NBER WORKING PAPER SERIES

AN OPTIMAL MECHANISM TO FUND THE DEVELOPMENT OF VACCINES AGAINST EMERGING EPIDEMICS

Christopher M. Snyder
Kendall Hoyt
Dimitrios Gouglas

Working Paper 30619
http://www.nber.org/papers/w30619

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2022

Early version presented at the "New Incentives for Developing, Manufacturing, Distributing, and Dispensing Medical Countermeasures" meeting, held in June 2019 in Washington DC. Snyder received support for initial work from the Nuclear Threat Initiative. The authors are grateful to Sophie Smith and Blake Dignes for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

An Optimal Mechanism to Fund the Development of Vaccines Against Emerging Epidemics
Christopher M. Snyder, Kendall Hoyt, and Dimitrios Gouglas
NBER Working Paper No. 30619
November 2022
JEL No. D47,H44,I18,L65,O31

ABSTRACT

We derive the optimal funding mechanism to incentivize development and production of vaccines against diseases with epidemic potential. In the model, suppliers' costs are private information and investments are noncontractible, precluding cost-reimbursement contracts, requiring fixed-price contracts conditioned on delivery of a successful product. The high failure risk for individual vaccines calls for incentivizing multiple entrants, accomplished by the optimal mechanism, a (w+1)-price reverse Vickrey auction with reserve. Our analysis determines the optimal number of entrants and required funding level. Based on a distribution of supplier costs estimated from survey data, we simulate the optimal mechanism's performance in scenarios ranging from a small outbreak, causing harm in the millions of dollars, to the Covid-19 pandemic, causing harm in the trillions. We assess which mechanism features contribute most to its optimality.

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

Kendall Hoyt
Geisel School of Medicine
Dartmouth College
1 Rope Ferry Rd
Hanover, NH 03755
kendall.l.hoyt@dartmouth.edu

Dimitrios Gouglas
Coalition for Epidemic Preparedness Innovations,
Oslo, Norway
dimitrios.gouglas@cepi.net

1. Introduction

Outbreaks of infectious diseases can emerge rapidly with little warning. These outbreaks can fizzle out (as with the 2003 SARS outbreak), kindle regional epidemics (as with the 2014 Ebola outbreak in West Africa), or ignite global pandemics (as with the Covid-19 outbreak still underway). When available, vaccines can contain the worst mortality, morbitidy, and economic losses from a disease outbreak. For example, Woolsey and Geisbert (2021) credit Merck's Ebola vaccine with curbing the West African epidemic and quelling subsequent flare ups. Castillo et al. (2021) conservatively estimate the social value of Covid-19 vaccines at over $17 trillion.

Typically, public organizations subsidize commercial firms to conduct late-stage clinical trials of vaccines and scale up manufacturing. For example Gavi, the Vaccine Alliance, paid Merck $178 million to stockpile 500,000 doses of its Ebola vaccine in the West African epidemic (Sagonowsky, 2019). In the Covid-19 pandemic, the U.S. government spent $18 billion on six vaccine candidates as part of its Operation Warp Speed program (Baker et al., 2020). For a variety of reasons—ranging from incentives to free ride on others' vaccination (Geoffard and Philipson, 1997) to the difficulty in extracting revenue from consumers with private information about disease risk (Kremer and Snyder, 2015)—the commercial vaccine market is unlikely to supply adequate investment and production without dedicated public support. There is a long history of public support for routine childhood vaccinations in both low-income countries (through UNICEF, founded in 1946) and high-income countries (for example, the U.S. Vaccines for Children program, launched in 1994). Public support for epidemic vaccines is a more recent idea. Initiatives like the U.S. Biomedical Advanced Research and Development Authority (BARDA), created in 2006, and the Coalition for Epidemic Preparedness Innovations (CEPI), an international nonprofit launched in 2017, provide resources to reduce the cost and risk of early-stage research and development.

Still lacking are robust mechanisms to incentivize pharmaceutical companies to scale-up manufacturing and to bring candidates through licensure. Without industry support for these late-stage activities, few epidemic vaccine candidates advance from the laboratory. The few that have lack the speed, scale, and accessibility required to maximize their social value. For example, Merck's aforementioned Ebola vaccine, effective against the Zaire variant, is ineffective against a Sudan variant currently breaking out in Uganda (Dahir, 2022). Vaccine candidates against the Sudan variant have been developed but none has been stockpiled to facilitate rapid clinical trials and limit spread

in an outbreak. The World Health Organization maintains a list of other infectious diseases having epidemic potential but lacking an approved vaccine including lassa fever, MERS, SARS, and Zika (WHO, 2018), which could also benefit from having a stockpile ready for the next outbreak. Even an identified success like Operation Warp Speed (Gross and Sampat (2022) call it "the most successful piece of the Covid-19 response") may have left room for improvement according to principles of optimal mechanism design. In view of the enormous insurance value of having an effective vaccine during a devastating pandemic, Athey et al. (2020) recommended spending an order of magnitude more, $70 billion to develop and procure 15–20 vaccine candidates.

Our analysis aims to assist public funders in identifying the appropriate level of funding under a wide range of outbreak scenarios and technical landscapes. Suppose, for example, that several vaccine candidates have shown promising results in early clinical trials. How many should be supported? Should support focus on the single most promising candidate? Or does the notoriously low probability of success in vaccine trials call for supporting several candidates? It may not be efficient to support all potential candidates since the least productive of them may have extremely high costs, and there is a reasonable chance that some vaccine candidates succeed without including every candidate.

In a perfect-information world, it would be a simple matter to attain the first best using a simple push-funding mechanism, auditing and reimbursing of participating suppliers' development and production costs. Suppliers could be allowed a profit margin sufficient to induce their participation, but this margin could be precisely controlled and kept to a modest level. Perfect information is unrealistic, however; real-world suppliers have considerable private information regarding their costs. A promise to defray their costs can generate subtantial moral hazard, leading suppliers to relax efforts to control costs and perhaps shift expenses from unreimbursed reimbursed lines of business, inflating measured costs. The large variation in this private information can be substantial: from a survey of vaccine suppliers' cost of completing late-stage clinical trials and producing a moderate stockpile, we will report estimates of a standard deviation of $120 million compared to a mean of about $240 million.

In this paper, we propose an optimal vaccine funding mechanism in a model in which incomplete information precludes the use of push funding, requiring the use of pull funding, that is, payment conditional on success, to incentivize suppliers to complete relevant milestones. (An equivalent, less extreme assumption is that push funding can be used for some initial milestones, and all available

such opportunities have been exploited, leaving remaining milestones to be incentivized with pull funding.) The mechanism accounts for suppliers' private information about their widely varying productivities. If the social benefit of having a vaccine stockpile is high—say because the disease is extremely harmful or poses a danger of spreading to pandemic proportions—the optimal mechanism should lean toward inducing multiple suppliers to participate, providing enough redundancy that even if several fail, the probability that at least one succeeds is reasonably high. Pulling in more suppliers inflates program costs for two reasons. First, the participation of multiple suppliers involves redundant investment. Second, the marginal supplier pulled in is less efficient than inframarginal suppliers, increasing the award that has to be paid to all participating suppliers under asymmetric information about supplier costs. Productive inframarginal suppliers receive an information rent that is increasing in the number of participating suppliers. All this expense is worthwhile if the social benefit of a vaccine is high enough. On the other hand, if the social benefit of a vaccine is low, the mechanism should adapt in the opposite direction, inducing the participation of fewer suppliers to economize on program costs.

The optimal mechanism has the form of a reverse Vickrey auction with multiple winners. The mechanism determines the number of winners endogenously as the intersection between the ascending schedule of submitted cost bids and a descending reserve-price schedule for which we derive a formula. Winners receive an award equal to the marginal losing supplier's bid or the reserve price for that many winners, whichever is lower. The logic from a standard Vickrey auction, a second-price auction with one winner, is that truthful bidding is ensured by severing the dependence of the award on the winner's bid. That logic extends to a $(w+1)$-price Vickrey auction with w winners and reserve, which is the optimal mechanism here.

We simulate the performance of the mechanism based on a distribution of supplier costs estimated from a survey of firms in the vaccine market collected by CEPI. Assuming that investment costs up to phase-3 clinical trials have been defrayed by push funding from other sources, we estimate the distribution of further costs involved in completion of a moderate stockpile to respond quickly to an epidemic outbreak of the disease in question. We estimate a lognormal distribution for the cost of phase-3 trials and a lognormal distribution for the manufacture of a moderate stockpile, and then combine those distributions into an all-in cost (also with a lognormal distribution). Results are displayed for a grid of key model parameters including the number of potential vaccine candidates n, a candidate's probability of success s, and the social benefit from the vaccine stockpile

b.

We use the simulations to gain an understanding of the number of winners the mechanism generally selects, the level of program spending, the net social benefit it generates, among other outcomes. We undertake a comparative-statics analysis of how those outcomes vary with the key model parameters. We also compare the performance of the mechanism to more or less efficient alternatives. We compare our second-best mechanism (optimal under a constraint of asymmetric cost information) to the first-best mechanism, which in effect allows the funder to perfectly observe supplier costs and offer discriminatory payments conditional on that information. To better understand the contribution of various features of our second-best mechanism, we also compare its performance to that of mechanisms shutting down those features one at a time, for example, removing the reserve price, or removing the Vickrey aspect of having the marginal losing bid constrain the award and just relying on the reserve price, or setting an exogenous number of winners in advance.

We find that the second-best mechanism can capture an increasing share of available surplus as b increases, up to 93% of the surplus captured by the first-best mechanism for the highest values of b considered. Shutting down features of the second-best mechanism can severely impair its performance.

The range of social benefits used in the epidemic simulations, from $500 million and $3 billion, are intended to apply to a situation in which a high-income country like the United States seeks a stockpile to quell an epidemic of a disease like Ebola or SARS abroad. The altruistic benefit of saving morbitity and mortality abroad and the self-interest in eliminating the small chance that the outbreak jumps to the home country and spreads there, substantial as they might be, pales in comparison to the benefit of quelling the spread of a pandemic to the home country. The general nature of the mechanism allows it to apply to either situation. Simulations of the mechanism in a pandemic setting have already been undertaken by a companion paper (Snyder et al., 2020), which applies the mechanism derived here to the Covid-19 pandemic. Here, we reprise some of the key simulation results from that paper and compare them to new results for smaller epidemics. Although the mechanism is the same, the new inputs to be substituted into the mechanism dictated by the pandemic scenario lead to a wholesale change in outcomes. Rather than selecting one or two winners, once benefits of quelling a pandemic amounting in the trillions of dollars are inputted, the mechanism sets such a high award that it becomes virtually certain all available candidates are induced to participate. Spending is exorbitant, with awards in the tens of billions of dollars, but this spending is eclipsed by

the losses wrought by the pandemic. Reassuringly, therefore, the mechanism automatically adapts to the scenario under consideration. In moderate epidemic scenarios, the mechanism economizes on costs, limiting the number of winners; in pandemic scenarios, it reflects the strong intuition that "all hands on deck" is optimal then.

Our paper contributes to the economic literature on innovation incentives. A series of papers analyzed alternative incentive mechanisms including Scotchmer (1999) on patent renewals, Kremer (1998) and Galasso et al. (2016) on patent buyouts, and Ridley et al. (2006) on priority review vouchers. Chari et al. (2012) and Weyl and Tirole (2012) explore hybrid mechanisms incorporating patents and prizes, optimally trading off rents paid to suppliers versus deadweight loss. Our paper focuses on prizes since markets for vaccines against emerging outbreaks may only be nascent, limiting the value of patents. Further limiting the value of patents, due to epidemiological externalities (Geoffard and Philipson, 1997) and other market frictions (Kremer and Snyder, 2015), vaccines tend to be procured by governments and central agencies. Prices tend to be set administratively rather than determined by a free market, limiting the rents provided by a patent. In this paper, we remain largely agnostic about whether the prize comes in the form of a lump-sum monetary payment, the grant of a priority review voucher, or is factored into a high price per dose for a stockpile. Our goal is to determine the size of the prize and the number of suppliers selected to invest in a risky project.

Our paper is part of the literature on innovation in pharmaceutical markets, including the seminal contribution of Acemoglu and Linn (2004) on the contribution of market size to pharmaceutical innovations. We focus on vaccines for several reasons. Vaccines typically are procured by governments, calling for a centralized procurement mechanism such as the one proposed in this paper. The complex facilities involved in vaccine production entail substantially higher capital investment than most other pharmaceuticals. Such capital investments can be a barrier to entry and source of substantial private information. Vaccines' unusually low probability of success may call for multiple innovators to ensure at least one success. A mechanism is required to determine how many should participate and how much they should be paid. Within the vaccine market, we train our focus even more narrowly on emerging outbreaks. Speed is a key factor in emerging outbreaks, calling for innovators to invest in parallel rather than sequentially as might economize on expenditures without much loss of benefits for endemic diseases that may not be expanding as rapidly.

