

UCLA

UCLA Previously Published Works

Title

Managing two-dose COVID-19 vaccine rollouts with limited supply: Operations strategies for distributing time-sensitive resources

Permalink

<https://escholarship.org/uc/item/8bb8m97h>

Journal

Production and Operations Management, 31(12)

ISSN

1059-1478

Authors

Mak, Ho-Yin
Dai, Tinglong
Tang, Christopher S

Publication Date

2022-12-01

DOI

10.1111/poms.13862

Peer reviewed

Managing two-dose COVID-19 vaccine rollouts with limited supply: Operations strategies for distributing time-sensitive resources

Ho-Yin Mak¹  | Tinglong Dai^{2,3}  | Christopher S. Tang⁴ 

¹McDonough School of Business, Georgetown University, 37th and O Streets, NW, Washington, District of Columbia, USA

²Carey Business School, Johns Hopkins University, 100 International Drive, Baltimore, Maryland, USA

³Hopkins Business of Health Initiative, Johns Hopkins University, 100 International Drive, Baltimore, USA

⁴UCLA Anderson School of Management, University of California, Los Angeles, 110 Westwood Plaza, Los Angeles, California, USA

Correspondence

Tinglong Dai, Carey Business School, Johns Hopkins University, Baltimore, MD 21202, USA.
Email: dai@jhu.edu

Handling Editor: Kalyan Singhal

Abstract

Distributing scarce resources such as COVID-19 vaccines is often a highly time-sensitive and mission-critical operation. Our research was prompted by a significant obstacle that the United States and other nations encountered during the early months of the COVID-19 vaccination campaign: Most COVID-19 vaccines require two doses given 3 or 4 weeks apart. Given the severely limited supply and mounting pressure on many countries to reduce hospitalizations and mortality, how to effectively roll out two-dose vaccines was a critical policy decision. In this paper, we first model and analyze inventory dynamics of the rollout process under three rollout strategies: (1) holding back second doses, (2) releasing second doses, and (3) stretching the lead time between doses. Then we develop an SEIR (susceptible, exposed, infectious, recovered) model that incorporates COVID-19 asymptomatic and symptomatic infections to evaluate these strategies in terms of infections, hospitalizations, and mortality. Among our findings, we show releasing second doses reduces infections but creates uneven vaccination patterns. In addition, to ensure second doses are given on time without holding back inventory, strictly less than half of the supply can be allocated to first-dose appointments. Stretching the between-dose lead time flattens the infection curve and reduces both hospitalizations and mortality compared with the strategy of releasing second doses. We also consider an alternative single-dose vaccine with lower efficacy and show that the vaccine can be more effective than its two-dose counterparts in reducing infections and mortality. We conduct extensive sensitivity analyses related to age composition, risk-based prioritization, supply disruptions, and disease transmissibility. Our paper provides important implications for policymakers to develop effective vaccine rollout strategies in developed and developing countries alike. More broadly, our paper sheds light on how to develop effective operations strategies for distributing time-sensitive resources in times of crisis.

KEYWORDS

COVID-19 vaccine rollout, healthcare operations management, public health operations, SEIR model, vaccine inventory

1 | INTRODUCTION

This research was motivated by a major challenge that COVID-19 vaccination campaigns faced during the first half of 2021. One year after the onset of the COVID-19 pandemic, global efforts aimed at ending the pandemic had evolved from

developing vaccines at a record pace to rapidly putting doses into people's arms. Experts predicted that manufacturing or deployment delays will cause the benefit of vaccination to "decline substantially" (Paltiel et al. 2021, p. 42). Adding to the urgency of distributing vaccines, SARS-CoV-2, the virus underlying the disease, continued to rapidly spread and mutate. Racing against time, the mass vaccination effort was arguably "the biggest logistics challenge in history" (Arnold, 2020, p. 36), requiring vaccine manufacturers to produce

billions of doses of newly developed vaccines in a matter of months.

Yet, expanding vaccine supply was challenging for two reasons. First, the vaccine production rate is limited by physical infrastructure, requisite raw materials (e.g., enzymes for converting DNA to mRNA for both the Pfizer and Moderna vaccines),¹ and packaging materials (e.g., glass vials and rubber stoppers). Second, the production yield can be uncertain. For example, uncertainty in AstraZeneca's vaccine production process, which requires the growth of cells inside bioreactors, caused major delays in its delivery to the E.U. in January 2021. In addition, developed countries have locked up much of the vaccine supply through 2023, leaving developing countries struggling with stringent supply constraints (Hopkins, 2021). Specifically, the vaccine shortage was particularly salient in Africa, which accounts for 18% of the world's population but has only received 2% of global COVID-19 vaccine supply by June 2021 (Muhumuza & Mutsaka, 2021).

In addition to limited supply, the authorization of three major COVID-19 vaccines (AstraZeneca, Moderna, and Pfizer) was for two-dose regimens, creating unprecedented *timing and inventory planning* issues. For example, the Pfizer vaccine required two doses with a recommended 3-week interval between them.² To ensure adherence to the two-dose regimen, the United States initially (from December 2020) followed a "hold-back policy" for each dose offered to a first-time recipient; that is, one extra dose is held in reserve until the recipient returns for the second dose (Gottlieb, 2021).

The hold-back policy ensures that two doses are reserved for each recipient. However, when facing severely limited vaccine supplies and rapidly rising COVID-19 case counts, the hold-back policy can cause a delay for other individuals to receive their first doses. To accelerate the vaccination process in order to reach herd immunity sooner, the United States and the United Kingdom had proposed unconventional solutions intended to address the logistical requirement of offering two doses of the COVID-19 vaccine for each recipient.

In the United States, then-President-elect Joe Biden proposed on January 8, 2021, that he would order a switch from the hold-back policy to a stock "release policy," which means that (1) no vaccine inventory would be held for *future* vaccine recipients who return to receive their second doses and (2) all vaccine inventory is used as either first doses for new recipients or second doses for returning recipients (O'Donnell & Spalding, 2021). Unlike the hold-back policy that can guarantee that the two-dose regimen is administered according to the recommended time interval, the release policy can increase access of the first doses to more recipients. However, without dedicated inventory for the second doses, shortages may arise when vaccine supply remains limited.

Across the Atlantic, the U.K. government announced a related and somewhat controversial policy on December 30, 2020. Instead of administering two vaccine doses 3 weeks

apart as recommended (in the case of the Pfizer vaccine), the lead time would be stretched to 12 weeks. Logistically, this "stretching policy" defers the required demand from second-dose recipients and frees up more stock to cover first-dose recipients. Hence, the stretching policy generates an effect similar to that of the release policy: It can expand the access to the first dose for more recipients, but it will delay the administration of the second dose, especially when the vaccine supply is limited. Delaying the time for recipients to receive their required second doses can increase the chance of infection for these recipients during the delay.

Both the U.S. and U.K. policies aimed to maximize the number of people developing at least partial immunity. As *The Guardian* reported, prioritizing the first dose "would mean each vulnerable person who has received the vaccine would be afforded less protection, [while] the number of people given at least some protection would double" (Rawlinson, 2020). Because clinical data suggest the marginal gain in protection of the first dose is significantly larger than that of the second dose, the U.K.'s approach appeared to increase overall protection to the population.

On the other hand, both policies were controversial for myriad reasons (Mancini & Gross, 2021). In particular, the issue of whether to hold back second doses garnered significant public attention in the United States: Some experts are in favor of holding back second doses, citing the risk of insufficient protection of single-dose inoculations in times of potential disruptions (Mak & Tang, 2021); others believe the supply rate of COVID-19 vaccines is expected to grow over time, at least in the United States, so holding back second doses is unnecessary (Dai & Yadav, 2021). These diverging views motivated us to model and compare the release and stretching policies against the benchmark hold-back policy. Essentially, both the release and stretching policies were intended to address the challenge of limited vaccine supply through an operational approach: Both policies call for *shifts in inventory-control policies* for vaccine stocks. Therefore, debates over these policies must involve the operations management perspective, which has been largely absent in the public discourse, in addition to the clinical, public health, and political perspectives. To facilitate such discussions, we developed a modeling framework to examine and compare these policies.

Due to limited vaccine supply (especially in many developing countries), allocating vaccines has become a time-sensitive and mission-critical issue. In our context, these rollout policies of a vaccine with a two-dose requirement involve the following trade-offs: (i) The hold-back policy can secure two doses for fewer recipients with high efficacy, yet it delays other recipients' access to the first doses (which can give them some immunity); (ii) the release policy can increase access to the first doses, but it may cause shortages or delays for administering the required second doses for returning recipients; and (iii) the stretching policy can also vaccinate more recipients, but it can cause even longer delays for the recipients receiving their required second doses.

Critically, through modeling the rollout process under the alternative policies, we point out that the popular belief that “the number of people vaccinated would double” (Rawlinson, 2020) is *incorrect*. In particular, the clinical requirement to administer the second dose imposes an *intertemporal* constraint on the rollout problem: As more first doses are given early on, supplies in later periods become more tightly constrained.

In light of these trade-offs, the limited supply of vaccines, and how the required second dose can boost immunity against COVID-19, the implications of these policies are not immediately clear, especially when the timing for individuals to receive the first and second doses would vary for each policy. These observations motivated us to develop a dynamic model to examine the following questions arising from the rollout process of a COVID-19 vaccine with a two-dose requirement:

- (i) Given the production rate of vaccines (constant or linearly increasing), what are the cumulative inoculation rates for the first and second doses under the hold-back policy, the release policy, and the stretching policy?
- (ii) Relative to the hold-back policy, will the release and stretching policies generate higher cumulative inoculation rates?
- (iii) In light of an age-stratified SEIR (susceptible, exposed, infectious, recovered) model that predicts the infection process under different vaccination policies, which policy is the most effective in reducing infections, hospitalizations, and mortality?
- (iv) How would a single-dose vaccine with a lower efficacy (e.g., the Johnson & Johnson COVID-19 vaccine) perform relative to its two-dose counterparts?

As an initial attempt to examine the above questions arising from the two-dose regimen of all approved vaccines and the limited supply of vaccines, we developed a dynamic model by examining the first- and second-dose inoculation rates of a single vaccine under all three policies (hold-back, release, and stretching). We also extend our model to capture the case of rolling out a new single-dose vaccine with a lower efficacy.

Through analytical comparisons, we find that under a constant or linearly increasing vaccine production rate, both release and stretching policies generate a higher cumulative vaccination population than the hold-back policy. Somewhat counterintuitively, we show that even if one intends to release all second doses, no more than half of the available doses can be allocated to first-dose appointments, in order to avoid any delays in administering those required second doses for returning recipients, no matter how quickly supply grows over time.

One can view the single-dose regimen as an extreme case of a stretching policy in which the timing of the required second dose is extended to the infinite future. Using this logic, we show that the single-dose regimen can generate a higher cumulative vaccinated population than the hold-back policy and the release policy.

We also developed a compartmental SEIR model that incorporates COVID-19 asymptomatic and symptomatic infections in a population that consists of high- and low-risk groups. Our SEIR model captures the evolution of the COVID-19 pandemic among these two groups who receive zero, one, or two doses of vaccine under different rollout strategies over time. We also calibrate our parameter values using the latest vaccine effectiveness studies (e.g., Haas et al., 2021). Our simulation study of our SEIR model reveals that both release and stretching policies can generate lower disease transmissions than the hold-back policy; however, these two policies create uneven vaccination patterns. Stretching the between-dose lead time flattens the infection curve, and reduces hospitalization and mortality (more than under the strategy of releasing second doses). We also consider an alternative single-dose vaccine (e.g., J&J vaccine) with a lower overall efficacy and show that it can be more effective than its two-dose counterparts in reducing infections, hospitalizations, and deaths. We conduct extensive sensitivity analyses related to age composition, risk-based prioritization, supply disruptions, and the basic reproduction number R_0 .

