the choice variables, and constraint (A101) itself. We have added superscript i to variables $X_r^i(k)$, $X_n^i(k)$, and γ^i satisfying the Kuhn-Tucker conditions to distinguish the optimum in this model of incomplete information from optima in other settings.

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

Case (i) Suppose $X_r^i(k) = X_n^i(k) = 0$. Then $\gamma^i = 0$ by (A130). Substituting $\gamma^i = 0$ into (A126)–(A127), dividing by $(1 + \lambda)f(k)$, and rearranging yields

$$s \left(\frac{\theta\beta h}{1 + \lambda} - c \right) (t_e - t_a) \leq k + \left(\frac{\phi + \lambda}{1 + \lambda} \right) \frac{F(k)}{f(k)} \quad (\text{A131})$$

$$\left(\frac{\theta\beta h}{1 + \lambda} - c \right) (t_e - t_a - t_\ell) \leq k + \left(\frac{\phi + \lambda}{1 + \lambda} \right) \frac{F(k)}{f(k)}. \quad (\text{A132})$$

Let \bar{k}^i be the infimum of the values of k satisfying (A132) with equality. The derivative of the right-hand side of (A132) with respect to k is

$$1 + \left(\frac{\phi + \lambda}{1 + \lambda} \right) \left[\frac{f(k)^2 - f'(k)F(k)}{f(k)^2} \right]. \quad (\text{A133})$$

The second term is positive if and only if $F(k)$ is logconcave. Since $F(k)$ is logconcave by assumption, derivative (A133) is positive, implying that (A127) is satisfied for all $k \geq \hat{k}^i$. By (A18), (A131) is a weaker condition than (A132). Hence, (A131) is satisfied for all $k \geq \hat{k}^i$. Thus $k \geq \hat{k}^i$ is a necessary and sufficient condition for solution $X_r^i(k) = X_n^i(k) = 0$. By (A34), the \hat{k}^i satisfying (A132) with equality here is the same \hat{k}^i given in equation (A70) in the text.

Case (ii) Suppose $X_r^i(k) > 0$ and $X_n^i(k) > 0$. By (A128)–(A129) both (A126) and (A127) must be satisfied with equality. Since both equal 0, the left-hand sides must be equal. Setting the left-hand sides of (A126) equal to the left-hand side of (A127) and rearranging yields

$$s(t_e - t_a) = t_e - t_a - t_\ell, \quad (\text{A134})$$

violating assumption (A18), ruling out this case.

Case (iii) Suppose $X_r^i(k) > 0$ and $X_n^i(k) = 0$. Since $X_r^i(k) > 0$, (A126) must be satisfied with equality by (A128). Thus,

$$\{s[\theta\beta h - (1 + \lambda)c](t_e - t_a) - (1 + \lambda)k\}f(k) - (\phi + \lambda)F(k) \quad (\text{A135})$$

$$= \gamma^i \quad (\text{A136})$$

$$\geq \{[\theta\beta h - (1 + \lambda)c](t_e - t_a - t_\ell) - (1 + \lambda)k\}f(k) - (\phi + \lambda)F(k) \quad (\text{A137})$$

Rearranging (A135)–(A137) gives

$$k + \left(\frac{\phi + \lambda}{1 + \lambda}\right) \frac{F(k)}{f(k)} \leq \frac{s}{1 - s} \left(\frac{\theta\beta h}{1 + \lambda} - c\right) t_\ell. \quad (\text{A138})$$

Let \hat{k}^i be the supremum of the values of k satisfying (A138). The positive sign on the derivative (A133) implies that the left-hand side of (A138) is increasing in k . Hence, (A138) is satisfied for all $k \leq \hat{k}^i$. By (A35), the \hat{k}^i satisfying (A138) with equality here is the same \hat{k}^i given in equation (A69) in the text.

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

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 (A111)–(A112). Consider a firm with type $k' \in [0, \hat{k}^i]$ that invests at risk in equilibrium according to the optimal contract. The firm's expected profit equals $\Pi(k') = (1 - s)\hat{k}^i + s\bar{k}^i - k'$ in equilibrium. The firm earns the same expected profit by deviating to any announcement $\tilde{k} \in [0, \hat{k}^i]$, because the contractual provisions are the same for all types in that interval. If the firm deviates to announcement $\tilde{k} \in (\hat{k}^i, \bar{k}^i]$, its expected profit equals $\pi(k', \tilde{k}) = s(\bar{k}^i - k') \leq \Pi(k')$ for all $k' \leq \hat{k}^i$. If the firm deviates to $\tilde{k} > \bar{k}^i$, its expected profit equals 0, which is strictly less than $\Pi(k')$ for $k' \leq \hat{k}^i$. Thus the incentive-compatibility constraint (A111) is satisfied for all $k' \leq \hat{k}^i$. We also just verified that individual-rationality constraint (A112) is satisfied for them as well. The reader can similarly verify that constraints (A111) and (A112) are satisfied for other types as well. *Q.E.D.*