Our mechanism is related to advance market commitments (AMCs) proposed to fund vaccines by Kremer and Glennerster (2004) and formally analyzed in Kremer et al. (2022). AMCs commit

funds for a subsidy on top of what the end user (say a developing country or a donor on its behalf) pays for a vaccine if and when research and development bear fruit. Our mechanism involves pull funding as does an AMC, and since we are largely agnostic about the form of the prize offered by our mechanism, could be structured to resemble an AMC; but doing so would involve some tensions. Unlike the original AMC proposal, the fund size is not predetermined in our mechanism but depends on suppliers' announced costs. If suppliers end up having lower than expected costs, the mechanism may induce the participation of more of them with potentially ambiguous effects on total program payments. Assuming as we do for much of the analysis that the mechanism is targeting a fixed stockpile, the award to suppliers can be expressed equivalently as a per-unit subsidy, as in the original AMC proposal, or a lump-sum payment. If applied to a setting in which quantity is endogenously determined rather than a fixed stockpile, expressing the award as a per-unit subsidy may have a danger of generating perverse incentives. More effective vaccines may quells an epidemic faster, leading fewer vaccines to be sold, reducing the total award.

Our paper is closest to a companion paper, Snyder et al. (2020), which applies the optimal mechanism to the case of Covid-19 vaccines. Here, we provide a full derivation of the optimal mechanism sketched in the companion paper's online appendix. We prove that the mechanism is optimal not just in the class of deterministic mechanisms but in the class of stochastic mechanisms. We analyze how much various features of the mechanism contribute to its performance in simulations that selectively omit the features. The estimation of the cost of stockpiles to combat a moderate outbreak and simulations of the optimal mechanism in that case are new in this paper. Section 6 on pandemic scenarios cites some of the simulation results from the companion paper, but that section also provides new material, including a formalization of the pandemic scenario that shows how it can be nested in a general model of epidemic outbreaks of any size. The section also provides new simulation runs for the pandemic scenario facilitating the closest possible comparison to the new simulation results for a moderate outbreak.

A series of papers (Ahuja et al., 2021; Athey et al., 2022) have also sought to construct the optimal portfolio of Covid-19 vaccines. Several simplifications allow them to bypass the full mechanism-design approach undertaken here: they assume the program designer has perfect information about supplier heterogeneity yet constrain the program to offer a uniform contract that does not leverage the perfect information. In this paper, we assume the designer knows the distribution of supplier heterogeneity but not private realizations. The contract ultimately offered to auction winners is

uniform, but uniformity is required to ensure truthful revelation of private information, not an exogenous contractual restriction.

The paper is organized as follows. Section 2 provides an informal description of the optimal mechanism, and Section 3 provides formal derivations. Section 4 discusses the estimation of the distribution of supplier costs using data from CEPI's supplier survey. Section 5 uses the estimated distribution to calibrate the performance of the optimal mechanism facing a moderate epidemic. Section 6 provides calibrations for a larger pandemic. Section 7 concludes. Technical details are relagated to a series of appendixes.

2. Description of Optimal Mechanism

This section describes the optimal funding mechanism in intuitive terms and provides an example to fix ideas. More formal derivations are postponed until the subsequent section.

For brevity, we refer to the optimal funding mechanism as a w-winner Vickrey auction with reserve (abbreviated WVAR, with suggested pronunciation "waiver"). The WVAR works as follows. Each of n potential suppliers submit bid c_i. After unsealing the bids, the funder decides on the number w of auction winners, which are eligible to invest and receive an award if the investment proves successful. The award is set according to a $(w+1)$-price version of a Vickrey auction with reserve price. Specifically, the bids c_i are ordered from smallest to largest. The lowest w bidders win the auction. Their award is set equal to the next higher bid (bid $w+1$) unless this exceeds a reserve price, in which case the award is set to the reserve price, computed as explained below.

Since c_i is private information, a supplier could always inflate its cost report. The virtue of the Vickrey auction is that it induces truth telling; regardless of how rivals behave, a supplier prefers to report its true cost c_i than any higher or lower cost. A supplier's bid only determines whether it wins; how much it wins depends not on its bid but on losing bids or a pre-specified reserve price. Truthful bidding ensures the supplier wins exactly when this is profitable, i.e., when the award exceeds its cost.

For example, suppose there are five potential suppliers with cost draws 100, 150, 275, 330, and 450 in ascending order (all monetary values in this section in millions). This cost schedule is graphed in red in Figure 1.

The mechanism juxtaposes the cost schedule against a demand curve reflecting the funder's will-

ingness to pay for the investment of the marginal supplier when w of them are selected as winners. Formally, funder demand equals $(1-s)^{w-1}sb$, the product of three factors: the probability $(1-s)^{w-1}$ that none of the $w-1$ other suppliers successfully completes their project, the probability s that supplier w succeeds, and the benefit b from a successful project. Figure 1 graphs funder demand assuming $b = 3000$ and $s = 50\%$ as the light blue curve.

In a textbook competitive market, the intersection between supply and demand determines equilibrium. Here, because the funder controls market design, it can retain more surplus by implementing a reserve price (a general component of optimal auction design; see Myerson (1981)). The dark blue curve graphs the optimal reserve price in our model based on the formula derived in the next section. This formula factors in both the funder's willingness to pay on the demand side and the supplier cost distribution on the supply side. The reserve price starts out well below demand, but the curves converge as the number of investing suppliers increases.

The equilibrium, drawn as an open circle, is the highest point at the intersection between the cost and reserve-price schedules. In this example, the mechanism optimally selects the two lowest-cost suppliers as winners, paying them an award with an expected value of 275, the cost of the next highest bid. The optimal reserve price, represented by the height of the blue dot above the open circle, is 321. In this example, the reserve price is not needed to constrain the award; competition among the bidding suppliers proves sufficient.

In this example, 275 represents the ex ante expected value of the award, averaged across success and failure. The mechanism we consider conditions payments on successful delivery of the stockpile. To preserve the expected value of the ex ante award in that case, the ex post award needs to be scaled up by the reciprocal of success probability s since winners are only paid with probability s. Given $s = 50\%$ in the example, the ex post award would need to be 550, double the ex ante expectation of 275.

There are several benefits from having the mechanism condition payments on success. Making a payment but getting no results may be a political risk, which the funder can avoid with conditional payments. Conditional payments help deter a "take the money and run" strategy, whereby a supplier pockets an up-front payment without making good on the investment commitment, saving itself some or all of the investment expense. Conditional payments can mitigate other moral-hazard and adverse-selection problems. While the model takes the probability of success s as exogenous, in richer settings in which it is endogenous, conditional payments can incentivize productive invest-

ment decisions and discourage participation of suppliers that are unlikely to succeed.[1]

As described, the mechanism procures multiple stockpiles if multiple auction winners happen to succeed. In the model, realizing benefit b requires only one stockpile, so stockpiles beyond the first are redundant. That might lead one to consider an alternative mechanism that only procures a stockpile from one successful supplier. This alternative will not improve on the mechanism as described, however, since there is no way to ensure suppliers will be successful in producing a stockpile without having them actually produce it; and once suppliers have produced their stockpiles, there are no savings from cancelling redundant ones. If the constraint that the mechanism only procure a single stockpile were imposed on the mechanism, a race or coin flip or other way to break ties among multiple successful suppliers would need to be devised and the award to the tie-break winner would need to be scaled up by the expected number of successful winners.

3. Derivation of Optimal Mechanism

This section provides a formal derivation of the optimal vaccine procurement mechanism, proving it has the form of a w-winner Vickrey auction with reserve (WVAR). The section also derives the formula for the reserve price.

We proceed in three steps. First, we lay out a formal model. Next, we state a result due to Chaturvedi and Martínez-de-Albéniz (2011) (hereafter CM), building on Myerson's (1981) seminal work on optimal auction design, characterizing the optimal mechanism for procuring multiple units from multiple suppliers with failure risk. CM's result provides the foundation for our last step, proving the optimality of WVAR in the specific setting of vaccine procurement as a corollary of CM's general characterization.

3.1. Model

A risk-neutral funder, labeled player 0, seeks to address a disease with epidemic potential by incentivizing the phase-3 testing of a vaccine and the manufacture of a stockpile of doses of some target size. Let $b \geq 0$ be the present discounted value to the funder of the stream of health and economic

[1] Up-front payments would have some advantages for reasons outside of the model. Suppliers are risk neutral and face no financial frictions in the model. In a richer model, up-front payments could provide effective insurance valued by risk-averse suppliers and alleviate financial constraints. Up-front payments can deliver the same expected award with lower nominal payments, which may look better politically.

benefits from achieving the stockpile endpoint via the mechanism. There are several interpretations of b. If the stockpile would not have been produced absent the funding mechanism, then b can be simply interpreted as the gross surplus from attaining the stockpile endpoint at all. If the stockpile would have been produced in any event, albeit more slowly, then b can be interpreted as the benefit from accelerating the arrival of the stockpile via the mechanism.

The funder, lacking the capability to undertake pharmaceutical R&D itself, must induce suppliers to make required investments. Assume a natural number n of risk-neutral suppliers, indexed by $i \in \{1,\ldots,n\}$, have the requisite R&D and manufacturing capability. For a supplier to reach the stockpile endpoint, it must succeed in two investment stages. In the first stage, it must manufacture the vaccine doses used in phase-3 trials and undertake those trials. If trials prove successful, in the second stage, the supplier must manufacture more of the vaccine to create a stockpile. Assume for simplicity that these two stages must be undertaken jointly. One justification for this assumption is that transfering the technology behind the complex vaccine-manufacturing process is too difficult to allow the manufacturer of doses used in phase-3 trials to "pass the baton" to a different firm that would then manufacture the stockpile.

Before investing, each supplier assesses its individual cost $c_i \geq 0$ of taking the vaccine through the combined investment stages. R&D investment is a complex dynamic process involving continuous expenditure that must be discounted according to when it is made and involving the option to terminate if prospects dim (Childs and Triantis, 1999). We capture this complexity by assuming c_i is the present discounted value of the payment required to induce supplier i to initiate investment. Investment costs may include monetary outlays as well as the opportunity cost of diverting manufacturing facilities from pharmaceuticals under production to this vaccine. We model c_i as an independently and indentically distributed (iid) draw, having probability density function (pdf) $f(c_i)$, cumulative distribution function (cdf) $F(c_i)$, and support $[0,\bar{c}]$, where $\bar{c}$ can be potentially infinite. Using the convention that boldface indicates a vector of variables across suppliers, we have $\mathbf{c} = (c_1,\ldots,c_n)$.

If supplier i, invests, it learns the value of an indicator $A_i \in \{0,1\}$ for the success of its combined investment. Assume A_i is an iid draw from a Bernoulli distribution, letting s denote the probability of success, $s = \Pr(A_i = 1)$. Let $\mathbf{A} = (A_1,\ldots,A_n)$.

While suppliers are heterogeneous ex post, after making their individual draws from cost and success distributions, suppliers are homogeneous from an ex ante perspective since the distributions

are identical and suppliers differ in no other dimensions.

Assume that different vaccine candidates are perfect substitutes, implying that the funder obtains the same benefit B whether one or more than one supplier succeeds. The funder may still want to induce multiple suppliers to invest because individual suppliers can fail when $s < 100\%$. The participation of more candidates increases the chance that at least one succeeds.

The funder enlists the participation of suppliers in a funding mechanism given by the triple $(\mathbf{q}(\mathbf{c}), \mathbf{y}(\mathbf{c}), \mathbf{x}(\mathbf{c}))$, where $\mathbf{q}(\mathbf{c}) = (q_1(\mathbf{c}), \ldots, q_n(\mathbf{c}))$ is a vector of probabilities $q_i(\mathbf{c}) \in [0,1]$ that suppliers invest; $\mathbf{y}(\mathbf{c}) = (y_1(\mathbf{c}), \ldots, y_n(\mathbf{c}))$ is a vector of ex ante payments $y_i(\mathbf{c}) \in \mathbb{R}$ to suppliers unconditional on success; and $\mathbf{x}(\mathbf{c}) = (x_1(\mathbf{c}), \ldots, x_n(\mathbf{c}))$ is a vector of ex post payments $x_i(\mathbf{c}) \in \mathbb{R}$ to suppliers conditional on success. Let $\mathbf{z}(\mathbf{c}) = (z_1(\mathbf{c}), \ldots, z_n(\mathbf{c}))$ denote the vector of expected payments to suppliers,

$$z_i(\mathbf{c}) = y_i(\mathbf{c}) + s x_i(\mathbf{c}). \tag{1}$$

Let $U_0(\mathbf{q}(\mathbf{c}))$ denote the funder's expected gross benefit from the mechanism. We have

$$U_0(\mathbf{q}(\mathbf{c})) = \left\{ 1 - \prod_{i=1}^{n} [1 - s q_i(\mathbf{c})] \right\} b. \tag{2}$$

The funder obtains benefit b if at least one supplier successfully completes investment, the probability of which is given by the factor in curly braces in equation (2). Assuming the funder does not internalize supplier profit, every dollar that the mechanism transfers from the funder to supplier counts as a dollar of lost funder surplus on net. The funder's expected net surplus, denoted $\Pi_0(\mathbf{c})$ thus equals

$$\Pi_0(\mathbf{c}) = U_0(\mathbf{q}(\mathbf{c})) - \sum_{i=1}^{n} z_i(\mathbf{c}), \tag{3}$$

which serves as the objective function maximized by the mechanism. Supplier i's expected profit is $\Pi_i(\mathbf{c}) = z_i(\mathbf{c}) - c_i q_i(\mathbf{c})$.