Although our work was motivated by the initial months of COVID-19 vaccination efforts in the E.U., United Kingdom, and United States, it has broader policy implications. Developing countries are set to face limited vaccine supply as developed countries have locked up most of the vaccine supply up to 2023 (Hopkins, 2021). This is the latest episode of the longstanding ineffectiveness of global health supply chains that calls for better operations management practices (Gallien et al., 2017; Kraiselburd & Yadav, 2013). Our results can help policymakers in various developing countries understand the pros and cons of different rollout policies. Furthermore, our modeling framework and our SEIR model can be applied to a variety of vaccination environments (when relevant data become available), including different variants and different vaccines with varying administration regimens and levels of efficacy.

The remainder of the paper is organized as follows. We review the relevant literature in Section 2. Section 3 describes our model and the three rollout policies. Using the hold-back policy as a benchmark, we examine the implications of the release policy in Section 4 and the stretching policy in Section 5. In Section 6, we examine an alternative single-dose vaccine. In Section 7, we develop an SEIR model to evaluate the three policies and the single-dose regimen. We conclude our paper in Section 8. All proofs are provided in Section OAI of the Supporting Information.

2 | LITERATURE REVIEW

Vaccine distribution is a vibrant line of research that has attracted substantial interest from the healthcare operations management community. Most of this stream of literature focuses on the influenza vaccine (for exceptions, see, e.g., Cummings et al., 2021; Hall et al., 2008). Whereas several papers in this literature (e.g., Arifoğlu et al., 2012; Arifoğlu

& Tang, 2022; Chick et al., 2008, 2017; Yamin & Gavius, 2013) study governments' vaccine contracting, procurement, and subsidies decisions, others (e.g., Cho & Tang, 2013; Dai et al., 2016) examine optimal contract design between vaccine manufacturers and healthcare providers who administer vaccines. Duijzer et al. (2018) use a susceptible–infectious–recovered (SIR) model to study the critical vaccine coverage leading to herd effect for a pandemic and consider the optimal allocation of vaccine doses across different populations. Unlike this literature, which focuses primarily on distributing *single*-dose vaccines, we study a unique problem of distributing a two-dose vaccine during a pandemic in which whether to hold back second doses becomes a consequential national policy. In doing so, our paper speaks to a number of related public health papers on various rollout strategies, including, for example, Paltiel et al. (2021a), Paltiel et al. (2021b), Tuite et al. (2021), and Voysey et al. (2021).

Another related stream of literature studies inventory management—a classical theme of the operations research and management science community. This literature emerges from the seminal work by Arrow et al. (1951), Dvoretzky et al. (1953), and Scarf (1960), and has become the foundation of much of today's work in manufacturing and supply chain operations (Porteus, 1990; Song et al., 2020). Despite the immense breadth and depth of the extant inventory literature, our paper represents a first attempt to address a novel time-sensitive resource allocation problem in which each customer requires exactly *two units* of the product at a *specified timing*. In addition, unlike the literature that focuses on a single organization's decision, we examine a national policy with public health implications. The only other work in the inventory literature that shares similar characteristics, to the best of our knowledge, is the paper by Natarajan and Swaminathan (2017), which considers the dynamic allocation of inventory of treatment drugs among patients with different (and evolving) health states. Compared with this work, which considers patients' health conditions in each state to evolve exogenously over time, our vaccine rollout problem involves the evolution of a pandemic that is governed by an epidemiology model.

Our paper also contributes to the rapidly growing operations management literature that addresses the COVID-19 pandemic. For example, Kaplan (2020) documents a number of basic mathematical models to support local-level COVID-19 decisions ranging from event crowd control to hospital capacity planning. As another example, Chen et al. (2020) use mobility data around nursing homes to establish the relationship between cross-facility staff movement and COVID-19 infections. Shumsky et al. (2021) study a healthcare provider's COVID-19 vaccine inventory management problem. We enrich this literature by studying a novel problem relevant to how health policymakers roll out COVID-19 vaccines—the endgame of the COVID-19 pandemic—and our findings have important implications for vaccination policies in future pandemics.

3 | MODELING A TWO-DOSE VACCINE ROLLOUT PROCESS

In this section, we present a parsimonious model of the vaccine rollout process in the spirit of scratch models as articulated by Kaplan (2020). To examine the implications of the three aforementioned policies—including (1) holding back second doses, (2) releasing second doses, and (3) stretching the between-dose lead time—our model focuses on the dynamic first- and second-dose inoculation rates over time by taking into consideration the limited availability of vaccine supply. Specifically, our model involves the following elements.

Vaccine supply rate $y(t)$

Consider a time horizon with a length of T (e.g., a year) during which the vaccine is rolled out. For any $t \in [0, T]$, let $y(t)$ be the deterministic supply rate of vaccine doses available at t ³; $y(t)$ is exogenously determined by the rate of manufacturing, quality and safety checks, and transportation. We assume $y(t)$ is (weakly) increasing in t ; that is, the supply ramps up over time. Let $Y(t)$ be the cumulative supply up to time t , where $Y(t) = \int_0^t y(\tau) d\tau$.

Vaccination rate $v_i(t)$ and *cumulative vaccination* $V_i(t)$

We denote by $v_1(t)$ and $v_2(t)$ the instantaneous rates of the first- and second-dose inoculations at time t , respectively. Thus, the cumulative vaccinations completed by time t , denoted by $V_i(t)$, satisfies

$$V_i(t) = \int_0^t v_i(\tau) d\tau \text{ for } i \in \{1, 2\}. \quad (1)$$

The cumulative vaccinations $V_1(t)$ (and $V_2(t)$) can be interpreted as the partially (and fully) protected populations who have received the first dose (and the second dose) by time $t \in [0, T]$.

Lead time L

A parameter of key concern is the lead time L ; that is, the recommended time interval between the two doses of the same vaccine. The lead time L depends on the specific vaccine and the vaccination policy. In the United States, Pfizer recommends $L = 3$ weeks and Moderna recommends $L = 4$. For simplicity of our analysis, we assume the second dose must be received exactly L periods after the first dose.⁴ However, in the United Kingdom, $L = 12$ weeks for the AstraZeneca, Moderna, and Pfizer vaccines under the dose-stretching policy.

Vaccination rollout policy $(v_1(\cdot), v_2(\cdot), L)$

As we will explain later, the hold-back, release, and stretching policies correspond to a particular choice of functions $v_1(\cdot)$, $v_2(\cdot)$, and the lead time L . In general, a feasible policy $(v_1(\cdot), v_2(\cdot), L)$ must satisfy the following constraints at any time t ⁵:

$$V_1(t) + V_2(t) \leq Y(t) \quad \text{for } t \in [0, T] \quad (2)$$

$$V_2(t) = 0 \quad \text{for } t \in [0, L] \quad (3)$$

$$V_2(t) \leq V_1(t - L) \quad \text{for } t \in [L, T]. \quad (4)$$

Due to limited vaccine supply, constraint (2) suggests cumulative first- and second-dose vaccinations offered are bounded by the cumulative supply at every instant t . In addition, to comply with the two-dose regimen requirement, constraints (3) and (4) indicate second doses cannot be offered until L time units have elapsed after first-dose inoculations. (If there is sufficient supply of vaccine at all times, then (4) will hold as equality; i.e., $V_2(t) = V_1(t - L)$.)

Cumulative protection $P_i(t)$

Building on the (partially and fully) protected population sizes ($V_1(\cdot)$ and $V_2(\cdot)$), we now define the cumulative partial protection rate (due to the first dose) $P_1(t)$ and full protection rate (due to both doses) $P_2(t)$ up to time t (in person-weeks):

$$P_i(t) = \int_0^t V_i(\tau) d\tau \quad \text{for } i = 1, 2. \quad (5)$$

The cumulative protection levels keep track of the number of people protected by vaccination *and* the duration over which protection has been developed. Hence, this measure reflects the amount of protection offered to the population as a whole. To elaborate, consider the following example.

Example 1. Consider two hypothetical scenarios: (i) One million people are given the first dose instantaneously at time $t = 0$ and none thereafter; and (ii) 50,000 people are given the first dose per day at a continuous rate over 20 days and none thereafter. In both cases, the cumulative first-dose vaccinated population sizes by Day 20 are the same at $V_1(t) = 1$ million when $t = 20$. However, the cumulative protection levels are different: $P_1(t)$ is equal to 20 million person-days in scenario (i), and is equal to 10 million person-days in scenario (ii). Thus, the amount of protection in scenario (i) is double that under scenario (ii).

The functions $y(t)$, $Y(t)$, $v_i(t)$, and $V_i(t)$ are meaningfully defined only for $t \in [0, T]$; for notational brevity, we adopt the convention that their values equal zero when $t < 0$. In addition to these measures, we further consider a broader class of performance metrics, including symptomatic infections, hospitalizations, and mortality, based on an SEIR model in Section 7 and Section OA2 of the Supporting Information. This involves a system of differential equations that utilizes $v_i(t)$ and $V_i(t)$ for $i = 1, 2$.

3.1 | Stock hold-back policy ($v_1^h(t)$, $v_2^h(t)$)

As explained in Section 1, the U.S. government followed the manufacturer's recommended lead time of L weeks by adopting the hold-back policy as follows. Specifically, each facility only administers half of the vaccine stock received for first doses, and holds back the other half for L weeks to ensure recipients will have guaranteed stock when they return for the second dose. By considering the supply rate $y(t)$, the hold-back policy ($v_1^h(t)$, $v_2^h(t)$) stipulates the following⁶:

$$v_1^h(t) = \frac{1}{2}y(t) \quad \text{and} \quad v_2^h(t) = v_1^h(t - L). \quad (6)$$

This policy ensures constraints (2)–(4) are satisfied at all time.

We now represent the cumulative protection functions $P_i^h(t)$ using the external supply function $Y(t)$. First, by the definition of $V_i^h(t)$, we have

$$V_1^h(t) = \frac{1}{2}Y(t) \quad \text{and} \quad V_2^h(t) = \frac{1}{2}Y(t - L). \quad (7)$$

Then, the amount of protection by time t can be given by

$$P_1^h(t) = \frac{1}{2} \int_0^t Y(\tau) d\tau \quad \text{and} \quad P_2^h(t) = \frac{1}{2} \int_0^{t-L} Y(\tau) d\tau. \quad (8)$$

3.2 | Stock-release policy ($v_1^r(t)$, $v_2^r(t)$)

Under the stock-release policy, vaccine stock is no longer put on reserve for second doses.⁷ The intention of the release policy is to maximize the rate of first-dose inoculations to accelerate the build-up of partial protection. Using our notation, the release policy stipulates that $v_1(t) + v_2(t) = y(t)$, such that all available stock is offered (as first or second doses) without delay. However, if the supply rate does not increase fast enough over time, then increasing the number of first-dose recipients earlier would constrain the number of first-dose recipients later.

To illustrate different rollout dynamics under the hold-back and the release policies with limited vaccine supply, let us consider a simple example.

Example 2. Consider a constant supply rate of 100 doses per day (i.e., $y(t) = 100$) and a lead time of $L = 21$ days (i.e., 3 weeks between dose time interval for the Pfizer vaccine).