3.2. Foundational Result

Following Myerson (1981), Chaturvedi and Martínez-de-Albéniz (2011) impose the following condition on the distribution of agents' private information.

Definition. *The distribution of c_i is* regular *if and only if $F(c_i)/f(c_i)$ is nondecreasing.*

Regularity is equivalent to logconcavity of the cdf, as the next proposition, proved in Appendix A, states.

Proposition 1. *A distribution with pdf $f(c_i)$ and cdf $F(c_i)$ is regular if and only if $F(c_i)$ is logconcave.*

Most common distributions have logconcave cdfs, including uniform, normal, lognormal, exponential, Weibull, gamma, beta, chi-squared, Student's t, F distributions, and power functions (Bagnoli and Bergstrom, 2005, Tables 1 and 3). Proposition 1 then implies that these common distributions are regular, thus satisfying the needed condition to apply CM's result.

We can now state the result providing the foundation for our optimal mechanism.

Proposition 2. *Assume the distribution of c_i is regular. Then $(\mathbf{q}^*(\mathbf{c}), \mathbf{y}^*(\mathbf{c}), \mathbf{x}^*(\mathbf{c}))$ represents an optimal mechanism in dominant strategy equilibrium if and only if for all $\mathbf{c}$ it satisfies*

$$\mathbf{q}^*(\mathbf{c}) = \underset{\mathbf{q}\in[0,1]^n}{\operatorname{argmax}} \left[U_0(\mathbf{q}) - \sum_{i=1}^{n} q_i M(c_i) \right] \tag{4}$$

$$x_i^*(\mathbf{c}) + s y_i^*(\mathbf{c}) = z_i^*(\mathbf{c}) = c_i q_i^*(\mathbf{c}) + \int_{c_i}^{\bar{c}} q_i^*(\mathbf{c}_{-i}, t_i) dt_i, \tag{5}$$

where

$$M(c_i) = c_i + \frac{F(c_i)}{f(c_i)}. \tag{6}$$

Proposition 2 is nearly an exact transcription of CM's Theorem 1, which they derive as a corollary of Myerson (1981). Thus, little further proof is required. The short proof in Appendix A fills in the details needed to bridge the gap between the exact statement of CM's Theorem 1 and the slight restatement here.

3.3. Concrete Characterization

With Proposition 2 as a foundation, we can proceed to derive a more concrete characterization of the optimal mechanism for vaccine procurement. We will present the derivation of the optimal mechanism restricting investment to be deterministic, specifying an indicator $q_i(\mathbf{c}) \in \{0,1\}$ rather than a probability $q_i(\mathbf{c}) \in [0,1]$. It turns out that this restriction is without loss of generality since a deterministic mechanism is optimal in the more general class allowing for stochastic investment. Establishing the optimality of deterministic mechanisms is more involved and so relegated to the appendix.

With a deterministic mechanism, suppliers can be divided into winners, i.e., those selected to invest in return for a payment, and losers. Let $w(\mathbf{c}) = \sum_{i=1}^{n} q_i(\mathbf{c})$ denote the number of winning suppliers. Let $u(w)$ denote the funder's expected gross benefit when w suppliers invest. Then $U_0(\mathbf{q}(\mathbf{c})) = u(w(\mathbf{c}))$, implying

$$u(w) = [1-(1-s)^w]b \tag{7}$$

by equation (2). Defining the expected gross benefit provided by the marginal supplier as $\Delta(w) = u(w)-u(w-1)$, we have

$$\Delta(w) = s(1-s)^{w-1}b. \tag{8}$$

Furthermore, $u(w) = \sum_{i=1}^{w} \Delta(i)$, implying that $U_0(\mathbf{q}(\mathbf{c}))$ can be expressed as

$$U_0(\mathbf{q}(\mathbf{c})) = \sum_{i=1}^{w(\mathbf{c})} \Delta(i). \tag{9}$$

Since $U_0(\mathbf{q}(\mathbf{c}))$ depends on $\mathbf{q}(\mathbf{c})$ only via the sufficient statistic $w(\mathbf{c})$, it is apparent that the term subtracted in objective (4) can be minimized leaving $U_0(\mathbf{q})$ unchanged by selecting the $w(\mathbf{c})$ lowest-cost suppliers for the winners. To see this, note that $M(c_i)$ is an increasing function since, by (6), it is the sum of c_i, which is increasing, and $F(c_i)/f(c_i)$, which is nondecreasing when the distribution is regular. To provide notation for the selection of the lowest-cost suppliers, arrange supplier costs in ascending order, $c_{[1]} \leq c_{[2]} \leq \cdots \leq c_{[n]}$, using bracketed subscripts to indicate order statistics. Substituting from equation (9) and using the insight about the selection of lowest-cost suppliers, objective function (4) can be rewritten

$$\sum_{i=1}^{w^*(\mathbf{c})} \left[\Delta(i) - M(c_{[i]})\right], \tag{10}$$

where

$$w^*(\mathbf{c}) = \underset{w\in\{0,1,\ldots,n\}}{\operatorname{argmax}} \left\{ \sum_{i=1}^{w} \left[\Delta(i) - M(c_{[i]})\right] \right\}. \tag{11}$$

The solution to equation (11) is straightforward. Differentiating equation (8) shows that $\Delta(i)$ is decreasing in i. Further, $M(c_{[i]})$ is nondecreasing in i. Hence, the bracketed terms summed in (11) are decreasing in i. The solution involves adding bracketed terms in sequence until they become negative. Hence,

$$w^*(\mathbf{c}) = \max\left\{ i \in \{0,\ldots,n\} \middle| \Delta(i) \geq M(c_{[i]}) \right\}. \tag{12}$$

Equivalently, upon inverting,

$$w^*(\mathbf{c}) = \max\left\{ i \in \{0,\ldots,n\} \middle| M^{-1}(\Delta(i)) \geq c_{[i]} \right\}. \tag{13}$$

Equation (13) embodies the procedure for selecting the number of winning suppliers shown in Figure 1. The funder demand schedule, drawn as the light blue curve, is given by $\Delta(i)$. Applying the inverse function M^{-1} gives the reserve-price schedule. Equation (13) states that the optimal number of winning suppliers is given by the intersection of the reserve-price schedule $M^{-1}(\Delta(i))$, drawn as the dark blue curve, and the cost schedule $c_{[i]}$, drawn as the red curve.

It remains to derive a concrete expression for the expected award payment to winning suppliers in equation (5). Now $q_i^*(\mathbf{c}_{-i}, t_i) = 1$ if both $t_i < c_{[w^*(\mathbf{c})+1]}$ and $t_i < M^{-1}(\Delta(w^*(\mathbf{c})))$, while $q_i^*(\mathbf{c}_{-i}, t_i) = 0$ if either $t_i > c_{[w^*(\mathbf{c})+1]}$ or $t_i > M^{-1}(\Delta(w^*(\mathbf{c})))$. Thus,

$$\int_{c_i}^{\bar{c}} q_i^*(\mathbf{c}_{-i}, t_i) dt_i = \min\left\{ c_{[w^*(\mathbf{c})+1]}, M^{-1}(\Delta(w^*(\mathbf{c}))) \right\} - c_i. \tag{14}$$

Substituting (14) into (5) yields award

$$z^*(\mathbf{c}) = \min\left\{ c_{[w^*(\mathbf{c})+1]}, M^{-1}(\Delta(w^*(\mathbf{c}))) \right\} \tag{15}$$

for winning suppliers and 0 for losing suppliers. Equation (15) is the payment in a $[w^*(\mathbf{c})+1]$-price Vickrey auction with reserve price $M^{-1}(\Delta(w^*(\mathbf{c})))$. This award can be visualized using Figure 1 as the highest point in the intersection between the reserve-price and supplier-cost schedules.

Proposition 3. *Suppose the distribution of c_i is regular. The optimal funding mechanism in the vaccine model is a $w^*(c)$-winner Vickrey auction with reserve price $M^{-1}(\Delta(w^*(c)))$, where $w^*(\mathbf{c})$, is determined by equation* (13)*. Winners invest in return for expected payment $z^*(\mathbf{c})$, the lesser of the reserve price and the marginal losing bid.*

The mechanism has intuitive appeal. A (w^*+1)-price Vickrey auction serves to select w^* winners while inducing suppliers to bid their true costs c_i. The reserve price protects the funder from overpaying if marginal losing bidder ends up with a high cost draw. In the absence of private cost information, the reserve price would be set to $\Delta(w^*)$, ensuring the funder does not pay more than the marginal value provided by the marginal winning supplier. In the presence of private cost information, suppliers obtain an information rent, which reducing the reserve price to $M^{-1}(\Delta(w^*))$ helps extract.

As in CM, the division of the expected award $z_i^*(\mathbf{c})$ into ex ante ($y_i^*(\mathbf{c})$) versus ex post ($x_i^*(\mathbf{c})$) payments is not pinned down in our model. We will specify them as ex post payments to address the concern coming from outside the model that suppliers would default on a supply commitment if paid up front (among other moral-hazard and adverse-selection problems). To derive an ex post payment $x_i^*(\mathbf{c})$ from an expected payment $z_i^*(\mathbf{c})$, by equation (5), the latter must be scaled up by the reciprocal of the success probability: $x_i^*(\mathbf{c}) = z_i^*(\mathbf{c})/s$.

4. Estimating Supplier-Cost Distribution

Calculating the optimal number of winners and their award requires information on suppliers' costs for the relevant investment stages. Since calculation is based on taking simulated draws from the distribution of suppliers' costs, we require information on the shape (i.e., mean and variance) of this distribution.

4.1. Previous Literature

Previous estimates of cost distributions for various stages of pharmaceutical R&D do not directly apply to our context (DiMasi and Grabowski, 2007; DiMasi et al., 2016; Sertkaya et al., 2016; Gouglas et al., 2018). We are aware of no published manufacturing cost estimates for the stockpile phase. Existing cost estimates for phase-3 trials, with few exceptions (DiMasi et al., 2016; Gouglas et al., 2018), apply to drugs rather than vaccines (DiMasi and Grabowski, 2007; Sertkaya et al., 2016). Some estimates cover only clinical trial management (Sertkaya et al., 2016), not manufacturing, a key component in our context. Other estimates adjust for the risk of project failure (DiMasi and Grabowski, 2007; DiMasi et al., 2016). Since we directly account for failure in our simulations, risk adjusting suppliers' costs in this case would amount to double counting.

4.2. CEPI Survey Data

We provide new estimates of the distribution of the cost of relevant investment stages using proprietary data from a CEPI-sponsored survey first analyzed in Gouglas et al. (2018). This survey of hundreds of firms identified from clinical trial registries has been ongoing since 2017. Firms report costs for all vaccine development stages. Previous work by Gouglas et al. (2018) analyzed a subset of the data covering early stages (preclinical development, phase-1, and phase-2 trials). Our new

contribution here is to analyze survey responses for the later stages (phase-3 trials and stockpiling) and use that to estimate the distribution of supplier costs.

Table 1 provides descriptive statistics for the final sample used in our estimation. The investment costs have distinct distributions. The mean cost for phase-3 trials is $83.1 million and for stockpiling is $164.8 million. That the mean cost of phase-3 trials is much higher than the median suggests its distribution is highly positively skewed, but the relationship between the mean and median of the stockpiling cost suggests it may have the opposite skew. Though the cost of phase-3 trials has a lower mean cost than stockpiling, its standard deviation is much higher.

4.3. Lognormal Specification

We follow a long economics literature in assuming that costs in each investment stage have a lognormal distribution. It is standard in the economics literature to assume investment costs across suppliers have a lognormal distribution, a two-parameter distribution with scale parameter m and shape parameter v. Lognormality has been applied in contexts ranging from new drug entry (Ellison and Ellison, 2011) to manufacturing R&D (Bøler et al., 2015) and has been empirically tested in a cross study of seven industrial sectors across six countries (Lee, 2002).

Let c_{ti} denote supplier i's investment cost in stage $t \in \{1,2\}$. In particular, c_{1i} is i's cost of completing phase-3 trials, and c_{2i} is i's cost of manufacturing a stockpile of required size. Given that c_{ti} has a lognormal distribution, its pdf is given by

$$f_t(c_{ti}) = \frac{1}{\sqrt{2\pi}v_t c_{ti}} \exp\left(-\frac{(\ln c_{ti} - m_t)^2}{2v_t^2}\right) = \phi\left(\frac{\ln c_{ti} - m_t}{v_t}\right)\frac{1}{v_t c_{ti}}, \tag{16}$$

where m_t denotes the lognormal distribution's scale parameter, v_t its shape parameter, and ϕ the standard normal pdf. The lognormal distribution's mean μ_t and variance σ_t^2 can be expressed functions of the scale and shape parameters using standard formulas:

$$\mu_t = \exp\left(m_t + \frac{v_t^2}{2}\right) \tag{17}$$

$$\sigma_t^2 = \left[\exp(v_t^2) - 1\right]\exp\left(2m_t + v_t^2\right). \tag{18}$$

We estimate the best-fitting lognormal distribution for the two relevant stages using maximum

likelihood.[2] Table 2 presents the estimates. The resulting shapes of the estimated cost distributions are shown in Figure 2, the dashed red curve represents phase-3 trials and the dashed blue curve represents stockpiling. The dashed curves exhibit all of the distinct features of the distributions deduced from the descriptive statistics in Table 1: the cost of phase-3 trials has a lower mean, positive skew, and a wider variance than the cost of stockpiling.