- (i) *Hold-back policy.* Under the hold-back policy, 50 first doses will be offered per day, whereas the remaining 50 doses of supply will be reserved for the same 50 recipients who return 21 days later for their second doses. This pattern repeats every day.
- (ii) *Release policy.* Under the release policy, from Day 1 to Day 21, 100 doses per day (i.e., the entire supply for each day) are administered to 100 people as their first

doses. However, on Day 22, those 100 recipients who received the first doses on Day 1 are now due for their second doses. To adhere to the required dosage schedule, it is necessary to allocate all 100 doses (i.e., the entire supply) on Day 22 for those 100 returning second-dose recipients. In doing so, the provider has no vaccine to offer to any first-dose recipient on Day 22. This situation will continue until Day 43, when there are no more recipients who are due for the second dose (because no first-dose recipient received a first dose from Day 22 to Day 42). Therefore, on Day 43, all 100 doses will be released to 100 first-dose recipients, and the pattern repeats as on Day 1. Overall, this “alternating pattern” repeats in 42-day (i.e., twice of the lead time $L = 21$ days) intervals.⁸

We now define the stock-release policy $(v_1^r(t), v_2^r(t))$ more formally. Initially, all available supply is released as first-dose vaccination over $t \in [0, L)$ so that $(v_1^r(t), v_2^r(t))$ satisfy

$$v_1^r(t) = y(t) \text{ and } v_2^r(t) = 0. \quad (9)$$

When $t \geq L$, the available supply will be split between first- and second-dose appointments. Let $B^r(t) = V_1^r(t - L) - V_2^r(t)$ denote the amount of backlog by time t ; that is, the difference between the number of people who have received their first doses L weeks ago and the number of people who have received their second doses. For clinical reasons, the backlog $B^r(t)$ should be minimized. Therefore, the rate of second doses offered should match as closely as possible the rate of first doses completed L periods ago. To capture this preference, the release policy $(v_1^r(t), v_2^r(t))$ for $t \in [L, T]$ must satisfy:

$$\begin{aligned} v_1^r(t) &= y(t) - v_2^r(t) \\ v_2^r(t) &= \begin{cases} \min\{v_1^r(t - L), y(t)\} & \text{if } B^r(t) = 0 \\ y(t) & \text{otherwise.} \end{cases} \end{aligned} \quad (10)$$

That is, when the backlog of second-dose appointments exists (i.e., $B^r(t) > 0$, or, equivalently, $V_2^r(t) < V_1^r(t - L)$), the entire supply will be allocated for the second dose. In the absence of a backlog, the second-dose inoculation rate follows the first-dose rate lagged by lead time L , up to the available capacity.

The backlog of second-dose appointments $B^r(t)$ will build up if the rate of recipients due for second doses, $v_1(t - L)$, exceeds the available supply, $y(t)$. When the supply rate ramps up slowly, most (or all) of the available supply might be allocated for second doses in order to avoid backlog. This scenario is not desirable for the purpose of steadily increasing the protected population.

The next result evaluates the fraction of supply that can be allocated for first-dose appointments. In preparation, let us define β as the minimum fraction of available supply that is allocated for first-dose appointments at all times, that is, $\beta = \min_{t \geq L} v_1(t)/y(t)$.

Lemma 1. *No backlogs exist if and only if β satisfies*

$$(1 - \beta) \cdot y(L + t) \geq \begin{cases} y(t) & \text{for all } t \in [0, L) \\ y(t) - y(t - L) & \text{for all } t \in [L, 2L) \\ y(t) - \beta \cdot y(t - L) & \text{for all } t \in [2L, T]. \end{cases} \quad (11)$$

Lemma 1 suggests that, to ensure a minimum fraction (β) of supply to be allocated for first doses (in order to continuously build up the protected population) at all times, the supply rate $y(t)$ must grow sufficiently quickly. The first two conditions of Lemma 1 hold when $(1 - \beta)$, the proportion of supplies made available for second doses, is high and/or supply accelerates quickly in the initial weeks (especially during the first L periods). To illustrate the third condition, consider the case in which $\beta = \frac{1}{2}$; that is, when at least half of the available supply is to be used on first doses. In that case, the third condition holds if $y(t)$ is convex.

If vaccine supply does not accelerate sufficiently quickly, the rate of first-dose inoculations cannot be sustained. To avoid backlogs for second-dose inoculations, first-dose offerings must be temporarily reduced. As we show in the next section, this reduction gives rise to an alternating pattern between offering first and second doses (see Example 2).

3.3 | Dose-stretching policy $(v_1^s(t), v_2^s(t))$

The dose-stretching policy is intended to optimize the use of vaccine supply and maximize first-dose inoculations. We define the dose-stretching policy $(v_1^s(t), v_2^s(t))$ by focusing on the effect of stretching the lead time from the recommended L to $L^s = mL$, where $m \geq 1$ is an integer.⁹ Using this definition, we may view the stretching policy as a generalization of the stock-release policy with a longer lead time $L^s = mL$, where $m > 1$. Thus, we omit the analysis of the policy for brevity.

4 | IS RELEASING SECOND DOSES THE RIGHT MOVE?

We now compare the performance between the hold-back policy $(v_1^h(t), v_2^h(t))$ and release policy $(v_1^r(t), v_2^r(t))$. To do so, we first derive the cumulative vaccinated population $V_i(t)$ and the cumulative protection $P_i(t)$ associated with these policies, by considering the case in which the supply rate grows linearly over time.¹⁰ In particular, we assume, without loss of generality, that $y(t) = 1 + \alpha t$, where $\alpha \geq 0$ and the initial supply rate is normalized to one (i.e., $y(0) = 1$).

4.1 | Benchmark: Hold-back policy

By using the supply rate $y(t) = 1 + \alpha t$ and the hold-back policy $(v_1^h(t), v_2^h(t))$ as defined in Section 3.1, we derive the

corresponding performance measures $V_i^h(t)$ and $P_i^h(t)$ for $i = 1, 2$ below.

Lemma 2. *Under the hold-back policy ($v_1^h(t), v_2^h(t)$) and a linear supply growth, the cumulative vaccinations and protection evolve over time as follows:*

$$\begin{aligned}
 V_1^h(t) &= \frac{t}{2} + \frac{\alpha}{4}t^2 \text{ for } t \in [0, T], \\
 V_2^h(t) &= \begin{cases} 0 & \text{for } t \in [0, L) \\ \frac{1}{2}(t-L) + \frac{\alpha}{4}(t-L)^2 & \text{for } t \in [L, T], \end{cases} \\
 P_1^h(t) &= \frac{1}{4}t^2 + \frac{\alpha}{12}t^3 \text{ for } t \in [0, T], \\
 P_2^h(t) &= \begin{cases} 0 & \text{for } t \in [0, L) \\ \frac{1}{4}(t-L)^2 + \frac{\alpha}{12}(t-L)^3 & \text{for } t \in [L, T]. \end{cases}
 \end{aligned} \tag{12}$$

The measures characterized in Lemma 2 later serve as the benchmark for ensuing comparisons.

$$\begin{aligned}
 \Delta V_1^r(t) \triangleq V_1^r(t) - V_1^h(t) &= \begin{cases} \frac{\alpha}{2}nL^2 + \frac{1}{2}(t-2nL) + \frac{\alpha}{4}(t-2nL)^2 & \text{for } t \in [2nL, (2n+1)L) \\ \frac{\alpha}{2}nL^2 + \frac{1}{2}L + \frac{\alpha}{2}L^2 - \frac{1}{2}[t-(2n+1)L] & \\ -\frac{\alpha}{4}[2(n+1)L-t]^2 & \text{for } t \in [(2n+1)L, (2n+2)L) \end{cases} \\
 \Delta V_2^r(t) \triangleq V_2^r(t) - V_2^h(t) &= \Delta V_1^r(t-L).
 \end{aligned} \tag{13}$$

4.2 | Implications of the release policy

Using the same approach as in Lemma 2, we can derive the expressions of the performance measures $V_i^r(t)$ and $P_i^r(t)$ associated with the stock-release policy ($v_1^r(t), v_2^r(t)$) as defined in Section 3.2. We omit the details for conciseness.

By considering the cumulative vaccinated population $V_i^r(t)$ associated with the release policy, we can determine the backlog $B^r(t)$ as defined in Section 3.2, where $B^r(t) = V_1^r(t-L) - V_2^r(t)$. The next result follows from Lemma 1 and highlights the conditions under which the backlog $B^r(t) = 0$ so that a smooth rollout of first-dose appointments (without causing delays in administering the second dose) can be guaranteed under the linear-supply-growth scenario $y(t) = 1 + \alpha t$.

Proposition 1. *Under a linearly growing supply, the stock-release policy ($v_1^r(t), v_2^r(t)$) can allocate β fraction of supply*

to first-dose appointments without incurring backlogs for second-dose appointments if and only if $\beta \leq \frac{1}{2} - \frac{1}{2(1+2\alpha L)}$.

To avoid backlogs at any time t (i.e., $B^r(t) = V_1^r(t-L) - V_2^r(t) = 0$ for any t) so that enough second doses are available in stock for those who have received the first dose L weeks ago, Proposition 1 suggests the stock-release policy can *never* guarantee half of available supplies are allocated to first-dose appointments (i.e., $\beta < 0.5$), no matter how quickly the supply rate grows (i.e., for any value of α). The minimum guaranteed fraction available for first-dose appointments is, intuitively, increasing in α . (Furthermore, as part of the proof of Proposition 1 in the Supporting Information, (OA2)-(OA3) suggest that when α approaches zero, that is, when supply grows slowly, first- and second-dose appointments will have to be almost shut down in alternating periods of lead time (L), in order to avoid backlogs.)

Next, through direct comparison between the performance measure associated with the hold-back (i.e., $V_i^h(t)$ and $P_i^h(t)$) as stated in Lemma 2) and the release policies, we have the following proposition.

Proposition 2. *Under linear supply growth, the stock-release policy ($v_1^r(t), v_2^r(t)$) can generate higher cumulative vaccinated populations than the stock hold-back policy ($v_1^h(t), v_2^h(t)$) by:*

Furthermore, over a full cycle $[0, 2nL)$, where $n \geq 1$, the stock-release policy achieves better cumulative protection than the stock hold-back policy by

$$\begin{aligned}
 \Delta P_1^r(2nL) \triangleq P_1^r(2nL) - P_1^h(2nL) &= \frac{n}{2}L^2 + \frac{\alpha}{2}n(n+1)L^3 \\
 \Delta P_2^r(2nL) \triangleq P_2^r(2nL) - P_2^h(2nL) &= \Delta P_1^r((2n-1)L) \\
 &= \left(\frac{n-1}{2} + \frac{1}{4}\right)L^2 + \frac{\alpha}{2}L^3\left(n^2 - \frac{5}{6}\right).
 \end{aligned} \tag{14}$$

From Proposition 2, we can verify that $\Delta V_i^r(t) > 0$ and $\Delta P_i^r(2nL) > 0$ for $i = 1, 2$. Hence, the net changes in cumulative inoculations as well as cumulative protection due to switching from the stock hold-back policy ($v_1^h(t), v_2^h(t)$) to

the stock-release policy ($v_1^r(t), v_2^r(t)$) are all strictly positive. Therefore, releasing second doses accelerates the buildup of protection, as long as the supply rate is (weakly) growing.

This improvement from releasing second doses, however, comes at the expense of potentially uneven offering patterns that alternate between first and second doses (Proposition 1) and the associated scheduling complexity, especially when the supply rate grows sluggishly. In practice, without assurance that second doses are reserved, many healthcare providers struggle with the uncertainty of scheduling second-dose appointments. In the state of Maryland, as recently as in February 2021, despite the Biden administration's decision not to hold back second doses at the federal level, Dennis Schrader, Acting Maryland Health Secretary, urged healthcare providers to hold back second doses (Miller & Cohn, 2021). In other words, the impact of the stock-release policy on infection control is dampened by its complexity in implementation.