4.4. Combining Stages

Since the model behind the optimal mechanism combines the two investment stages into a single project, the component lognormal cost distributions must be aggregated into one representing their sum. Assume c_{1i} and c_{2i} are independent. There is no general guarantee that the sum of lognormals, whether or not they are independent, is lognormal. Following the argument that the lognormal is an attractive functional form for investment-cost distributions, however, we will impose the lognormal form on the combined cost distribution. A decision remains, however, because there is no single accepted way of approximating a sum of lognormals with a lognormal. We adopt the Fenton (1960) approximation, which Beaulieu et al. (1995) find preforms better than leading alternatives under many circumstances.

The Fenton approximation matches the first two moments of the lognormal distribution. For random variables that are independent—as we assume c_{1i} and c_{2i} are—the sum of the means equals the mean of the sum (first moment) and the sum of the variances equals the variance of the sum (second moment). Let μ and σ^2 denote the mean and variance of the combined cost distribution and $\hat{\mu}$ and $\hat{\sigma}^2$ denote their respective estimates. The Fenton approximation sets

$$\hat{\mu} = \hat{\mu}_1 + \hat{\mu}_2 \tag{19}$$

$$\hat{\sigma}^2 = \hat{\sigma}_1^2 + \hat{\sigma}_2^2. \tag{20}$$

One can combine equations (17)–(20) to obtain estimates of the scale and shape parameters, $\hat{m}$ and $\hat{v}$, of the lognormal distribution for combined investment as functions of the parameters $\hat{m}_t$ and $\hat{v}_t$,

[2]The maximum-likelihood estimates can be computed directly via straightforward formulas. The scale-parameter estimate is simply the mean of the log-transformed cost, $\hat{m}_t = (1/n)\sum_{i=1}^{n} \ln c_{it}$, and the square of the shape-parameter estimate is its variance, $\hat{v}_t^2 = (1/n)\sum_{i=1}^{n} (\ln c_{it} - \hat{m}_t)^2$, where n denotes the sample size. We employ a statistical routine to obtain standard errors and log-likelihoods. Maximum-likelihood estimates of lognormal parameters are consistent but biased in small samples, with a negligible bias in samples as large as ours.

$t \in \{1,2\}$, estimated via maximum likelihood for the component cost distributions. We relegate the relevant formulas to Appendix B. Those formulas yield parameter estimates $\hat{m} = 5.377$ and $\hat{v} = 0.471$ for combined costs. By construction, the estimated mean for the combined distribution is $\hat{\mu} = 241.7$ and $\hat{\sigma} = 120.4$, as can be seen by substituting the derived moments from Table 2 into equations (19)–(20). All these estimates are measured in units of million dollars. The estimated lognormal distribution of combined costs is drawn as the solid black curve in Figure 2.

5. Simulation Results for Epidemic Scenario

Having characterized the optimal funding mechanism and empirically estimated the suppliers' cost distribution, it remains to specify three parameters—b, n, and s—to simulate performance of the mechanism. Since b will vary by context, we will preserve the generality of the analysis by reporting results for a range of b from \$500 million to \$3 billion. We begin with a baseline scenario in which $n = 5$ and $s = 0.5$, but later comparative-statics exercises will also analyze a range of values of those parameters.

For each parameter configuration, we conduct a million simulations. Each simulation draws a random cost for each of the n suppliers and feeds those into the optimal WVAR mechanism specified by Proposition 3 for the posited parameters and empirical cost distribution. As indicated in Figure 1, the WVAR mechanism arrives at the number of winners w^* as the intersection between the cost schedule (suppliers' announced costs arranged in increasing order) and the reserve-price schedule, which as shown in Section 3.3 equals $M^{-1}(\Delta(w))$. Although $M^{-1}(\Delta(w))$ does not have a closed-form solution, Appendix B provides details on the use of standard numerical methods to compute it.

5.1. Outcome Means

Table 3 reports a variety of outcome variables in the baseline scenario with $n = 5$ and $s = 0.5$. For the lowest reported value of b, \$500 million, funder net surplus averages \$82 million across simulations. Defining the surplus capture rate as funder net surplus as a proportion of b, we see that the surplus capture rate averages a mere 16%. The reason for this low capture rate is that the low benefit cannot support too high an award before funder net surplus becomes negative. The resulting low ex post award, averaging \$317 million, generates few winners. Indeed, in 14% of simulations, there is no

winner, pulling the mean number of winners below 1, to 0.89. With such few winners and uncertain success per winner, the stockpile is completed in only 44% of simulations.

For the highest reported value of b, \$3 billion, funder net surplus averages over \$1.9 billion, representing a 65% capture rate. The higher benefit justifies a higher award, averaging \$478 million, generating more winners, 2.58 on average. The award is sufficiently high that it essentially zeroes out the risk that b is foregone because no supplier enters. Indeed, only six of the million simulations generate no winner, which only results when the single-supplier reserve price fails to cover the lowest of five supplier cost draws. The stockpile is completed in 82% of the simulations for $b = \$3$ billion, nearly twice as many as for $b = \$500$ million.

Table 3 provides an indication of which feature of WVAR determines the optimal award. In a majority of simulations, the marginal losing bid determines the award. In roughly a third of the simulations generating a winner, the award is determined by the reserve price. Thus, both features of WVAR matter. The marginal losing bid seems to exert stronger discipline on the award than the reserve price, a finding confirmed below in more formal analysis.

5.2. Outcome Distributions

The focus thus far on means of outcome variables masks the fact that the optimal mechanism does not generate a certain outcome but a distribution of possibilities, including some that are quite disappointing. The realization of a disappointing outcome does not necessarily indicate the mechanism is a failure; it is a natural consequence of the uncertainty in costs and project success.

Figure 3 captures the distribution of outcomes in a series of histograms for various values of b. For brevity, the figure focuses on a single, key outcome variable: funder net surplus. The histogram in each panel shows a similar pattern. A mass of simulations generate no net surplus, whether because the mechanism does not select a winner to invest or because the selected winners do not end up being successful. Both events become less likely as b increases, evidenced by the shorter bars at zero funder net surplus. As b increases, the mechanism selects more winners, increasing the chance of a positive surplus realization.

The positive outcomes in each histogram are not massed on a single value but exhibit a distribution. From an ex post perspective, the highest realizations of funder net surplus occur when several suppliers draw low costs, resulting in a low award, but the pool of winning suppliers generates only one success, requiring only one ex post award payment to be made. While such an outcome is for-

tunate from an ex post perspective, the optimal mechanism from an ex ante perspective invites the possibility of multiple successes as a by-product of the effort to raise the probability of realizing at least one success and obtaining b. Lower but still positive realizations of funder net surplus can occur when suppliers draw higher costs or realize more successes. The distribution of positive outcomes shifts right as b increases. Part of the rightward shift is an automatic consequence of the funder's obtaining a greater benefit from successful stockpile completion, arising even if mechanism terms were not adjusted. But the terms of the mechanism adjust to increase the probability of successful completion, accounting for part of the rightward shift.[3]

5.3. Comparative Statics

The baseline scenario analyzed so far has fixed $n = 5$ and $s = 0.5$. Figures 4 and 5 explore the comparative-static effect of varying those parameters on a selection of outcome variables. Figure 4 varies the number of suppliers, n, from one to 20. Increasing n increases funder net surplus. More suppliers exert greater competitive pressure, which the mechanism exploits to lower the award while at the same time increasing the number of winners and the probability of stockpile completion. The marginal effect of adding a supplier falls with n. Starting from one supplier, adding another produces a big jump in outcomes; but when there are 10, the changes from adding 10 more are fairly negligible.

Our previous discussion of Table 3 already hinted at the comparative-static effect of an increase in b, since that table reported results for various b. However, the graphs in Figure 4 allow for better visualization of the comparative-static effects. We see in the first panel that funder net surplus is increasing and convex in b. The convex shape is due to two forces. First, holding constant the terms specified by the mechanism, an increase in b scales up benefits but not costs, so the net increases more than proportionately. Second, the mechanism terms do not remain constant as b increases but adjust to increase the number of winners and probability of completion, adding further to funder net surplus.

Figure 5 is similar to Figure 4, displaying the same outcomes in the analogous panels, but varies the success probability s rather than the number of suppliers n. For ease of comparison, the same curve with the medium tone is maintained in both figures, representing the benchmark scenario of

[3] In addition to the spike at zero funder net surplus, the histograms exhibit a spike among positive outcomes. This spike reflects a set of simulations in which the award payment is set by the reserve price for a range of cost draws.

$n = 5$ and $s = 0.5$. The mechanism performs poorly for a success probability as low as $s = 0.05$. Funder net surplus is close to zero even for the highest reported value of b. The ex post award has to be high since it is scaled by the reciprocal of a probability that is close to zero. Even with high awards, the mechanism rarely generates any winners, and the probability of stockpile completion is close to zero.

At the other extreme is the curve for $s = 0.95$. With investment nearly certain to produce success, there is less need for redundancy. The outcomes from the optimal mechanism bear this out. A modest ex post award is enough to generate one winner, which is what the optimal mechanism selects in nearly all simulations and nearly always results in stockpile completion. In essence, the WVAR reverts into a standard, second-price Vickrey auction with one winner.

Looking at values of s between the extremes, the increase from $s = 0.5$ to $s = 0.75$ has a larger effect on funder net surplus than the increase from $n = 5$ to $n = 10$ in the previous figure.

The effect of s on number of winning suppliers in Panel C is nonmonotonic. For very low s, the mechanism typically selects no winner. For very high s, the mechanism typically selects exactly one winner. The number of winners is greatest for intermediate s, high enough to justify selecting some winners, yet not so high that any value of redundancy is eliminated.

5.4. Alternative Mechanisms

Table 4 compares the outcome from the WVAR mechanism to alternative mechanisms, both more and less efficient. To facilitate easy comparison, the table uses the mean across simulations of the surplus capture rate as its index of performance. Recall the surplus capture rate equals funder net surplus as a proportion of b.

The first column of results shows the surplus capture rate in the first-best mechanism: the mechanism maximizing funder net surplus when it has complete information about suppliers' costs and can offer discriminatory awards based on this information. The first best can be understood with reference to Figure 1. The intersection between the suppliers' cost schedule and the funder's demand schedule determines the first-best number of winning suppliers, w^{**}. The funder pays each winning supplier $[i] \leq w^{**}$ an ex post award equal to its cost $c_{[i]}$ scaled by the reciprocal of the probability of success s.[4] The first-best mechanism can be equivalently viewed as a perfect push mechanism,

[4]It is intuitive that the number of winners maximizing the funder's net surplus in the first best should maximize net social surplus. The funder's benefit constitutes social benefit; since its payments to suppliers in the first best lead it to internalize social costs, the number of winners maximizing the funder's net surplus should maximize net social surplus.

which makes ex ante payments to each of the w^{**} winning suppliers to exactly defray each one's cost $c_{[i]}$.

Table 4 shows that the first-best outperforms WVAR, capturing a higher rate of surplus for all reported all levels of b. This is not surprising since WVAR, the second-best mechanism, operates under a private-information constraint absent in the first best. WVAR cannot avoid offering winning suppliers an information rent, lowering funder net surplus. The gap between the first and second best falls in percentage terms as b rises. The reason is that the information rent is related to suppliers' cost, which does not scale with b, so falls as a percentage of b as b increases. At the highest reported value of b, the surplus capture rate under WVAR (0.65) is over 90% of the rate under the first best (0.70).

The remaining columns in Table 4 report the performance of mechanisms that are impaired in one way or other relative to WVAR, so have a lower rate of surplus capture. Retaining the reserve price but dispensing with the marginal losing bid as a constraint on the award reduces the surplus capture rate. The reduction is more substantial for higher values of B. Higher values of B are associated with more generous reserve prices and thus depend even more on the competitive pressure provided by the marginal losing bid to restrain the award. This finding is consistent with the result in Table 3 that the marginal losing bid was the effective constraint on the award in a majority of simulations.

Next, we consider a mechanism that dispenses with the reserve price but reinserts the marginal losing bid as a constraint on the award. Without a reserve-price schedule, there is no longer a well-defined intersection between it and the supplier-cost schedule determining the optimal number of winners. As a substitute, we analyze different predetermined numbers of winners from one to four.[5] If the funder guesses well and ends up predetermining a number of winners close to the mean selected by the WVAR, the surplus capture rate can come close to that under WVAR. This is not a foregone conclusion because the ability to flexibly tailor the number of winners to announced costs rather than predetermining the number in principle is a valuable mechanism-design element. While WVAR beats any of the alternative mechanisms predetermining the number of winners, it only beats by one or two percentage points mechanisms that predetermine a number of winners close to

To see this more formally, note funder demand is the schedule $\Delta(w)$. At the w^{**} determined by the intersection between funder-demand and supplier-cost schedules, $\Delta(w^{**}) \geq c_{[w^{**}]}$, implying that the marginal winning supplier generates weakly greater social benefit than social cost. Further, $\Delta(w^{**}+1) < c_{[w^{**}+1]}$, implying that another supplier would add more to social cost than social benefit.

[5]We stop at four because, without a reserve price to constrain the award, having more than four winners in the baseline scenario with $n = 5$ would result in an infinite equilibrium award.

WVAR's mean across simulations. The main benefit of the reserve price is to prevent the funder from making the wrong choice in predetermining the number of winners. Too many or too few and the surplus capture rate can fall substantially, in some cases becoming negative. For example, when $b = \$500$ million, fixing the number of winners at $w = 4$ generates surplus capture rate of –1.22, which translates into a mean loss of about \$610 million. When $b = \$3$ billion, fixing the number of winners at $w = 1$ leads to a surplus capture rate of 0.44, substantially lower than 0.65 under WVAR.

6. Pandemic Scenario

WVAR is a general mechanism for optimally funding vaccines against epidemics ranging from local outbreaks to global pandemics. We have tried to keep the analysis general, but for concreteness have made some choices in the analysis that ended up tailoring it to the case of a moderate outbreak. The benefits of quelling a pandemic would be much higher than the \$500 million to \$3 billion used for b in the simulations in the previous section. Other than making the figures less legible, there would be no difficulty in extending the figures' axes to higher values of b. The main difficulty in automatically rescaling the simulations up to pandemic proportions is that the cost functions in Figure 2 are estimated using data from a survey of the costs of producing a moderate stockpile. Orders of magnitude more vaccine capacity and output would be needed to combat a pandemic, requiring costs to be rescaled, if not reestimated, to provide realistic values at a pandemic scale.

The analysis was further tied to a moderate outbreak via a subtle modeling choice. We assumed that any successful supplier would be able to build sufficient capacity to complete the required stockpile on its own. The full benefit b is realized as long as one supplier is successful. Multiple successes provide no benefit beyond an increased probability that b is realized. The assumption that one supplier can fully satisfy program needs would likely be strained in a pandemic. A campaign to reach herd immunity by vaccinating the majority of a large country's or the world's population would require enormous capacity to produce billions of vaccine courses, not a moderate stockpile. Successful suppliers beyond the first could contribute their additional capacity to vaccinating more people and more quickly, increasing social benefits compared to one success.

This section discusses our contemporaneous work in Snyder et al. (2020) that applied the WVAR mechanism to funding vaccines in the Covid-19 pandemic. We complement that non-technical piece here by providing formal details on extensions to the model to accommodate the possibility that

benefits increase in the aggregate capacity of successful suppliers. We outline the steps Snyder et al. (2020) took to estimate the benefit and cost functions, summarize the key simulation results, and contrast those results with the simulation results for a moderate epidemic from the previous section.

6.1. Generalizing Model

The model from Section 3.1 needs to be generalized slightly to accommodate the case of a pandemic. Let $\bar{k}$ be the capacity constraint for an individual supplier, i.e., the maximum output it can produce over specified period of time, here taken to be the duration of the pandemic. To simplify the analysis, we assume that each successful supplier produces this amount of vaccine.

Let $B(K)$ denote the funder's gross surplus from K aggregate units of output across suppliers. In effect, the model of Section 3.1 took $B(K)$ to have the specific form

$$B(K) = \begin{cases} b/K & K \in [0, \bar{k}) \\ b & K > \bar{k}. \end{cases} \tag{21}$$

This section generalizes $B(K)$, allowing it to be any nonnegative, weakly increasing, weakly concave function.

The function $u(w)$, representing the funder's expected gross benefit when w suppliers invest, previously given by equation (7), becomes

$$u(w) = \sum_{i=1}^{w} \beta(i, w, s) B(i\bar{k}), \tag{22}$$

where

$$\beta(i, w, s) = \frac{w!}{i!(w-i)!} s^i (1-s)^{w-i} \tag{23}$$

is the binomial probability function, equal to the probability of i successes from w draws from a Bernoulli distribution with probability of success s.

Letting $U_0(\mathbf{q}(\mathbf{c})) = u(w(\mathbf{c}))$, the analysis proceeds exactly as in Section 3.3. It is immediate that the WVAR mechanism continues to be optimal in this generalized model as well.

6.2. Estimating Pandemic Benefits and Costs

To estimate gross benefits from a Covid-19 vaccine at the global level, Snyder et al. (2020) combined estimates of economic output losses and mortality losses for each country in the world. Economic output losses were based on International Monetary Fund (2020) projections of growth decline in countries caused by Covid-19. Mortality losses were based on projections by the Imperial College Covid-19 Response Team of deaths by country (Walker et al., 2020). Each death was translated into 12 years of lost life (YLLs), the average estimated by Hanlon et al. (2020) accounting for age and comorbidities. A YLL was valued at three times the per-capita GDP in the country in which the life was lost, derived from the World Health Organization's standards for cost-effective health interventions.[6]

Assuming avoided harm in a country is proportional to vaccination coverage up to complete avoidance at 100% coverage, individual countries can be ordered in terms of harm severity to generate a global demand curve for avoiding Covid-19 harm. Snyder et al. (2020) allowed for heterogeneity within countries in the form of vulnerable populations suffering greater than average harm, but the baseline assumes homogeneity within countries. The authors generated a monthly version of the demand curve for avoided harm by dividing total harm by the duration of the pandemic in months and further assuming this demand declines 5% each month reflecting the possible arrival of a treatment or herd immunity ending the pandemic early. Assuming vaccines provide durable protection over the remainder of the pandemic, monthly demand for avoided harm can be translated into global demand for a vaccine. Vaccine demand constructed in this way declines as the pandemic wanes as fewer months remain during which protection is warranted. Vaccine demand provides the needed schedule of marginal vaccine benefits.

Snyder et al. (2020) estimated costs from an updated CEPI survey, targeted to potential Covid-19 vaccine suppliers. Suppliers were asked about the investment costs for late-stage clinical trials and reservation of capacity for large-scale manufacturing as well as production costs. Ex ante investments were assumed to be undertaken at risk (before the candidate is known to be successful in clinical trials) to accelerate the arrival of the vaccine, but production costs are only expended

[6]The World Health Organization regards a health intervention as cost effective if the cost per disability adjusted life year (DALY) saved is less than three times that country's per-capita GDP (Marseille et al., 2015). Since Hanlon et al.'s (2020) estimate of YLLs already allocates shorter lifespans to people with comorbidities, one YLL was assumed to translate into one DALY without need for further downward adjustment to reflect a proportion of years lived with a disability.

conditional on success.

The authors used this survey data to estimate lognormal distributions for investment and production costs for 750 million doses of a Covid-19 vaccine. The combined investment cost was not far from that estimated in the epidemic scenario in Section 4. The two scenarios diverge regarding production costs. Whereas the epidemic scenario folded the cost of manufacturing the moderate stockpile into investment costs, production costs are accounted for separately in the pandemic scenario and dominate cost calculations. The mean \$13.3 cost per dose scales to a \$9.98 billion production cost when multipled by the 750 million doses a successful firm is assumed to produce. The standard deviation of production costs is quite high at \$13.5 per dose, leading to a wide cost distribution, endowing suppliers with considerable private information.

6.3. Results

Table 5 reports the mean of selected outcomes from simulations of the optimal mechanism in Covid-19 scenarios for various parameters. The first three rows excerpt results reported in Exhibit 3 of Snyder et al. (2020). None of the parameter sets exactly matches the baseline established in this paper for the epidemic scenario. Thus, for ease of comparison, we ran an additional set of simulations inputting $n = 5$ potential suppliers and probability of success $s = 0.5$, the parameter values used in the benchmark of the epidemic scenarios in this paper. The results are reported in the last row of the table.

In all rows, the optimal mechanism selects all available firms to invest in virtually all simulations. Focusing on the last row, on average across the million simulations, 4.9 suppliers are selected. The enormous costs of producing at pandemic scale are eclipsed by the vastly higher social benefits. By contrast, in simulations of the baseline epidemic scenario reported in Section 5, even considering the highest reported value of b ($b = \$3$ billion), only about half the five potential firms are selected as winners, eligible to receive an award conditional on success.

The mean ex post award in the last row of Table 5 is \$39.6 billion. The award is this large for two reasons. First, the mean cost of producing at pandemic scale is enormous, around \$10 billion. In addition, all five suppliers are nearly always selected, meaning that the award is nearly always set by the reserve price rather than being disciplined by competition, further inflating program expense. While the mean award in the baseline epidemic scenario at \$478 million is only about twice the mean combined cost, in the pandemic scenario, the mean award is four times the mean of combined

investment and production costs.

Total program expenditure averages averages $2.6 trillion across simulations in the last row, more than a thousand times the $1.9 billion in the epidemic scenario with the highest reported b. While we could have inputted a thousand times higher value of b into the epidemic scenario, we chose more moderate values of b to approximate the value of quelling an emergent epidemic perhaps in a foreign country that may only have a tiny chance of becoming a global concern. The value of quelling an epidemic that has already grown to pandemic proportions easily rises into the trillions, as Table 5 suggests.

7. Conclusion

In this paper, we derived the optimal funding mechanism for a vaccine against an emergent outbreak. The mechanism can cover a variety of settings, including one country's procurement of a vaccine to help quell outbreaks abroad, whether out of altruism or out of self interest, to prevent the possible spread of the disease to the home country. Alternatively, the mechanism could be part of a multilateral procurement such as five finance ministries undertook to increase global coverage of second-generation pneumococcal vaccines (Kremer et al., 2020) or the COVAX mechanism to procure Covid-19 vaccines for participating countries.

The optimal mechanism, which we dubbed WVAR, is a $(w+1)$-price reverse Vickrey auction with reserve price. It combines the procurer's internal assessment of the social benefit of the vaccine with suppliers' cost announcements to arrive at the number of suppliers selected to participate and the award that eligible suppliers who are successful should receive. The higher the vaccine benefits, the more eligible participants that are selected by the mechanism, which it induces to participate with a higher award. The higher suppliers' cost draws, the less likely any given number of them have costs below the reserve threshold, leading fewer of them to be selected to participate.

We peformed two sets of calibrations, one set for a model and benefits tailored to a moderate outbreak and another set for a model and benefits tailored to a pandemic like Covid-19. Though the underlying mechanism remained constant, the different inputs into it led to different outcomes. For a moderate outbreak, avoiding excessive expenditures is a key concern, achieved by setting a moderate award that often winnowed out several of the highest-cost suppliers. In pandemic scenarios, with benefits in the trillions not billions, the mechanism induced the participation of virtually every

available supplier with little concern for expenditures, confirming the intuition that "all hands on deck" is optimal in a global crisis of that scale.

In our analysis of alternative mechanisms, we found that WVAR captured an increasing share of first-best surplus the higher the benefit b considered. As b increases, the relative importance of the asymmetric-information "overhead" falls. For example, in calibrations for a moderate epidemic with b set to \$3 billion, WVAR obtained 90% of net first-best surplus. The design of the optimal mechanism matters: eliminating certain features of WVAR substantially impairs its performance. For example, substantial net surplus is lost if the marginal losing bid is dispensed with as a cap on awards. If the procurer fixes the number of participating firms rather than letting the number be determined endogenously by the intersection of the cost and reserve-price schedules, the performance of the resulting mechanism is mixed. If the procurer somehow manages to guess the number that is correct on average across simulations, the outcome will be close to the optimum. However, without the formulas provided, it may be hard to guess that the correct number of winning bidders averages two when $s = 0.5$, $n = 5$, and $b = 2000$ but rises to three keeping the parameters the same but increasing b to 3,000.

The mechanism is agnostic concerning how the award should be structured, whether a lump-sum payment made unconditional on success, a scaled-up lump sum conditional on success, or yet a further scaled-up lump sum to the first successful supplier to market. If the endpoint is a stockpile of a given quantity, the award can equivalently be construed as a lump sum or a per-dose price. A danger of setting the award as a per-dose subsidy is that it may provide perverse incentives in settings outside of the model in which quantity is discretionary. Then, a very effective vaccine may end up quickly quelling an epidemic with few doses, penalizing the supplier for efficacy. In such applications, a lump-sum award tied in part to the social value of quelling the epidemic may work better.

Appendix A. Proofs of Propositions

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

Proof of Proposition 1

Suppose a distribution has pdf $f(c_i)$ and cdf $F(c_i)$. Assume $F(c_i)$ is logconcave. Then

$$\frac{d^2 \ln F(c_i)}{dc_i^2} = \frac{f'(c_i)F(c_i) - f(c_i)^2}{F(c_i)^2} < 0, \tag{A1}$$

in turn implying $f(c_i)^2 - f'(c_i)F(c_i) > 0$. But the preceding inequality implies

$$\frac{d}{dc_i}\left[\frac{F(c_i)}{f(c_i)}\right] = \frac{f(c_i)^2 - f'(c_i)F(c_i)}{f(c_i)^2} > 0, \tag{A2}$$

proving that the cost distribution is regular. The arguments can be reversed to prove that regularity implies logconcavity. *Q.E.D.*

Proof of Proposition 2

CM's model nests every dimension of ours except for two. First, CM assume $U(\mathbf{q})$ is concave in $\mathbf{q}$, but our specification in equation (2) is neither concave nor convex. One can show that CM's proof of Theorem 1 does not hinge on the concavity assumption, which is only material for some of their subsequent results. Second, we construe $q_i(\mathbf{c})$ as probabilities, bounding them above by 1, whereas CM construe them as quantities with no explicit upper bound. Again, one can show that the constrained choice set does not materially affect their proof.