Proposition 2 reveals that the release policy dominates the hold-back policy. Will the stretching policy create additional value over the release policy by extending the lead time? We analytically examine this issue in the section that follows.

5 | WILL STRETCHING THE DOSE TIME INTERVAL HELP?

We now examine the implication of the dose-stretching policy ($v_1^s(t), v_2^s(t)$). Recall from Section 3.3 that the stretching policy is essentially a stock-release policy with a longer lead time that has $L^s = mL$, where $m > 1$. Hence, we can extend our analysis presented in Section 4.2 to determine the corresponding performance measures $V_i^s(t)$ and $P_i^s(t)$ when $L^s = mL$ with $m > 1$. We omit the details to avoid repetition. In addition, through direct comparison between the performance measure associated with the release (i.e., $V_i^r(t)$ and $P_i^r(t)$) as examined in Section 4.2) and the stretching policies, we obtain the following result:

Proposition 3. *Over a cycle $[0, 2mL)$, the dose-stretching policy $L^s = mL$ with $m > 1$ can increase the cumulative first-dose inoculations and reduce the second-dose inoculations over the stock-release policy by*

$$\begin{aligned}\Delta V_1^s(2mL) &\triangleq V_1^s(2mL) - V_1^r(2mL) = \frac{\alpha}{2}L^2m(m-1) > 0, \\ \Delta V_2^s(2mL) &\triangleq V_2^s(2mL) - V_2^r(2mL) = -\frac{3\alpha}{4}L^2(m-1)^2 < 0.\end{aligned}\tag{15}$$

In addition, the dose-stretching policy offers higher cumulative partial and lower full protection than the stock-release

policy by

$$\begin{aligned}\Delta P_1^s(2mL) &\triangleq P_1^s(2mL) - P_1^r(2mL) \\ &= \frac{m(m-1)L^2}{2} + \alpha[2m^3 - m(m+1)]L^3 > 0, \\ \Delta P_2^s(2mL) &\triangleq P_2^s(2mL) - P_2^r(2mL) \\ &= -\frac{3(m-1)^2L^2}{4} - \frac{\alpha(m-1)(m^2+1)L^3}{2} < 0.\end{aligned}\tag{16}$$

Proposition 3 highlights the trade-off between prioritizing the first and second doses. Through stretching the lead time (i.e., by increasing m), more supplies are made available for the first dose; thus, the cumulative population having completed the first dose and the cumulative partial protection both increase. However, deferring the second dose leads to lower levels of cumulative second-dose completions and the level of full protection.

To better explore the trade-off between first and second doses, let us consider the case when the first and second doses lead to efficacy levels of $1 - \theta_1$ and $1 - \theta_2$, respectively, that is, the rate that an otherwise unimmunized person being infected by COVID-19 drops from λ without vaccination to $\theta_1\lambda$ after the first dose, and $\theta_2\lambda$ after the second.¹¹ For notational brevity, let $\zeta = (\theta_2 - \theta_1)/(1 - \theta_1)$. If the first dose provides most of the protection, we have $\zeta < 1$ (e.g., for the case of Pfizer, where the efficacy rates are $1 - \theta_1 = 0.423$ and $1 - \theta_2 = 0.047$ after the first and second doses, we have $\zeta = 0.376/0.577 = 0.652$). Then, the efficacy-weighted protection level $P(t)$ measures the amount of protection provided by the vaccination program weighted by the strength of protection:

$$\begin{aligned}P(t) &= (1 - \theta_1)\lambda P_1(t) + (\theta_1 - \theta_2)\lambda P_2(t) \\ &= \lambda(1 - \theta_1)(P_1(t) + \zeta P_2(t)).\end{aligned}\tag{17}$$

Proposition 4. *Switching from the stock-release policy ($v_1^r(t), v_2^r(t)$) to the dose-stretching policy ($L^s = mL$) can increase the efficacy-weighted protection $P(t)$ over the cycle $[0, 2mL]$ if $\zeta \leq \frac{2m}{3(m-1)}$.*

Proposition 4 suggests dose stretching (i.e., by increasing m) increases efficacy-weighted protection $P(t)$ as long as ζ is below the threshold value $\frac{2m}{3(m-1)}$. The threshold value drops as m increases, indicating dose stretching becomes less favorable as m increases. Recall from Section 3.3 that $m = 4$ under the stretching policy for the Pfizer vaccine when the U.K. government stretched the lead time from 3 weeks to 12. The threshold value in Proposition 4 is $\frac{2m}{3(m-1)} = 0.89$ when $m = 4$. Combine this observation with the fact that $\zeta = 0.652$ for the Pfizer vaccine, the condition $\zeta \leq \frac{2m}{3(m-1)}$ is satisfied

so that the stretching policy adopted by the U.K. government can generate a higher efficacy-weighted protection $P(t)$ over the release policy.

6 | IMPLICATIONS OF AN ALTERNATIVE SINGLE-DOSE VACCINE

Thus far, our model and analysis have been motivated by three approved vaccines developed by Pfizer, Moderna, and AstraZeneca that require two doses. Another COVID-19 vaccine, developed by Johnson & Johnson, was approved for emergency use on February 27, 2021. Unlike Pfizer, Moderna, and AstraZeneca, Johnson & Johnson ran its Phase III trials based on a single-dose regimen, which achieved an efficacy rate of 66%.¹² Without the need to deal with the second-dose timing issue, how would the single-dose vaccine perform relative to those two-dose vaccines? Clearly, if the single-dose vaccine has a higher efficacy, it will dominate the two-dose vaccine. However, the relative performance becomes unclear when the single-dose vaccine (e.g., the Johnson & Johnson vaccine) has a lower efficacy than the two-dose vaccine.

We now compare the rollout dynamics $V_i(t)$ and $P_i(t)$ of the single-dose regimen against the hold-back and stock-release policies in this section. In preparation, observe that when offered as a single-dose vaccine, all stock can be immediately released (because the second dose is no longer needed). Hence, this single-dose regimen can be viewed as an extreme case of the stretching policy with lead time $L^s = mL$, where $m = \infty$. Hence, we can apply the analysis of the stretching policy (see Section 5) to examine the performance of a single-dose vaccine.

Under linear supply growth, the vaccination rate associated with a single-dose regimen satisfies $v^d(t) = y(t) = 1 + \alpha t$. Using the same approach as presented in Sections 4 and 5 to compare the vaccinated population and protection levels for the single-dose regimen with the hold-back and stock-release policies under the two-dose regimen, we obtain the following proposition.

Proposition 5. *Compared with the stock hold-back policy under a two-dose regimen, the single-dose regimen increases the cumulative first-dose inoculations and cumulative first-dose protection over a full cycle $[0, 2nL)$ by:*

$$\begin{aligned} V_1^d(2nL) - V_1^h(2nL) &= nL + \alpha n^2 L^2 > 0, \\ P_1^d(2nL) - P_1^h(2nL) &= n^2 L^2 + \frac{2\alpha}{3} n^3 L^3 > 0. \end{aligned} \quad (18)$$

Compared with the stock-release policy, the single-dose regimen increases the cumulative first-dose inoculations and cumulative first-dose protection over a full cycle $[0, 2nL)$

by:

$$\begin{aligned} V_1^d(2nL) - V_1^r(2nL) &= nL + \frac{\alpha}{2}(2n^2 - n) > 0, \\ P_1^d(2nL) - P_1^r(2nL) &= \frac{n}{2}L^2 + \frac{\alpha}{2} \left[\frac{4}{3}n^3 - n(n+1) \right] L^3 > 0. \end{aligned} \quad (19)$$

From Proposition 5, we can see that the single-dose regimen speeds up the buildup of protection. (Because only one dose is administered under the single-dose regimen, our results focus only on pertinent measures V_1^d and P_1^d .)

However, as articulated above, when the single-dose vaccine has a lower efficacy rate than its two-dose counterparts, the performance of the single-dose vaccine over the two-dose vaccine is unclear. Specifically, even though the two-dose regimen builds up protection more slowly, it can reach a higher level of protection due to its higher efficacy. To evaluate this trade-off, we use the “efficacy-weighted” protection level $P(t)$ for the two-dose vaccine as defined in (17) in Section 5 to define the corresponding $P^d(t)$ for the single-dose vaccine.

Recall from Section 5 that for the two-dose regimen, one and two doses reduce the risk of infection by factors of $1 - \theta_1$ and $1 - \theta_2$, respectively, and $\zeta = \frac{1 - \theta_1}{\theta_1 - \theta_2}$ reflects the relative contribution of the first dose. Let $1 - \theta^d$ be the efficacy rate of the single-dose vaccine. To make a fair comparison with a two-dose vaccine, we define $\zeta^d = \frac{1 - \theta^d}{1 - \theta_1}$. (For example, the Pfizer vaccine has $\theta_1 = 0.423$ and the Johnson & Johnson vaccine has $\theta^d = 0.34$ so that $\zeta^d = 0.66/0.577 = 1.14$.) In addition, by using the results presented in Proposition 5, we can define the corresponding $P^d(t) = P_1^d(t)$ for the single-dose vaccine. The next proposition follows from direct comparisons.

Proposition 6. *Switching from the two-dose regimen under the stock-release policy to the single-dose regimen can increase the efficacy-weighted protection over the cycle $[0, 2nL)$ by*

$$\begin{aligned} \Delta P^d(2nL) \triangleq P^d(2nL) - P^r(2nL) &= \lambda(1 - \theta_1)L^2 \\ &\times \left[2n^2(\zeta^d - 1) + \frac{n-1}{2} + \left(\frac{1}{2} - \frac{3\zeta}{4} \right) (2n^2 - 2n + 1) \right] \\ &+ \lambda(1 - \theta_1) \frac{\alpha L^3}{12} [(16\zeta^d - 16)n^3 - 12n + 6 \\ &+ (1 - \zeta)(8n^3 - 6n^2 + 6n - 6)]. \end{aligned} \quad (20)$$

A sufficient condition for this improvement to be nonnegative is that $\zeta^d \geq \frac{11}{8}$ and $\zeta \leq \frac{2}{3}$.

Recall from above that for the two-dose regimen of the Pfizer vaccine, we have $\zeta = 0.652$, and for the single-dose regimen of the Johnson & Johnson vaccine, $\zeta^d = 1.14$. Even though these values do not satisfy the sufficient condition in Proposition 6, we can compute:

$$\begin{aligned} \Delta P^d(2nL) &= \lambda(1 - \theta_1)L^2(0.5n^2 + 0.58n - 0.54) \\ &+ \lambda(1 - \theta_1)\frac{\alpha L^3}{12}(6.56n^3 - 1.68n^2 - 10.32 + 4.32). \end{aligned} \quad (21)$$

We can check that the first term is positive for all $n \geq 1$ and that the second term is negative when $n = 1$ and positive for $n \geq 2$. When $n = 1$, the overall expression is positive when $\alpha L \leq 5.36$. Because the supply rate is unlikely to grow more than 5.36 times over lead time L , we are more likely to observe that the Johnson & Johnson single-dose regimen offers better cumulative efficacy-weighted protection than the Pfizer vaccine under the stock-release policy for any n . Overall, the above analysis suggests that the single-dose regimen can underperform in the short run, but will eventually outperform the higher efficacy two-dose regimen.