Along other dimensions CM's model differs from ours, their model nests ours. In particular, CM allow for supplier heterogeneity, writing success probabilities s_i and cost pdfs $f_i(c_i)$ with a subscript i. Our restatement of CM's Theorem 1 here, which drops those subscripts, follows immediately from their more general result. *Q.E.D.*

Proof of Proposition 3

It remains to show that the optimal mechanism allowing for stochastic investment can be taken to be deterministic without loss of generality . We will do this by showing that the relevant objective function is linear in investment probabilities q_i and thus has a corner solution at either 0 or 1, i.e., is deterministic.

We proceed by redefining some of the ancillary functions introduced in Section 3, allowing them to reflect stochastic investment. The main notational complication arising with stochastic investment is that it is no longer sufficient to keep track of the number of winning suppliers since winning suppliers may have different investment probabilities, which have to be tracked. Let $u(\ell, \mathbf{q})$ denote the funder's gross benefit derived from the ℓ lowest-cost suppliers in the mechanism, excluding higher-cost suppliers:

$$u(\ell, \mathbf{q}) = b\left[1 - \prod_{i=0}^{\ell}(1 - sq_{[i]})\right], \tag{A3}$$

defined for $\ell \in \{0,1,\ldots,n\}$, setting $q_{[0]} = 0$ by definition. Let $\Delta(\ell,\mathbf{q})$ denote the marginal contribution of the highest-cost of the included firms to the gross benefit in equation (A3):

$$\Delta(\ell,\mathbf{q}) = u(\ell,\mathbf{q}) - u(\ell-1,\mathbf{q}) = bsq_{[\ell]}\prod_{i=0}^{\ell-1}(1-sq_{[i]}). \tag{A4}$$

Let $\hat{\Delta}(\ell,\mathbf{q})$ be the same marginal contribution but reflecting the marginal firm's investing with certainty rather than with probability $q_{[\ell]}$ specified by the mechanism:

$$\hat{\Delta}(\ell,\mathbf{q}) = bs\prod_{i=0}^{\ell-1}(1-sq_{[i]}). \tag{A5}$$

The two marginal contributions are related by $\Delta(\ell,\mathbf{q}) = q_{[\ell]}\hat{\Delta}(\ell,\mathbf{q})$.

With those ancillary functions in hand, we can write

$$U_0(\mathbf{q}) = u(n,\mathbf{q}) = \sum_{i=1}^{n}\Delta(i,\mathbf{q}) = \sum_{i=1}^{n} q_{[i]}\hat{\Delta}(i,\mathbf{q}). \tag{A6}$$

Substituting into objective function (4) and rearranging yields

$$\mathbf{q}^*(\mathbf{c}) = \underset{\mathbf{q}\in[0,1]^n}{\operatorname{argmax}} \sum_{i=1}^{n} q_{[i]}\left[\hat{\Delta}(i,\mathbf{q}) - M(c_{[i]})\right]. \tag{A7}$$

We will show that the maximand in (A7) is linear in the investment probabilities.

Consider an arbitrary investment probability, $q_{[\ell]}$. The sum in (A7) can be expanded, emphasizing three groups of relevant terms as follows:

$$\sum_{i=1}^{\ell-1} q_{[i]}\left[\hat{\Delta}(i,\mathbf{q}) - M(c_{[i]})\right] + q_{[\ell]}\left[\hat{\Delta}(\ell,\mathbf{q}) - M(c_{[\ell]})\right] + \sum_{i=\ell+1}^{n} q_{[i]}\left[\hat{\Delta}(i,\mathbf{q}) - M(c_{[i]})\right]. \tag{A8}$$

The first summation is not a function of $q_{[\ell]}$. The middle term is the product of $q_{[\ell]}$ and the factor in square brackets, which is not a function $q_{[\ell]}$. Thus, the middle term is linear in $q_{[\ell]}$. Substituting from (A5) and rearranging, the last term can be shown to equal

$$bs(1-q_{[\ell]})\left[\sum_{\ell+1}^{n} q_{[i]} \prod_{\substack{j\le i-1 \\ j\neq\ell}} (1-sq_{[j]})\right] - \sum_{\ell+1}^{n} M(c_{[i]}). \tag{A9}$$

Since the factor in square brackets is not a function of $q_{[\ell]}$, we have that (A9) is linear in $q_{[\ell]}$. *Q.E.D.*

Appendix B. Additional Formulas

This appendix provides additional formulas omitted from the text for space considerations.

Fenton Approximation of Combined Distribution

Substituting from equations (17) and (18) into (19) and (20) and solving yields

$$\hat{v} = \sqrt{\ln\left(1 + \frac{\hat{\xi}_1}{\hat{\xi}_0^2}\right)} \tag{B1}$$

$$\hat{m} = \ln\hat{\xi}_0 - \frac{\hat{v}^2}{2}, \tag{B2}$$

where

$$\hat{\xi}_0 = \exp\left(\hat{m}_1 + \frac{\hat{v}_1^2}{2}\right) + \exp\left(\hat{m}_2 + \frac{\hat{v}_2^2}{2}\right) \tag{B3}$$

$$\hat{\xi}_1 = [\exp(\hat{v}_1^2) - 1]\exp(2\hat{m}_1 + \hat{v}_1^2) + [\exp(\hat{v}_2^2) - 1]\exp(2\hat{m}_2 + \hat{v}_2^2). \tag{B4}$$

Inversion for Reserve Price

Letting Δ be funder demand for the marginal winning firm's investment, the associated reserve price is given by $M^{-1}(\Delta)$, the inverse of the function $M(c)$ defined in equation (13). Lacking a closed-form solution for this inverse, we resort to numerical methods. Letting $c = M^{-1}(\Delta)$, we have $M(c) = \Delta$, implying

$$\psi(c, \Delta) = 0, \tag{B5}$$

defining $\psi(c, \Delta) = M(c) - \Delta$. Since $M(c)$ is increasing, $\psi(c, \Delta)$ is increasing in c as well.

A variety of numerical methods can be used to solve for c in equation (B5). We will use the Newton-Raphson method, which involves iterating

$$c_{k+1} = c_k - \frac{\psi(c_k, \Delta)}{\psi_1(c_k, \Delta)} \tag{B6}$$

until convergence. In equation (B6), ψ_1 denotes the partial derivative with respect to its first argument,

$$\psi_1(c, \Delta) = M'(c) = \frac{2f(c)^2 - f'(c)F(c)}{f(c)^2}. \tag{B7}$$

For the lognormal, one can differentiate equation (16) and rearrange to show

$$f'(c) = f(c)h(c), \tag{B8}$$

where

$$h(c) = \frac{m - \ln c - v^2}{cv^2}. \tag{B9}$$

Substituting equations (B7)–(B9) into (B6) and rearranging yields

$$c_{k+1} = \frac{c_k f(c_k) - [c_k h(c_k) + 1]F(c_k) + f(c_k)\Delta}{2f(c_k) - h(c_k)F(c_k)}. \tag{B10}$$

Equations (16) and (B9) provide closed-form solutions for $f(c_k)$ and $h(c_k)$.

References

Acemoglu, A. and J. Linn. (2004) "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *Quarterly Journal of Economics* 119 (3): 1049–1090.

Ahuja, A., S. Athey, A. Baker, E. Budish, J. C. Castillo, R. Glennerster, et al. (2021) Preparing for a Pandemic: Accelerating Vaccine Availability," *American Economic Association Papers and Proceedings* 111: 331–335.

Athey, S., M. Kremer, C. M. Snyder, and A. Tabarrok. (2020) "In the Race for a Coronavirus Vaccine, We Must Go Big. Really, Really Big. *New York Times*, May 4. [Internet.] Available at https://www.nytimes.com/2020/05/04/opinion/coronavirus-vaccine.html.

Athey, S., J. C. Castillo, E. Chaudhuri, M. Kremer, A. Simoes Gomes, and C. M. Snyder. (2022) "Expanding Capacity for Vaccines Against Covid-19 and Future Pandemics: A Review of Economic Issues," *Oxford Review of Economics and Policy* 34 (4): forthcoming.

Bagnoli, M. and T. Bergstrom. (2005) "Log-concave Probability and its Applications," *Economic Theory* 26 (2): 445–469.

Baker, S., S. Koons, and D. Kocieniewksi. (2020) "Here. It. Comes." *Bloomberg Businessweek*, November 2, 42—47.

Beaulieu, N. C., A. A. Abu-Dayya, and P. J. McLane. (1995) "Estimating the Distribution of a Sum of Independent Lognormal Random Cariables. *IEEE Transactions on Communications* 43 (12): 2869–2873.

Bøler, E. A., A. Moxnes, and K. H. Ulltveit-Moe. (2015) "R&D, International Sourcing, and the Joint Impact on Firm Performance," *American Economic Review* 105 (2): 3704–3739.

Butler, D. (2017) "Billion-dollar Project Aims to Prep Vaccines Before Epidemics Hit," *Nature News* 541 (7638): 444–445.

Castillo, J. C., A. Ahuja, S. Athey, A. Baker, E. Budish, T. Chipty, et al. (2021) "Market Design to Accelerate COVID-19 Vaccine Supply," *Science* 371 (6534): 1107-–1109.

Chari, V. V., M. Golosov, M., and A. Tsyvinski. (2012) "Prizes and Patents: Using Market Signals to Provide Incentives for Innovations," *Journal of Economic Theory* 147 (2): 781–801.

Chaturvedi, A. and V. Martínez-de-Albéniz. (2011) "Optimal Procurement Design in the Presence of Supply Risk," *Manufacturing & Service Operations Management* 13 (2): 227—243.

Childs, P. D. and A. J. Triantis. (1999) "Dynamic R&D Investment Policies," *Management Science* 45 (10): 1359–1377.

Dahir, A. L. (2022) "Uganda Races to Contain a Deadly Ebola Outbreak," *New York Times*, October 1.

DiMasi, J. A. and H. G. Grabowski. (2007) "The Cost of Biopharmaceutical R&D: Is Biotech Different?" *Managerial and Decision Economics* 28 (4–5): 469–479.

DiMasi, J. A., H. G. Grabowski, and R. W. Hansen. (2016) "Innovation in the Pharmaceutical Industry: New Estimates of R&D Costs," *Journal of Health Economics* 47: 20–33.

Doughman, E. (2020) "Global Vaccine Market Revenue to Reach $59.2 Billion by 2020," *Pharmaceutical Processing World*, May 20. Available from https://www.pharmaceuticalprocessingworld.com/global-vaccine-market-revenue-to-reach-59-2-billion-by-2020.

Ellison, G. and S. F. Ellison. (2011) "Strategic Entry Deterrence and the Behavior of Pharmaceutical Incumbents Prior to Patent Expiration," *American Economic Journal: Microeconomics* 3 (1): 1–36.

Fan, V. Y., D. T. Jamison, and L. H. Summers. (2018) "Pandemic Risk: How Large Are the Expected Losses?" *Bulletin of the World Health Organization* 96 (2): 129–134.

Fenton, L. F. (1960) "The Sum of Lognormal Probability Distributions in Scatter Transmission," *IEEE Transactions on Communication Systems* 8 (1): 57–67.

Galasso, A., M. Mitchell, and G. Virag. (2016) "Market Outcomes and Dynamic Patent Buyouts," *International Journal of Industrial Organization* 48: 207–243.

Geoffard, P.-Y. and T. Philipson. (1997) "Disease Eradication: Private Versus Public Vaccination," *American Economic Review* 87 (1): 222–230.

Gouglas, D., T. T. Le, K. Henderson, A. Kaloudis, T. Danielsen, N. C. Hammersland, J. M. Robinson, P. M. Heaton, and J.-A. Røttingen. (2018) "Estimating the Cost of Vaccine Development Against Epidemic Infectious Diseases: A Cost Minimization Study," *Lancet Global Health* 6 (12): e1386–1396.

Gross, D. P. and B. V. Sampat. (2022) "Crisis Innovation: Historical Evidence, Insights, and Open Questions," *NBER Reporter*, October (3), 12–15.

Guyton, J., H. McClellan, and F. Li. (2010) "Synthesis of Business Models and Economic and Market Incentives for Vaccines and Therapeutics," in *The Public Health Emergency Medical Countermeasures Enterprise: Innovative Strategies to Enhance Products from Discovery through Approval*, Appendix E. Washington DC: National Academies Press.

Hanlon, P., F. Chadwick, A. Shah, R. Wood, J. Minton, G. McCartney, G. et al. (2020) "COVID-19—Exploring the Implications of Long-term Condition Type and Extent of Multimorbitity on Years of Life Lost: A Modelling Study," *Wellcome Open Research* 5:75 [Preprint]. April 23. Available from https://wellcomeopenresearch.org/articles/5-75.

International Monetary Fund. (2020) *World Economic Outlook, April 2020: The Great Lockdown.* Washington DC: International Monetary Fund staff report. Available from https://www.imf.org/en/Publications/WEO/Issues/2020/04/14/weo-april-2020.

Kremer, M. (1998) "Patent Buyouts: A Mechanism for Encouraging Innovation," *Quarterly Journal of Economics* 113 (4): 1137–1167.

Kremer, M. (2001) "Creating Markets for New Vaccines Part I: Rationale," in Innovation Policy and the Economy, vol. 1. Cambridge MA: MIT Press, pp. 35–72.

Kremer, M. and R. Glennerster. (2004) *Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases*. Princeton, New Jersey: Princeton University Press.

Kremer, M., J. Levin, and C. M. Snyder. (2020) "Advance Market Commitments: Insights from Theory and Experience," *American Economic Association Papers and Proceedings* 110: 269–273.

Kremer, M., J. Levin, J., and C. M. Snyder. (2022) "Designing Advance Market Commitments for New Vaccines," *Management Science* 68 (7): 4786–4814.

Kremer, M. and C. M. Snyder. (2015) "Preventives Versus Treatments," *Quarterly Journal of Economics* 130 (3): 1167—1239.

Lee, C.-Y. (2002) "Industry R&D Intensity Distributions: Regularities and Underlying Determinants," *Journal of Evolutionary Economics* 12 (3): 307–341.

Marseille, E., B. Larson, D. S. Kazi, J. G. Kahn, and S. Rosen. (2015) "Thresholds for the Cost-effectiveness of Interventions: Alternative Approaches," *Bulletin of the World Health Organization* 93 (2): 118–124.

Myerson, R. B. (1981) "Optimal Auction Design," *Mathematics of Operations Research* 6 (1): 58–73.

Ridley, D. B., H. G. Grabowski, and J. L. Moe. (2006) "Developing Drugs for Developing Countries," *Health Affairs* 25 (2): 313—324.

Sagonowsky, E. (2019) "Following Merck Approval, Gavi Authorizes $178M to Build Ebola Vaccine Stockpile," *FiercePharma*, December 6.

Scotchmer, S. (1999) "On the Optimality of the Patent Renewal System," *RAND Journal of Economics* 30 (2): 181–196.

Sertkaya, A., H.-H. Wong, A. Jessup, and T. Beleche. (2016) "Key Cost Drivers of Pharmaceutical Clinical Trials in the United States," *Clinical Trials* 13 (2): 117–126.

Snyder, C. M., K. Hoyt, D. Gouglas, T. Johnston, and J. Robinson. (2020). "Designing Pull Funding for a COVID-19 Vaccine," *Health Affairs* 39 (9): 1633–1642.

Snyder, C. M., W. Begor, and E. R. Berndt. (2011) "Economic Perspectives on the Advance Market Commitment for Pneumococcal Vaccine," *Health Affairs* 30 (8): 1508—1517.

UNICEF. (2022) "Child Health and Survival." [Internet.] Available from https://www.unicef.org/child-health-and-survival.

Walker, P. G. T., C. Whittaker, O. Watson, M. Baguelin, K. E. C. Ainslie, S. Bhatia, et al. (2020) "The Global Impact of COVID-19 and Strategies for Mitigation and Suppression." London: Imperial College COVID-19 Response Team, working paper, March 26. Available from https://www.imperial.ac.uk/media/imperialcollege/medicine/sph/ide/gidafellowships/Imperial-College-COVID19-Global-Impact-26-03-2020v2.pdf.

Weyl, E. G. and J. Tirole. (2012) "Market Power Screens Willingness-to-pay," *Quarterly Journal of Economics* 127 (4): 1971–2003.

World Health Organization. (2018) *Annual Review of Diseases Prioritized under the Research and Development Blueprint.* Geneva: WHO meeting report. Available from http://www.emro.who.int/pandemic-epidemic-diseases/news/list-of-blueprint-priority-diseases.html.

Woolsey, C. and T. W. Geisbert. (2021) "Current State of Ebola Virus Vaccines: A Snapshot," *PLoS Pathogens* 17 (12): e1010078.

Table 1: Descriptive Statistics for Cost Data from CEPI Survey

Investment cost	Notation	Obs.	Mean (mil. $)	Median (mil. $)	Std. dev. (mil. $)
Phase-3 trials ($t = 1$)	c_{1i}	99	83.1	30.0	125.3
Stockpiling ($t = 2$)	c_{2i}	27	164.8	200.0	47.7

Source: CEPI survey of suppliers' investment costs. *Notes:* Stockpile row omits 63 observations reflecting only the cost of maintaining, not manufacturing, the stockpile.

Table 2: Maximum-Likelihood Estimates of Lognormal Cost Distributions

	Investment stage, t	
	Phase-3 trials ($t = 1$)	Stockpiling ($t = 2$)
Parameter estimates		
Scale, $\hat{m}_t$	3.789	5.057
	(0.105)	(0.062)
Shape, $\hat{v}_t$	1.044	0.324
	(0.074)	(0.044)
Log-likelihood	–519.9	–144.4
Number of observations	99	27
Derived moments		
Mean, $\hat{\mu}_t$	76.2	165.5
Std. dev., $\hat{\sigma}_t$	107.1	55.0

Notes: Results are maximum-likelihood estimates. Standard errors reported below estimates in parentheses. Derived moments computed using equations (17)–(18). Estimates, standard errors, and derived moments in units of million dollars.

Table 3: Means of Selected Outcomes in Baseline Simulations of Optimal Mechanism

					Mechanism indicators			
b (mil. $)	Funder net surplus (mil. $)	Surplus capture rate	Ex post award (mil. $)	Number winners	No winner	Award set by losing bid	Award set by reserve	Stockpile completed
500	82	0.16	317	0.89	0.14	0.50	0.36	0.44
1,000	366	0.37	378	1.52	0.00	0.63	0.37	0.63
1,500	726	0.49	416	1.91	0.00	0.67	0.33	0.72
2,000	1,114	0.56	439	2.16	0.00	0.68	0.32	0.76
2,500	1,520	0.61	459	2.39	0.00	0.64	0.37	0.81
3,000	1,942	0.65	478	2.58	0.00	0.61	0.39	0.82

Notes: Entries are means of selected outcome variables over the million simulations run for the value of b heading row. Simulations run for baseline case with $n = 5$ firms and probability of success $s = 0.5$. Surplus capture rate equals funder net surplus as proportion of b. Ex post award is expressed per winning supplier. In simulations with no winner—when the lowest cost draw is not covered by the single-supplier reserve price—the expected award is set equal to that reserve price. Entries in "no winner" column reflect a positive number of cases but few enough that can round to zero as a proportion of the million simulations.

Table 4: Surplus Capture Rates Under Alternative Mechanisms

b	First best	Second best (WVAR)	Reserve without Vickrey	Vickrey without reserve with fixed number of winners, w			
				$w = 1$	$w = 2$	$w = 3$	$w = 4$
500	0.24	0.16	0.13	0.15	–0.05	–0.49	–1.22
1,000	0.44	0.37	0.27	0.32	0.35	0.19	–0.14
1,500	0.55	0.48	0.37	0.38	0.48	0.41	0.22
2,000	0.62	0.56	0.44	0.41	0.55	0.53	0.40
2,500	0.66	0.61	0.49	0.43	0.59	0.60	0.51
3,000	0.70	0.65	0.53	0.44	0.62	0.65	0.58

Notes: Entries are means of the surplus capture rate over the million simulations run for each alternative mechanism. Surplus capture rate equals funder net surplus as proportion of b. Simulations run for baseline case with $n = 5$ firms and probability of success $s = 0.5$. Simulations for different mechanisms use same set of random cost draws. Column of results for WVAR repeats results from Table 3 for reference.

Table 5: Means of Outcomes in Simulations of Optimal Mechanism in Covid-19 Pandemic

Pandemic scenario parameters	Funder net surplus (tril. $)	Ex post award (bil. $)	Number winners
$n = 10$, $s = 0.3$	2.8	37.7	9.8
$n = 10$, $s = 0.5$	3.6	31.0	9.6
$n = 5$, $s = 0.3$	1.8	41.5	4.9
$n = 5$, $s = 0.5$	2.6	39.6	4.9

Notes: Entries are means of selected outcome variables over the million simulations run for the pandemic scenario with the parameters heading row. Variable n denotes number of potential firms and s probability of firm success. First three rows reported in Exhibit 3 of Snyder et al. (2020); last row is a new set of simulation runs.

Figure 1: Working of the Funding Mechanism in an Example

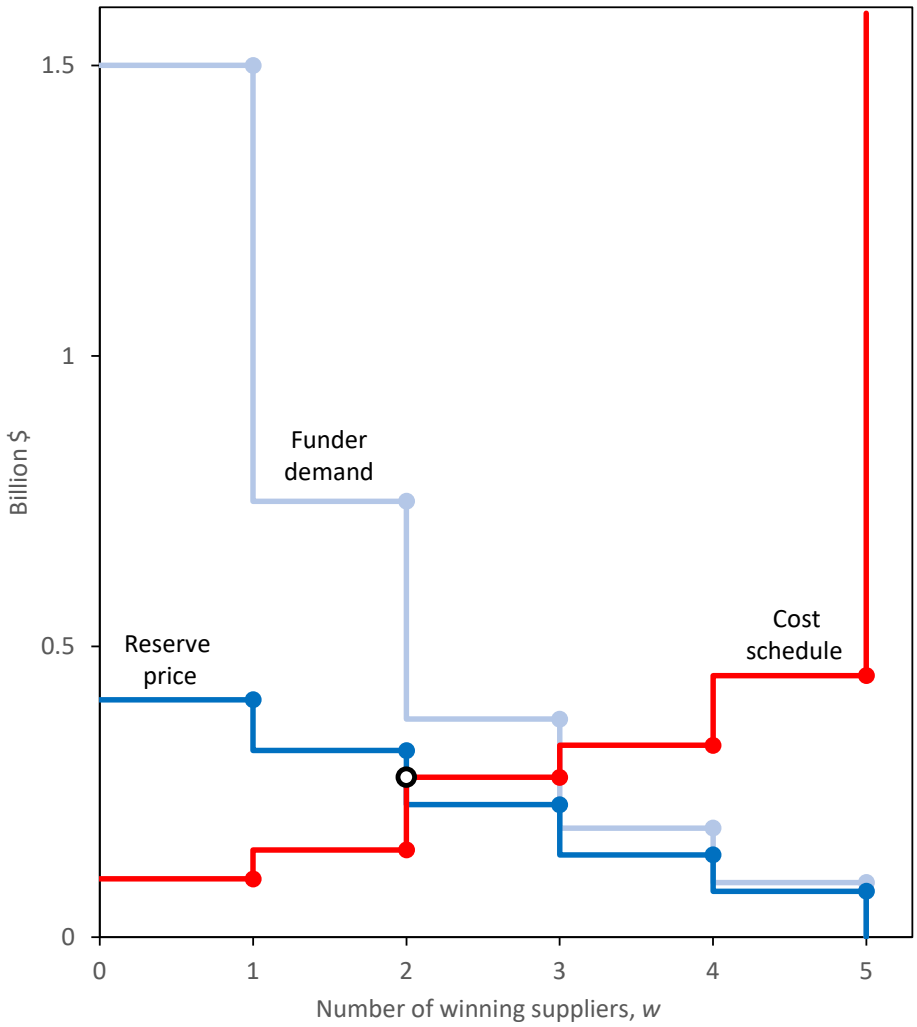

Notes: Light blue funder-demand curve equivalent to marginal social benefit $(1-s)^{w-1}sb$ from increasing the number of winning firms to w. Funder demand is drawn assuming health and economic benefit $b = \$3$ billion and success probability $s = 50\%$. Dark blue reserve-price curve derived from funder demand by applying the transformation in equation (15). Red cost curve graphs the illustrative cost draws mentioned in the text in size order. Open circle on the intersection between reserve price and costs denotes the outcome from the mechanism.

Figure 2: Lognormal Distributions Estimated for Investment Costs

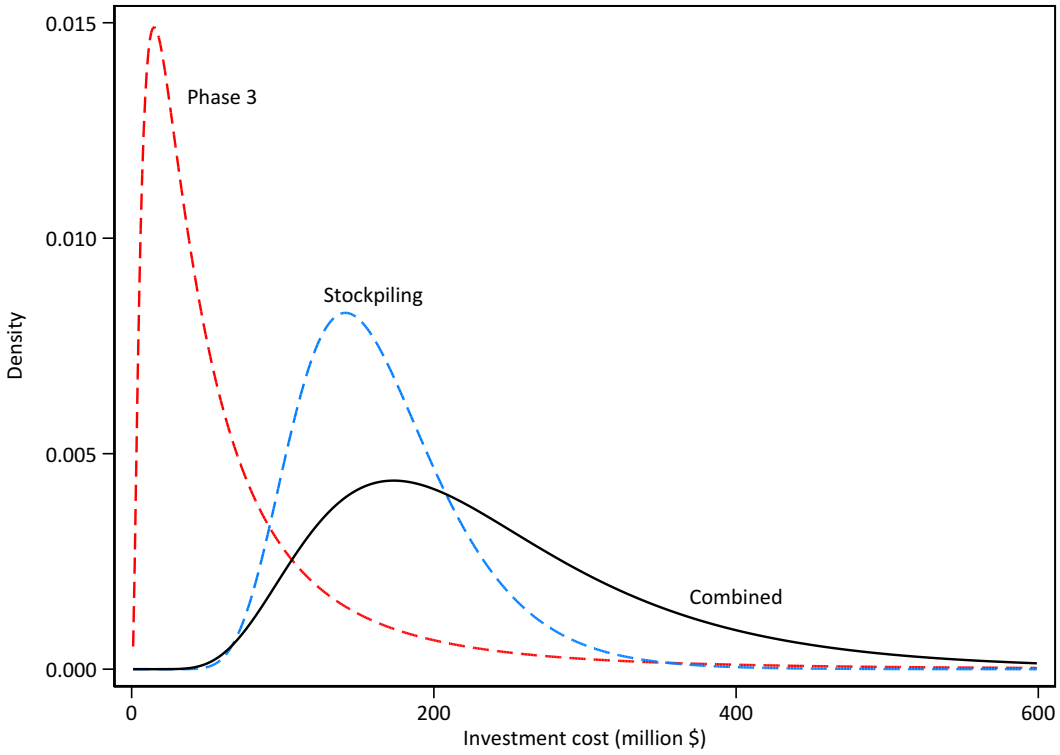

Notes: Dashed curves are lognormal distributions based on maximum-likelihood estimates provided in Table 2. Lognormal distribution of investment cost for combined stages obtained via Fenton approximation.