So far, we have obtained three key analytical results. First, we show that the release policy can improve the cumulative vaccination population and the cumulative protection in Proposition 2. Second, we prove that the stretching policy can generate a higher efficacy-weighted protection $P(t)$ than the release policy in Proposition 4 when $\zeta \leq \frac{2m}{3(m-1)}$. Third, we show that a single-dose vaccine with a lower efficacy can outperform a two-dose vaccine in terms of a higher efficacy-weighted protection when $\zeta^d \geq \frac{11}{8}$ and $\zeta \leq \frac{2}{3}$ in Proposition 6. Besides the performance measures $V(t)$ and $P(t)$, one should evaluate and compare the symptomatic infection rate, hospitalization rate, and the mortality rate of different groups of people across all policies over time. To do so, we present a variant of the SEIR epidemic model to examine these issues numerically in the next section.

7 | EVALUATING ROLLOUT POLICIES VIA AN SEIR MODEL

In this section, we simulate the pandemic trajectories under the four rollout policies: (1) the stock hold-back policy ($v_1^h(t), v_2^h(t)$), (2) the stock-release policy ($v_1^r(t), v_2^r(t)$), (3) the dose-stretch policy ($L^s = mL$ with $m > 1$), and (4) the alternative single-dose vaccine. We begin by presenting our SEIR epidemic model in Section 7.1. In Section 7.2, we simulate the rollout policies for the “baseline case” in which the parameter values are listed in Table OAI of the Supporting Information. To analyze the robustness of our simulation results, we conduct sensitivity analyses with respect to age composition in Section 7.3, risk-based vaccine prioritization in Section 7.4, vaccine supply disruption in Section 7.5, and the basic reproductive number (R_0) in Section 7.6. We caution that, as our simulation parameters are calibrated to data

obtained on the original strain of the vaccines and the earlier (e.g., Alpha) variants, the results do not necessarily reflect the effects of future variants. Yet, the same methodology developed in our paper supports new analysis using actual or estimated parameters associated with future variants.

7.1 | SEIR model

We present a compartmental SEIR epidemic model to capture the evolution of the COVID-19 pandemic and the vaccine rollout dynamics. Using this SEIR model, we evaluate different rollout policies by comparing three performance metrics, including symptomatic infections, hospitalizations, and mortality. These metrics provide an epidemiological perspective and complement the operational perspective presented in Sections 4 to 6.

7.1.1 | Risk groups and vaccination states

To capture various age-dependent factors such as hospitalization rate, mortality rate due to COVID, and frequencies of social contact, we consider a population of size N (e.g., the U.S. adult population) that consists of a high-risk group H (e.g., the population aged 65 and above) of size N_H and a low-risk group L (e.g., the population under the age of 65) of size N_L ; this setting is consistent with the literature on modeling compartments with age groups (e.g., Longini et al., 1978). Based on the U.S. demographics, we normalize the adult population to 100 and set N_H and N_L to 13 and 87 (according to the profile of the adult population in the United States), respectively. To incorporate the vaccination status of the two-dose regimen, we denote by (i, k) each individual’s type, where $i \in \{H, L\}$ is the risk group and $k \in \{0, 1, 2\}$ is the number of vaccine doses that the individual has received.

7.1.2 | Compartments

In line with Keskinocak et al. (2020), an individual of type (i, k) is in one of the following eight compartments at any time t : susceptible ($S_{ik}(t)$), exposed ($E_{ik}(t)$), infectious and presymptomatic ($IP_{ik}(t)$), infectious and asymptomatic ($IA_{ik}(t)$), infectious and symptomatic ($IS_{ik}(t)$), hospitalized ($H_{ik}(t)$), recovered ($R_{ik}(t)$), or deceased ($D_{ik}(t)$). By considering this expanded set of states instead of the standard infectious and recovered compartments, our model captures key features of the COVID-19 pandemic such as *asymptomatic transmission* and the demand on the healthcare system due to *hospitalization*. With two risk groups (i), three vaccination states (k), and eight compartments for each individual type (i, k) , our SEIR model comprises a total of 48 compartments. Figure 1 illustrates our SEIR model for group $i \in \{H, L\}$.¹³

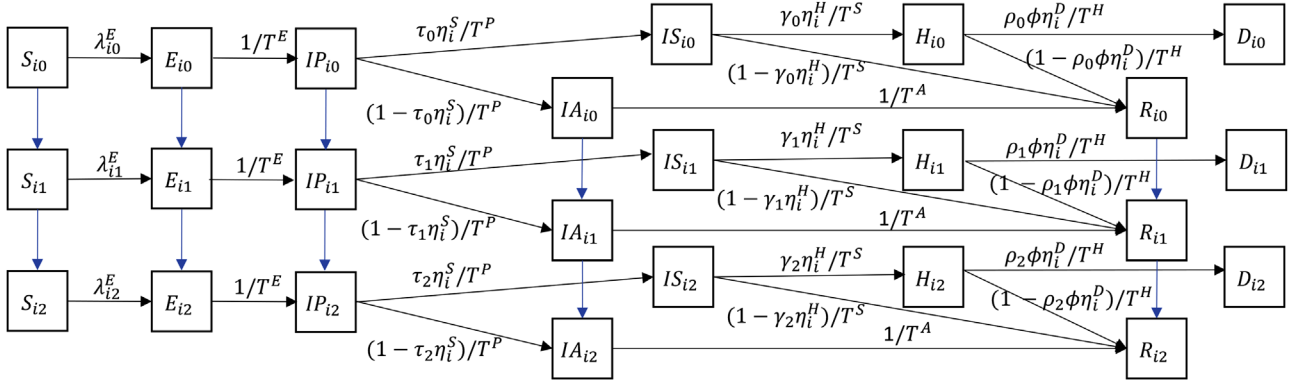


FIGURE 1 An expanded SEIR model for group $i \in \{H, L\}$. The downward arrows (in blue) indicate the change in vaccination status. $\phi = 1$ if $\sum_{i,k} H_{ik} \leq K$ and $\phi = \kappa$ otherwise. [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

7.1.3 | Transition between compartments

For each individual of type (i, k) , we explain in Section OA2 of the Supporting Information about how we determine the transition rate from one compartment to a different compartment as depicted in Figure 1. In particular, COVID transmissions depend on the numbers of infectious and susceptible individuals, their transmissibility and susceptibility (which depend on their vaccination status), and the frequencies of social contact between the two age groups as governed by a “contact matrix” (Bubar et al. 2021).

After contracting the virus, individuals will go through the exposed and presymptomatic stages, at the end of which they will either develop symptoms or remain asymptomatic, governed by the age- and vaccination-dependent probabilities. Those who are symptomatically infected will be hospitalized (and if so, recover or die) according to the age- and vaccination-dependent probabilities. We calibrate different age- and vaccination-dependent parameter values using the CDC data and the recent literature and summarize the values in Table OA1 of the Supporting Information.

In addition to the transition rates among different compartments, our SEIR model depicted in Figure 1 captures the impact of the healthcare system utilization on mortality rate. Specifically, as more patients are hospitalized, critical facilities (e.g., intensive care units) become overloaded and mortality rates increase substantially (Wilde et al., 2021). Details of our mortality rate formulation are provided in Section OA2 of the Supporting Information.

7.1.4 | The role of vaccination

To evaluate different vaccine rollout policies via the SEIR model for an individual type (i, k) as depicted in Figure 1, we can use the vaccination rates $v_k(t)$ ($k = 1, 2$) under each rollout policy presented in Section 4 to 6 as the key inputs for our SEIR model. In doing so, we can determine the corresponding transition rates between compartments associated with different vaccination status (e.g., from $S_{i0}(t)$ to $S_{i1}(t)$ after

receiving the first vaccine dose). Based on the system dynamics captured by our SEIR model, we simulate the following three performance metrics for each risk group $i = H, L$, under all four aforementioned rollout policies: (1) the number of active symptomatic infections: $IS_i(t) = \sum_{k=0}^2 IS_{ik}(t)$, (2) the number of active hospitalizations: $H_i(t) = \sum_{k=0}^2 H_{ik}(t)$, and (3) the cumulative number of deaths: $D_i(t) = \sum_{k=0}^2 D_{ik}(t)$. In the following subsections, we compare the simulated performance metrics $IS_i(t)$, $H_i(t)$, and $D_i(t)$ via our SEIR model across all four rollout policies.

We simulate our SEIR model over a 1-year horizon. The parameter values and the settings of the simulation are summarized in S. Consistent with our theoretical analysis in the previous sections, we consider a supply function that is growing linearly unless otherwise specified.

7.2 | Baseline case

We begin our numerical examples by simulating the performance of the rollout policies under the baseline case; all the parameter values are listed in Table OA1 of the Supporting Information.

Recall from Section 4 to 6 that $v_1(t)$ and $v_2(t)$ denote the rate of vaccine doses to be given as first and second doses at time t , under the hold-back, release, stretch, or single-dose policies. To determine who receives the vaccines, we consider a simple proportional allocation rule between the high- and low-risk groups, for example, $\pi \cdot v_k(t)$ and $(1 - \pi) \cdot v_k(t)$ of dose- k supplies ($k = 1, 2$) are allocated to the high- and low-risk groups at any time t , respectively. In the baseline case, we have $\pi = 1$ (i.e., 100% of supply is given to eligible high-risk individuals until they have all been vaccinated). This is consistent with the strict age-based rollout practice in countries such as the United States and the United Kingdom. In Section 7.3, we shall consider an alternative case with a lower value of π for comparison.

Figure 2 shows the SEIR simulation results for the baseline case when we scale the population size $N = 100$. The panels plot the number of active symptomatic cases (left),

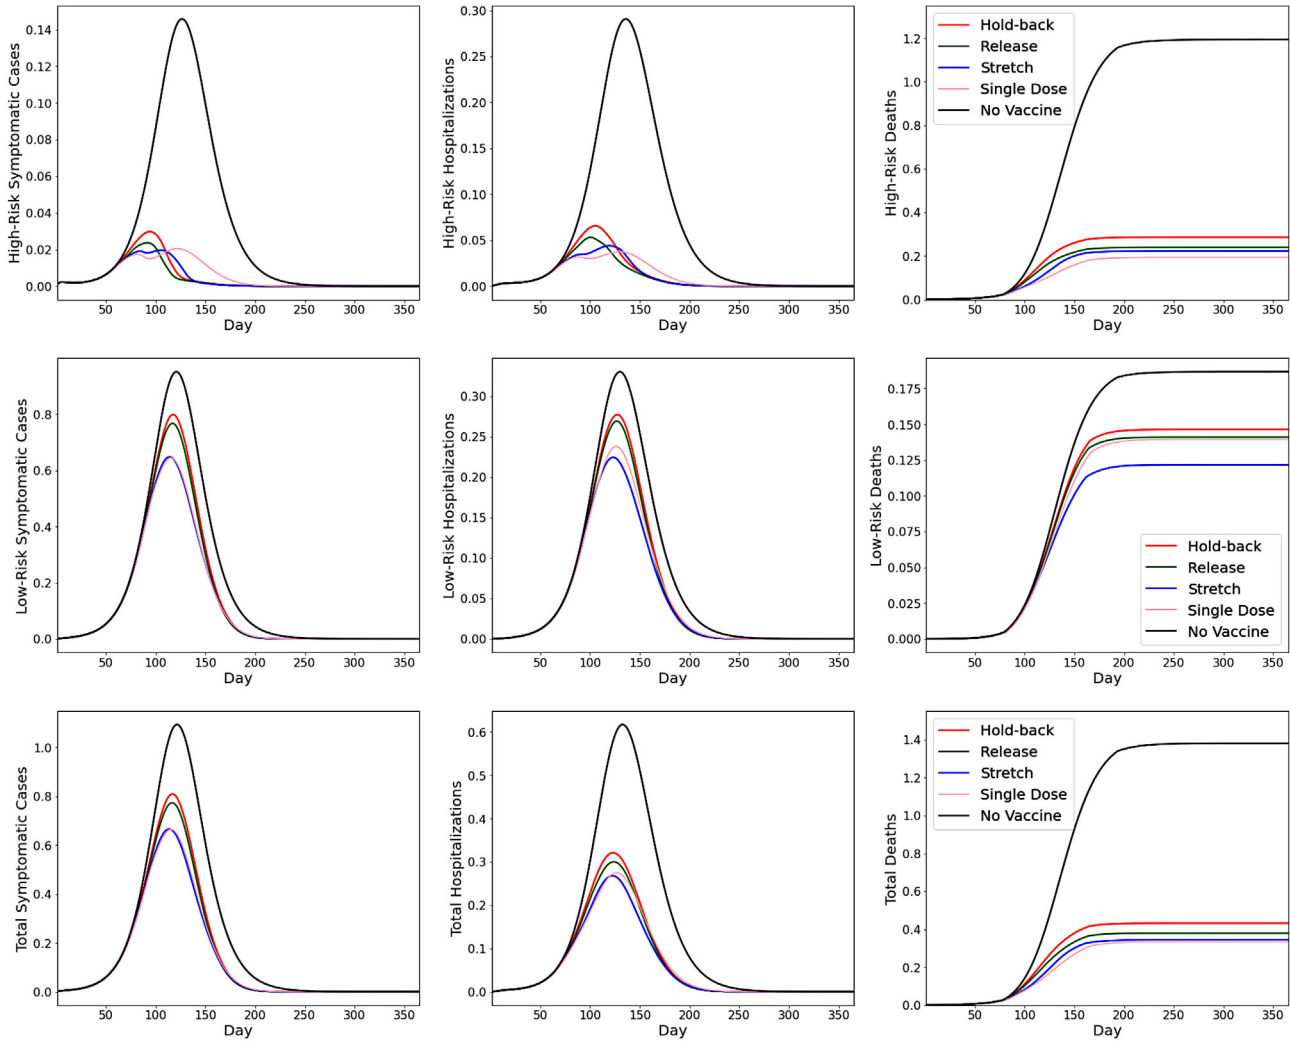


FIGURE 2 SEIR simulation results under baseline case [Color figure can be viewed at wileyonlinelibrary.com]

hospitalizations (middle), and cumulative deaths (right), for the high-risk group (top), low-risk group (middle), and total population (bottom). Because the size of the population is normalized to 100, the vertical axes represent the percentage of the total population in each status of interest. In each panel, the four rollout policies are compared against the no-vaccination benchmark (solid line in black).

From the top panels of Figure 2, we find that vaccination is highly effective in reducing symptomatic cases, hospitalizations, and deaths for the high-risk group. Under all four rollout policies, each of these key metrics are reduced by at least 75%, compared with the no-vaccination benchmark. In comparison, the corresponding reductions are smaller for the low-risk group, due to prioritizing the high-risk group and the overall larger population of the low-risk group (thus taking longer to get vaccinated). For the whole population, due to strict priority of vaccinating the high-risk group (i.e., $\pi = 1$), the reduction in deaths (i.e., 69%, 72%, 75%, and 76% for the hold-back, release, stretch, and single-dose policies, respectively) are more substantial than the reductions in hospitalizations (i.e., 48%, 51%, 57%, and 51%, respec-

tively) and in symptomatic infections (i.e., 26%, 29%, 39%, and 39%, respectively).

Consistent with the implications from our analytical findings, Figure 2 shows that both the stretching policy and the single-dose regimen are effective in accelerating vaccine impact and reducing symptomatic infections, hospitalizations, and mortality. The release policy also outperforms the hold-back policy in all three metrics, although the difference is modest in the baseline case.

7.3 | Age composition

We now evaluate how the demographic composition may influence the choice of rollout policies. Since hospitalization and mortality risks of COVID are highly dependent on age, the pandemic impact on relatively younger (e.g., in India and South Africa, with a median ages of about 28) and older (e.g., in the United States, with median age of about 38) populations can differ significantly. All else being the same, we now change N_H to half as in the baseline case (so that 6.5% of the

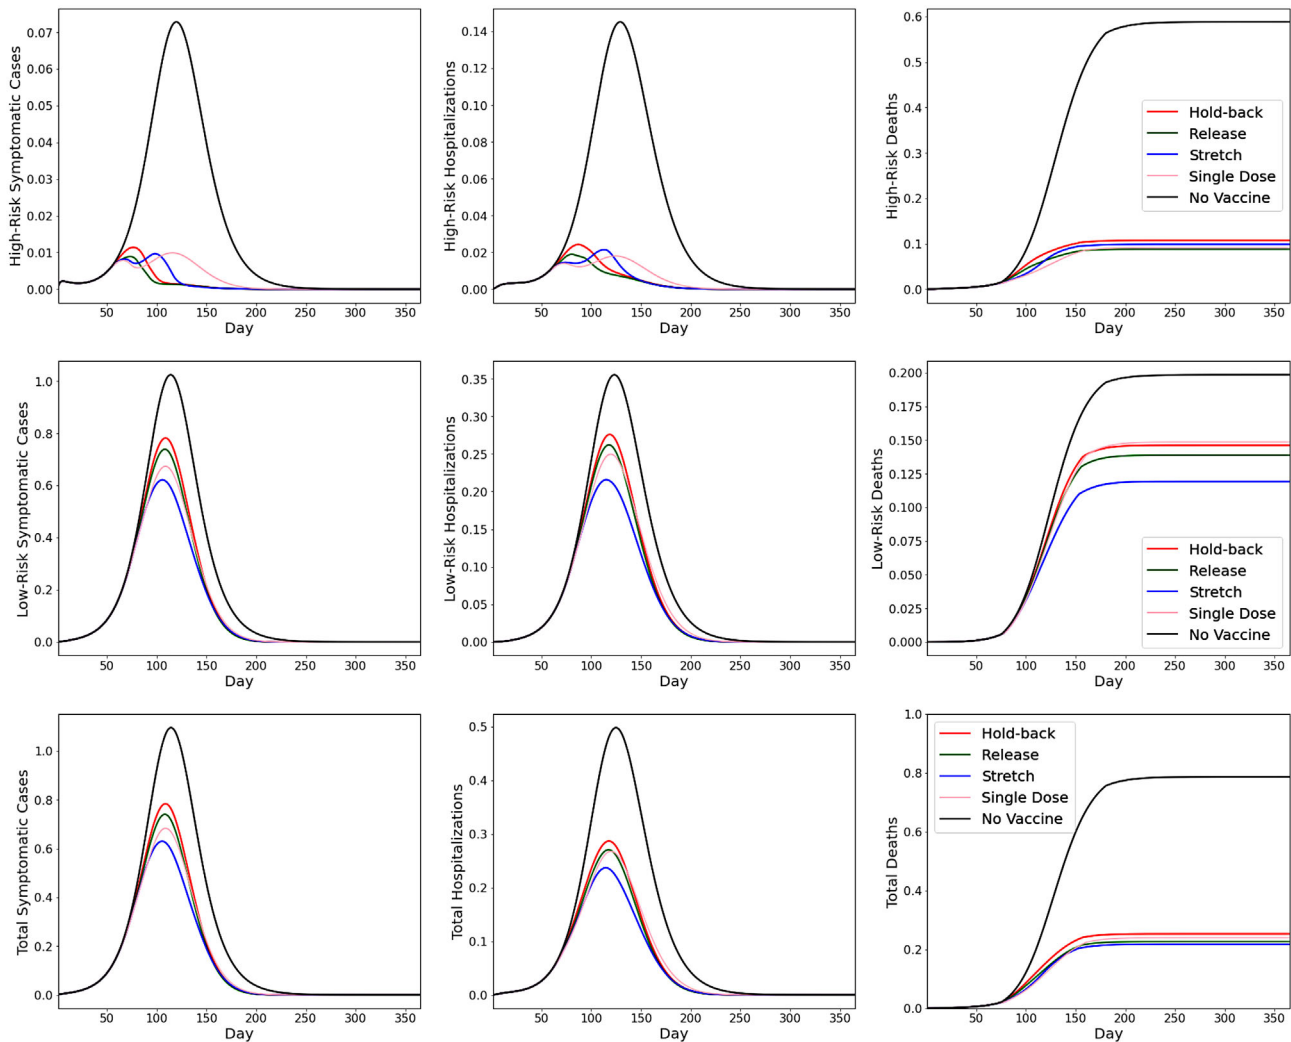


FIGURE 3 SEIR simulation results with younger population [Color figure can be viewed at wileyonlinelibrary.com]

total population is above 65, compared with 93.5% below 65) so that the resulting population is younger. Figure 3 depicts our results.

Figure 3 resembles the baseline case as reported in Figure 2 in terms of, for example, the effectiveness of the stretching policy and the single-dose regimen. Observe from Figure 3 that even in a younger population with a 6.5% high-risk population, vaccination reduces mortality significantly—the four policies achieve a mortality reduction from 68% (hold-back policy) to 72% (stretching policy). Despite similar percentage reductions in mortality compared to the baseline case (Figure 2), the absolute number of deaths is considerably lower because we consider a younger population. However, this is not the case for the other two key metrics, as the younger population records similar absolute levels of symptomatic cases and hospitalizations as the older population. This is due to the fact that most symptomatic cases and about half of hospitalizations are associated with the low-risk group due to its higher transmission rate (albeit lower mortality rate).

Interestingly, Figure 3 shows that, relative to the stretching policy, the single-dose regimen does not perform as well when the population is younger. Specifically, the single-dose regimen reduces peak symptomatic cases, peak hospitalizations, and cumulative mortality by 38%, 46%, and 69%, respectively. However, it is outperformed by the stretching policy, which reduces the three key metrics by 43%, 52%, and 72%, respectively. In terms of preventing mortality, the single-dose regimen performs worse than the release policy (71% reduction) and is similar to the hold-back policy (68% reduction). The reason is that, when the high-risk group is relatively small, the main advantage of the single-dose regimen (i.e., rollout speed) becomes less appealing, whereas its weakness (i.e., lower efficacy, especially in preventing infections) becomes more pronounced.