Figure 3: Histograms of Funder Net Surplus Across Simulations

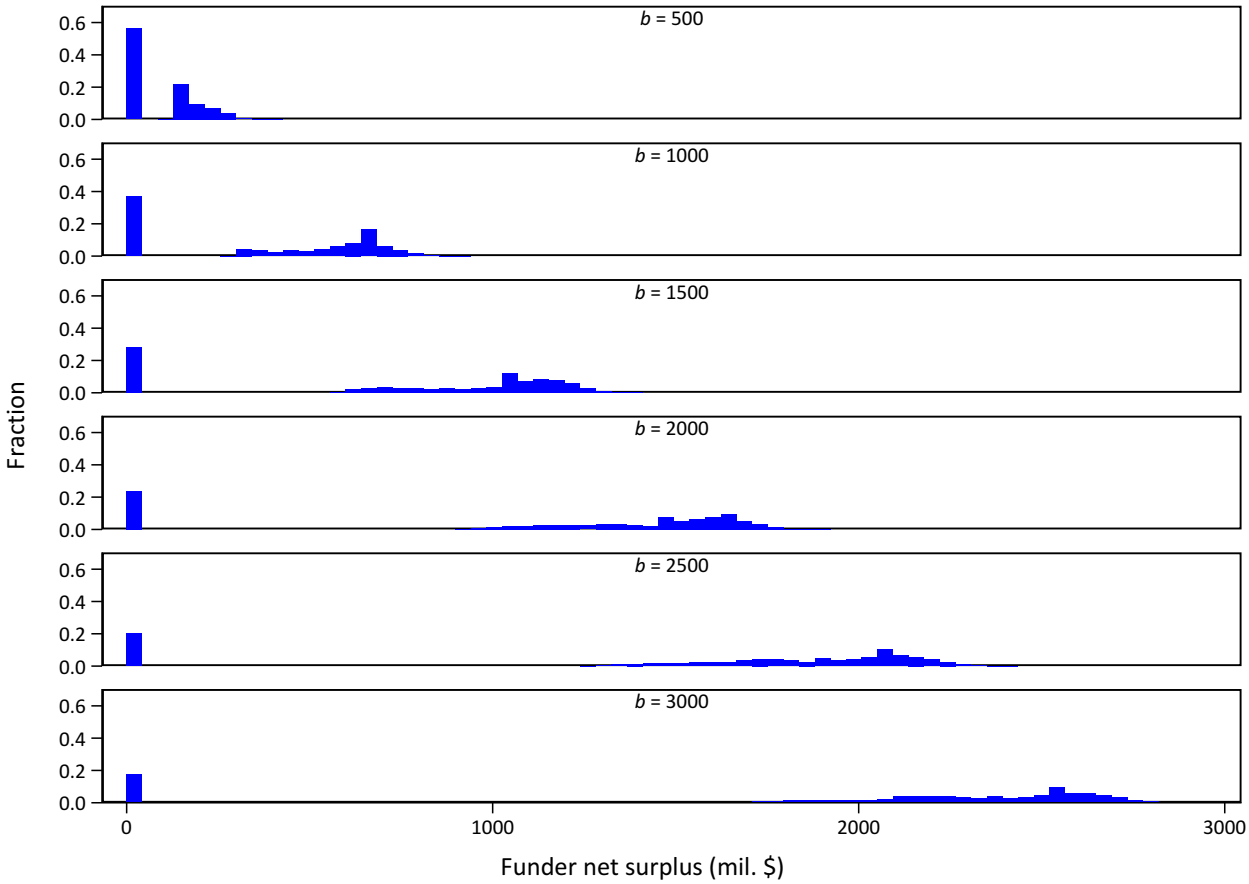

Notes: Histograms for one million simulations for each value of benefit b from 500 to 3000 (all results in million $) shown in different panels. Simulations run for baseline case with $n = 5$ firms and probability of success $s = 0.5$.

Figure 4: Means of Simulation Outcomes Varying Number of Suppliers

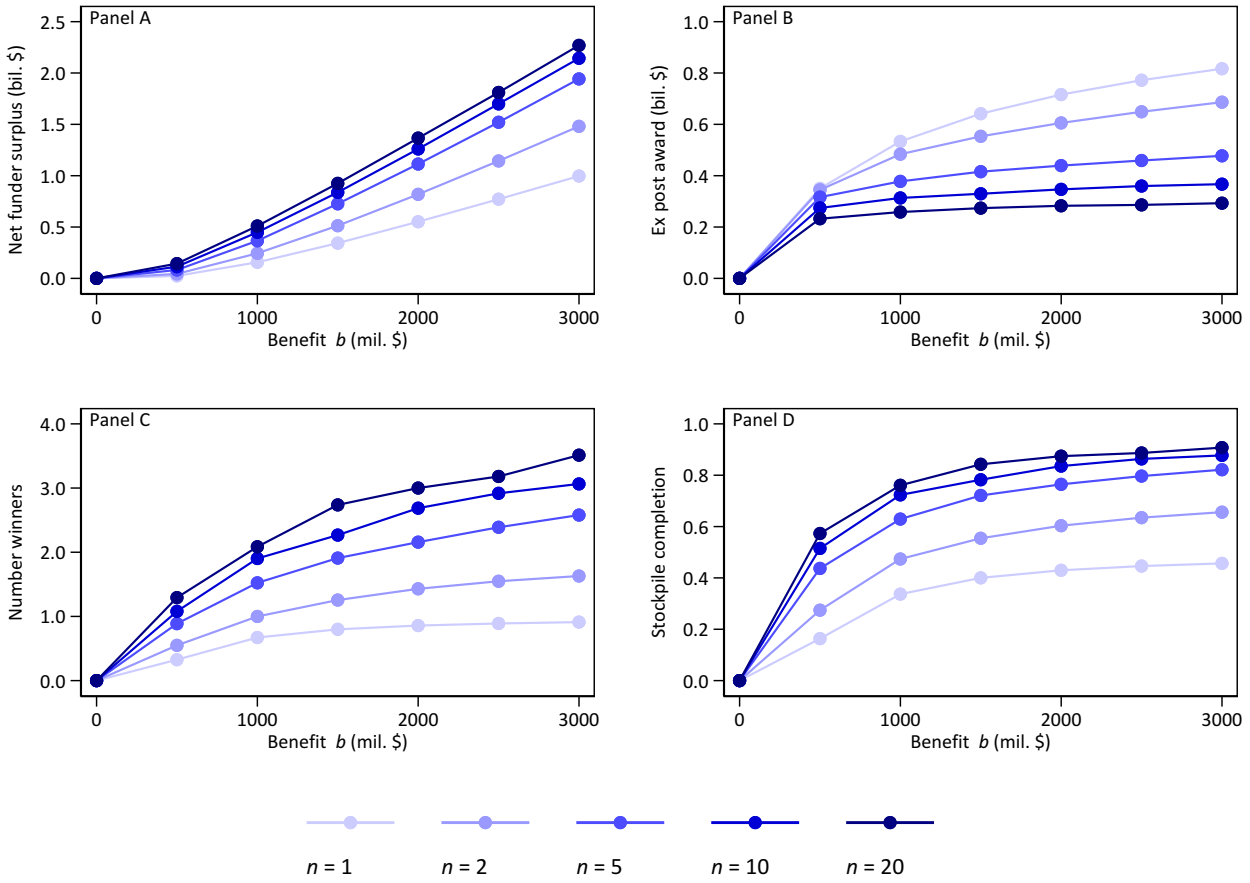

Notes: Each dot represents mean of the results across one million simulations run for the indicated parameters. Darker lines represent higher number of suppliers, n. Probability of success set to baseline value, $s = 0.5$. In Panel B, the mean ex ante award is the average across simulations of the per-supplier payment conditional on success specified by the mechanism, regardless of whether and how many suppliers received it.

Figure 5: Means of Simulation Outcomes Varying Probability of Success

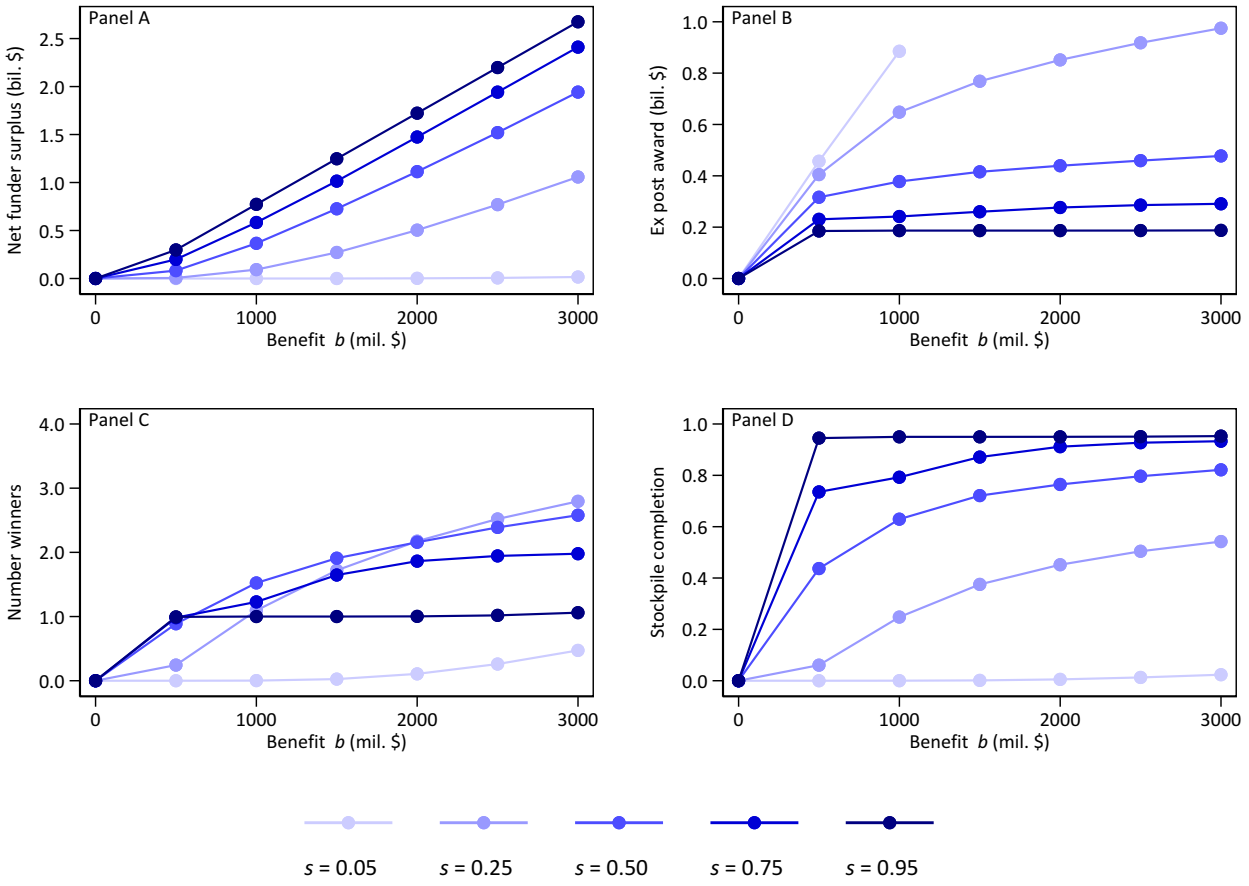

Notes: Each dot represents mean of the results across one million simulations run for the indicated parameters. Darker lines represent higher success probability, s. Number of suppliers set to baseline value, $n = 5$. Curve for $s = 0.5$ is identical to the curve of the same shade for $n = 5$ in Figure 4. In Panel B, the mean ex ante award is the average across simulations of the per-supplier payment conditional on success specified by the mechanism, regardless of whether and how many suppliers received it. In that panel, curve for $s = 0.05$ truncated above an ex post award of $1 billion to preserve legibility of other curves and for consistency with axes in Figure 4.