This observation has important practical implications. Prior to the halt of the Johnson & Johnson vaccine in the United States due to blood clot concerns, its “one-and-done” logistical convenience was considered highly promising for younger age groups (Dai et al., 2021). Various countries

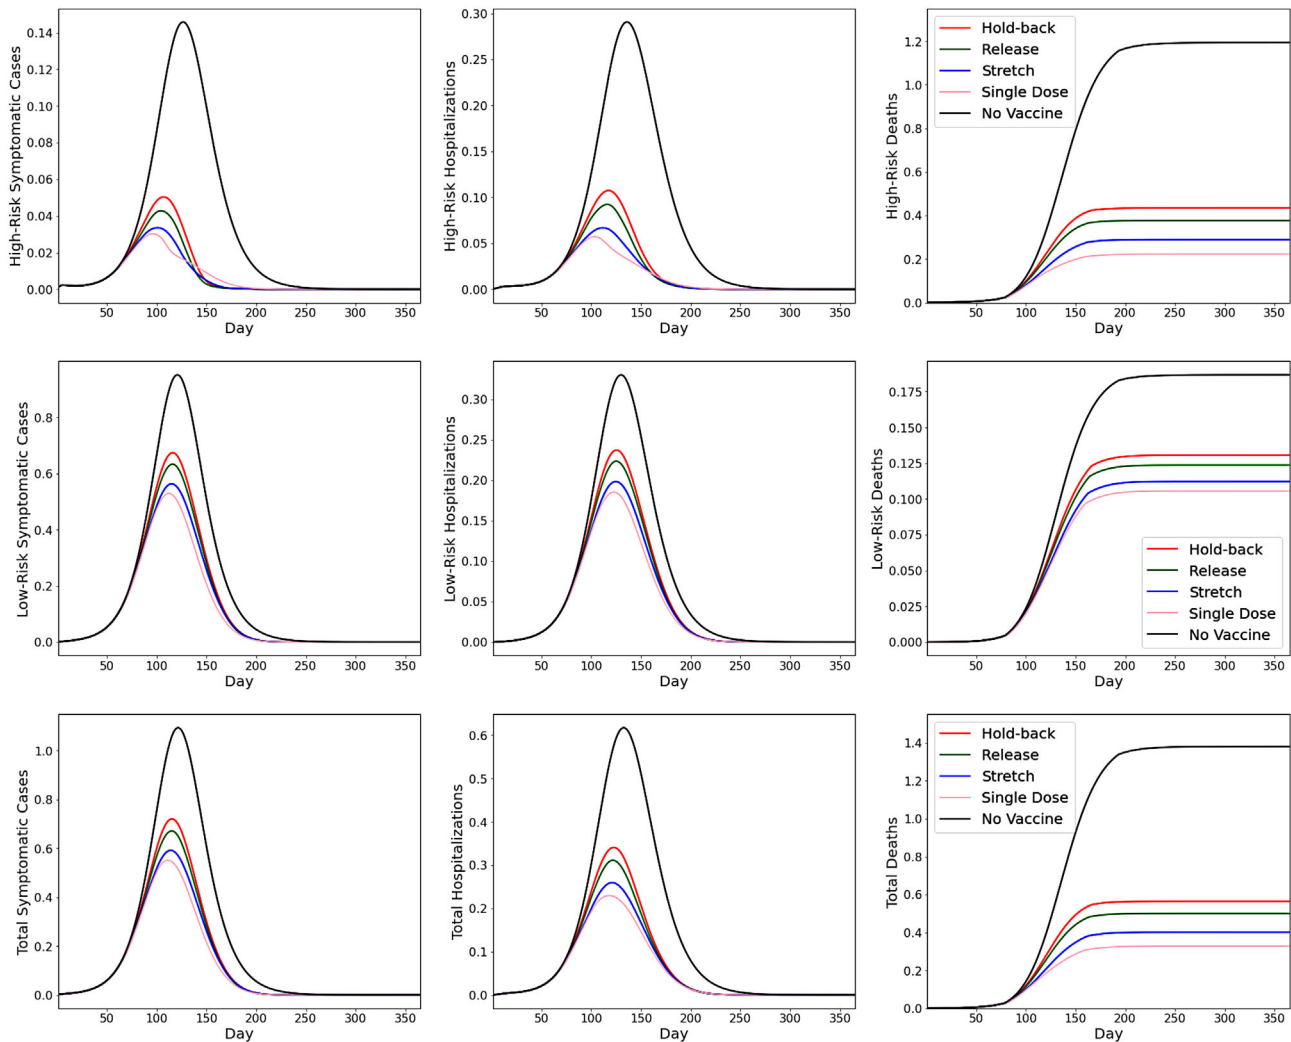


FIGURE 4 SEIR simulation results with $\pi = 0.5$ [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

have since refocused the rollout of Johnson & Johnson to older age groups on the basis of lower blood clot risks (Strauss, 2021). Our findings suggest that doing so is sensible logistically and epidemiologically as well, since the single-dose regimen would perform better in an older population.

7.4 | Risk-based prioritization

In the baseline case (Figure 2), we have set $\pi = 1$ (i.e., 100% of supply is given to eligible high-risk individuals until they have all been vaccinated). In this subsection, we compare this with an alternative case that has $\pi = 0.5$ (i.e., half of the supply is allocated to the high-risk group of age above 65 and the remaining half is allocated to the low-risk group of age below 65). Because the low-risk group has higher frequencies of social contacts than the high-risk group, the alternative case is intended to examine the impact of allocating some vaccine to the low-risk group earlier as a mechanism to curb the infection rates.

Figure 4 plots the simulation results for the alternative case with $\pi = 50\%$ priority for the high-risk group. Despite a different vaccine priority rule, Figure 4 resembles Figure 2 in that the single-dose regimen and stretch policies are the most effective in reducing symptomatic infections, hospitalizations, and mortality. In addition, we observe from Figure 4 that the effects of vaccination are abated for the high-risk group and improved for the low-risk group under 50% priority, compared with the baseline case (Figure 2), which strictly prioritizes the high-risk group. Because most deaths arise from the high-risk group, this leads to higher overall mortality (bottom-right panel) under the release and hold-back policies. On the other hand, because most symptomatic cases arise from the low-risk group that has a larger population size and more social contacts, a less stringent risk-group priority leads to fewer symptomatic cases, and a similar or lower number of hospitalizations.

Our findings reflect an important trade-off. Notwithstanding the importance of protecting the most vulnerable people by prioritizing the high-risk group, a strict priority (e.g., the appointment system in the United Kingdom, where younger

age groups cannot make bookings until the majority of older age groups receive their vaccines) can be inefficient in containing the spread of the pandemic due to the more frequent social contacts of younger age groups. This is because the vaccines (e.g., Pfizer and AstraZeneca) have significant effects on not only preventing infections, but also reducing infectiousness of those who get infected despite vaccinations, as found by Public Health England (Harris et al. 2021) in the U.K. rollout. Many countries and regions, including various states of the United States, also prioritize the vaccination of younger individuals in high-contact essential professions, such as public transit drivers, delivery drivers, and postal service workers, together with the elderly; vaccinating them helps prevent secondary infections.

Furthermore, we notice that the differences (in all three key metrics) between the rollout policies become more pronounced under less stringent (50%) priority. In particular, the single-dose regimen and the stretching policy, which are effective in accelerating vaccine coverage, become relatively more beneficial when supply is split between the risk groups. This is intuitive, because splitting supply leads to fewer doses available for the high-risk group, which in turn heightens the importance of rollout efficiency. On the whole, we find that both the stretching policy and single-dose regimen offer good trade-offs under the 50% priority case: These policies lead to lower symptomatic cases (46% and 50% reduction from the no-vaccination benchmark, respectively, compared with 39% and 39% under strict priority), and a similar number of hospitalizations (58% and 63% reduction from the no-vaccination benchmark, respectively, compared with 57% and 55% under strict priority) through preventing transmissions, and record similar mortality figures (71% and 76% reduction from the no-vaccination case, compared with 75% and 76% under strict priority) thanks to accelerated build-up of protection for the high-risk group.

7.5 | Supply disruptions

So far, our analytical model and simulation analysis are based on the assumption that the supply process is smooth. However, the supply of the COVID-19 vaccines is subject to disruptions to the production¹⁴ or distribution process.¹⁵ We conduct additional simulations to examine the performance of all four rollout policies in the presence of potential supply disruptions.

We simulate a scenario with production hiccups occurring at the initial stage of vaccine rollouts. In particular, the supply rate is assumed to be growing linearly, as explained earlier. However, the supply rate $y(t)$ suffers from two disruptions occurring on Days 81–110 and Days 141–170, during which the supply rate is decreased to half of the initial rate. Figure 5 depicts our simulated results.

Figure 5 shows that under supply disruptions, the results obtained from the baseline case continue to hold: The single-dose regimen and stretch policies are most effective. Because rollouts prioritize the high-risk group, it is natural that the

improvements for the low-risk group are limited with lumpy supply. Overall, substantial reduction in mortality can still be achieved, especially under the stretching policy or the single-dose regimen policy.

We also observe that the qualitative insights from the baseline case with smooth supply (Figure 2) remain intact; that is, the single-dose regimen and stretch policies help prevent mortality better (by 47% and 44% compared with the no-vaccination case) than do the release and hold-back policies (39% and 34% reductions, respectively). These figures suggest that the performance gaps between the policies become more pronounced during supply disruptions, because making full use of the limited (and unreliable) supply to build up protection becomes more essential.

7.6 | Basic reproduction number (R_0)

Finally, we evaluate the impact of the speed at which the pandemic propagates, as measured by the basic reproduction number R_0 . In the COVID-19 pandemic, the R_0 number could change due to the rise of more transmissible variants or the relaxation of social distancing measures. In Figure 6, we conduct our simulation based on the case when $R_0 = 1.8$, which is higher than the baseline case that has $R_0 = 1.5$ (Figure 2).

Figure 6 reveals that whereas all key metrics have worsened under faster transmissions, the results obtained from the baseline case continue to hold qualitatively. Figure 6 also offers insights similar to those from the case with supply disruptions (Figure 5). In particular, vaccine rollouts mostly improve the key metrics for the high-risk group and show limited benefits for the low-risk group. Overall, we observe significant improvements in mortality reduction—by 57%, 52%, 42%, and 36% under the single-dose, stretch, release, and hold-back policies, respectively. Comparing these figures with those in the baseline case (which range from 69% to 76% reduction across the policies), we observe that under a higher R_0 , choosing the right rollout policy has a greater impact on reducing mortality.

The reason that the high R_0 and the supply disruption cases are qualitatively similar is that vaccination outcomes depend on the *relative* speeds of the pandemic and of vaccine rollout. When supply is lumpy or when R_0 is high, vaccination rollout lags behind the pandemic spread, and thus the rollout policies that accelerate protection (i.e., single-dose and stretching policies) perform relatively better in preventing mortality of the high-risk group and thus the total population.

8 | CONCLUSIONS

Rolling out a two-dose COVID-19 vaccine such as that developed by Pfizer offers a rare and compelling research context for examining the impact of different inventory control policies (i.e., hold-back, releasing, and stretching) on public health. We develop a modeling framework of a vaccine rollout process and complement it with an SEIR model to capture




FIGURE 5 SEIR simulation results with supply disruptions [Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

the disease transmission process. We compare—both analytically and numerically—various vaccine rollout policies, including (1) holding back all the second doses, (2) releasing all the doses, and (3) stretching the lead time between two doses. We also examine the rollout of a single-dose vaccine with a lower efficacy level than its two-dose counterparts.

Our key findings are threefold. First, one might expect releasing all the second doses would lead to an oversized effect in slowing down the pandemic. Our analytical results show that compared with the hold-back policy, releasing doses does help slow down infections and end the pandemic sooner; we numerically show that such an impact can be modest. This result is in part due to the fact that releasing more first doses also necessitates more second doses in the future; we show that even if one intends to release all second doses, *fewer than half* of the available doses can be used for first-dose appointments to avoid delays in administering the required second doses, regardless of how fast supply grows over time. Second, stretching the between-dose lead time helps reduce mortality as well as infections and hospitalizations (more than under the strategy of releasing second

doses). Third, a single-dose vaccine, even with an overall efficacy significantly lower than its counterparts, can be more effective in infection control.

It should be acknowledged that the logistics parameters of vaccine regimens (i.e., one or two doses, recommended lead time) depend on the design of trials and the data presented for regulatory approval. For instance, the fact that Pfizer requires two doses but Johnson & Johnson requires only one is primarily due to how the respective trials were designed and thus how the specific regimens were approved, but is not necessarily reflective of the differences in the *science* of the vaccines.¹⁶ In this sense, our results underline the importance of taking the *operational* aspect of the rollout process into consideration when designing clinical trials.

Our model has several limitations that can serve as subjects for future research. First, COVID-19 vaccine rollouts inspire future research into a wide range of operational and logistics issues specific to vaccine administration, including cold chain management, inventory management, and manufacturing capacity planning (Dai & Song, 2021). One particularly interesting area of research entails incorporating and




FIGURE 6 SEIR simulation results with $R_0 = 1.8$ [Color figure can be viewed at wileyonlinelibrary.com]

managing the uncertainty associated with the number of doses that can be extracted from each vial. Second, as Mak and Tang (2021) note, no-shows for second-dose appointments may worsen when the time interval between two doses is stretched from 3 weeks to 12 weeks for a variety of reasons, including, for example, the perception that the second dose provides limited incremental benefit due to the stretching. To incorporate potential no-shows, one may extend our SEIR model by adding a new parameter capturing the reduction in the second-dose vaccination rate. Third, we consider in this paper a time horizon long enough that the pandemic subsides. One related matter that we do not consider is the possibility of waning vaccine effects, which could affect infections beyond the time horizon. Lastly, several highly contagious new variants of SARS-CoV-2, the virus causing COVID-19, have been found in various countries. Examining other rollout policies that can expedite herd immunity even more quickly to reduce the spread and mutation of the virus would be of interest.

Our modeling framework represents an initial attempt to examine different vaccine rollout policies under limited

supply. In addition to helping inform vaccine rollout policy-making, our modeling framework can be applied to analyze similar issues facing local vaccination planners and health systems leadership in both the current and future epidemics. At the global level, our work draws from the experience of the E.U., the United Kingdom, and the United States. Yet, our results have important implications for vaccine planning in various countries, especially because developing countries are expected to face extremely limited vaccine supply up to 2023 (Hopkins, 2021). More broadly, our paper sheds light on how to develop effective operations strategies for distributing time-sensitive resources in times of crisis.

ACKNOWLEDGMENTS

The authors are grateful to Stephen Chick, Eric Denardo, Edward Kaplan, Steven Lippman, Elisa Long, John Mamer, David Paltiel, Andrew Pollard, Prashant Yadav, seminar and session participants at INFORMS 2021 Annual Meeting, INSEAD, The 3rd Johns Hopkins Symposium on Healthcare Operations, McGill University, and Rutgers University for their helpful comments and suggestions.

ORCID

Ho-Yin Mak  <https://orcid.org/0000-0002-4289-8029>

Tinglong Dai  <https://orcid.org/0000-0001-9248-5153>

Christopher S. Tang  <https://orcid.org/0000-0001-9597-7620>

ENDNOTES

- ¹Throughout the paper, for brevity, we refer to the Pfizer/BioNTech vaccine as the Pfizer vaccine. Similarly, we refer to the AstraZeneca/Oxford vaccine as the AstraZeneca vaccine.
- ²The U.S. Food and Drug Administration (FDA) has only authorized the use of two specified doses of both Pfizer and Moderna vaccines at specified intervals—a 3-week interval for the Pfizer vaccine and a 4-week interval for the Moderna vaccine. See <https://www.fda.gov/news-events/press-announcements/fda-statement-following-authorized-dosing-schedules-covid-19-vaccines> for details.
- ³Besides tractability, the deterministic production rate is a reasonable assumption for the following reason: The production process for the Pfizer and the Moderna vaccines is relatively stable, because the process involves converting DNA into mRNA rather than growing cells in a bioreactor (Boseley, 2021).
- ⁴The recommended lead time $L = 21$ days for the Pfizer vaccine, and CDC guidelines in 2021 allow for a grace period of 4 days earlier than the recommended date, see <https://www.cdc.gov/vaccines/covid-19/info-by-product/clinical-considerations.html> for details. In February 2022, the CDC modified its recommended lead time to 8 weeks for 12- to 64-year-olds, particularly males aged 12 to 39 (<https://khn.org/morning-breakout/cdc-advising-some-to-wait-8-weeks-between-doses-of-mrna-vaccines/>). Despite these new guidelines, our modeling framework and approach are applicable to any lead time L and any number of doses.
- ⁵To isolate the effect of different vaccine rollout policies, we do not model the issue of setting up a physical infrastructure (i.e., location and the size of vaccination sites) or mobilizing personnel (controlled by FEMA, state/county governments, or local healthcare providers) to administer vaccines. Instead, we abstract away those issues and assume all policies will be operated under the same environment.
- ⁶We use the superscripts h , r , and s to denote the hold-back, release, and stretching policies, respectively.
- ⁷Though small amounts of smoothing stock will still be kept in reality, we ignore this fact in our model for simplicity.
- ⁸This alternating pattern has been observed in various parts of the world. For example, in early February 2021, COVID-19 vaccine sites in Los Angeles County postponed first-dose appointments and only administered second doses to returning recipients; see <http://cbsloc.al/3bo5uiS> for details.
- ⁹In the case in which $m = 1$, the stretching policy reduces to the release policy. For practicality of implementation, we restrict m to being integer-valued. For example, the U.K.'s dose-stretching policy increases the lead time from $L = 3$ weeks to $L^s = 12$ weeks, that is, $m = 4$.
- ¹⁰For tractability, we consider the linear case. However, one can use the same approach to examine the case of a nonlinear production rate numerically.
- ¹¹For example, the Pfizer vaccine offers 57.7% and 95.3% efficacy levels after the first and second doses; thus, $\theta_1 = 0.423$ and $\theta_2 = 0.048$ (Haas et al., 2021).
- ¹²In late January 2021, Johnson & Johnson announced the Phase III trial results of its single-dose COVID-19 vaccine with a relatively lower efficacy (66% worldwide and 72% in the U.S. trial). However, this vaccine has major operational advantages. In particular, it has a less stringent temperature requirement than the Pfizer vaccine and only requires a single dose. Thus, it circumvents the trade-off between offering first and second doses.
- ¹³We assume the time horizon to be short enough so that individuals do not move from the low- to high-risk group and vice versa (their underlying health status does not change). Therefore, the two groups interact only through sharing the same vaccine supply and hospitalization capacity.

¹⁴For example, the AstraZeneca vaccine faces an uncertain yield due to uneven cell growth in bioreactors.

¹⁵Disruption can occur due to weather conditions as experienced in the United States during the snow storm in February.

¹⁶The authors acknowledge Andrew Pollard of the University of Oxford for this important insight.

REFERENCES

- Arifoğlu, K., Deo, S., & Irvani, S. (2012). Consumption externality and yield uncertainty in the influenza vaccine supply chain: Interventions in demand and supply sides. *Management Science*, 58(6), 1072–1091.
- Arifoğlu, K., & Tang, C. S. (2022). A two-sided incentive program for coordinating the influenza vaccine supply chain. *Manufacturing & Service Operations Management*, 24(1), 235–255.
- Arnold, C. (2020). The biggest logistics challenge in history. *New Scientist*, 248(3309), 36–40.
- Arrow, K. J., Harris, T., & Marschak, J. (1951). Optimal inventory policy. *Econometrica*, 19(3), 250–272.
- Boseley, S. (2021). Why has AstraZeneca reduced promised vaccine supply to EU and is UK affected? *Guardian* (Jan 26). <http://bit.ly/azredeuuk>
- Bubar, K. M., Reinholt, K., Kissler, S. M., Lipsitch, M., Cobey, S., Grad, Y. H., & Larremore, D. B. (2021). Model-informed COVID-19 vaccine prioritization strategies by age and serostatus. *Science*, 371(6532), 916–921.
- Chen, M. K., Chevalier, J. A., & Long, E. F. (2020). Nursing home staff networks and COVID-19. *Proceedings of the National Academy of Sciences*, 118(1), e2015455118.
- Chick, S. E., Mamani, H., & Simchi-Levi, D. (2008). Supply chain coordination and influenza vaccination. *Operational Research*, 56(6), 1493–1506.
- Chick, S. E., Hasija, S., & Nasiry, J. (2017). Information elicitation and influenza vaccine production. *Operational Research*, 65(1), 75–96.
- Cho, S.-H., & Tang, C. S. (2013). Advance selling in a supply chain under uncertain supply and demand. *Manufacturing & Service Operations Management*, 15(2), 305–319.
- Cummings, K., Behzad, B., & Martonosi, S. (2021). Centers for Disease Control and Prevention as a strategic agent in the pediatric vaccine market: An analytical approach. *Manufacturing & Service Operations Management*, 23(6), 1398–1412.
- Dai, T., Cho, S.-H., & Zhang, F. (2016). Contracting for on-time delivery in the US influenza vaccine supply chain. *Manufacturing & Service Operations Management*, 18(3), 332–346.
- Dai, T., Mak, H. Y., & Tang, C. S. (2021). The great promise of a one-dose vaccine. *Barron's* (February 26, 2021). <http://bit.ly/barronsnjn>
- Dai, T., & Song, J. S. (2021). Transforming COVID-19 vaccines into vaccination: Challenges and opportunities for management scientists. *Health Care Management Science*, 24(3), 455–459.
- Dai, T., & Yadav, P. (2021). Release second doses to speed vaccinations. *USA Today* (January 13) 7A.
- Duijzer, L., van Jaarsveld, W., Wallinga, J., & Dekker, R. (2018). Dose-optimal vaccine allocation over multiple populations. *Production and Operations Management*, 27(1), 143–159.
- Dvoretzky, A., Kiefer, J., & Wolfowitz, J. (1953). On the optimal character of the (s, S) policy in inventory theory. *Econometrica*, 21(4), 586–596.
- Gallien, J., Rashkova, I., Atun, R., & Yadav, P. (2017). National drug stock-out risks and the global fund disbursement process for procurement. *Production and Operations Management*, 26(6), 997–1014.
- Gottlieb, S. (2021). Pharmacies can get shots in arms. *Wall Street Journal* (January 3). <http://on.wsj.com/3aim4Be>
- Hall, S. N., Jacobson, S. H., & Sewell, E. C. (2008). An analysis of pediatric vaccine formulary selection problems. *Operational Research*, 56(6), 1348–1365.
- Harris, R. J., Hall, J. A., Zaidi, A., Andrews, N. J., Dunbar, J. K., & Dabrera, G. (2021). Effect of vaccination on household transmission of SARS-CoV-2 in England. *New England Journal of Medicine*, 385(8), 759–760.
- Haas, E. J., Angulo, F. J., McLaughlin, J. M., Anis, E., Singer, S. R., Khan, F., Brooks, N., Smaja, M., Mircus, G., Pan, K., Southern, J., Swerdlow, D. L.,

- Jodar, L., Levy, Y., & Alroy-Preis, S. (2021). Impact and effectiveness of mRNA BNT162b2 vaccine against SARS-CoV-2 infections and COVID-19 cases, hospitalisations, and deaths following a nationwide vaccination campaign in Israel: An observational study using national surveillance data. *Lancet*, 397(10287), 1819–1829.
- Hopkins, J. S. (2021). Developed countries lock up Covid-19 vaccines through 2023. *Wall Street Journal* (June 2). <https://on.wsj.com/3pL3qbb>
- Kaplan, E. H. (2020). OM Forum—COVID-19 scratch models to support local decisions. *Manufacturing & Service Operations Management*, 22(4), 645–655.
- Keskinocak, P., Oruc, B. E., Baxter, A., Asplund, J., & Serban, N. (2020). The impact of social distancing on COVID19 spread: State of Georgia case study. *PLoS One*, 15(10), e0239798.
- Kraiselburd, S., & Yadav, P. (2013). Supply chains and global health: An imperative for bringing operations management scholarship into action. *Production and Operations Management*, 22(2), 377–381.
- Longini, I. M., Ackerman, E., & Elveback, L. R. (1978). An optimization model for Influenza A epidemics. *Mathematical Biosciences*, 38(1–2), 141–157.
- Mak, H. Y., & Tang, C. S. (2021). More vaccine supplies are being released. It's a gamble. *Barron's* (Jan 12, 2021). <http://bit.ly/2ndvaxgamble>
- Mancini, D. P., & Gross, A. (2021). UK decision to space vaccination doses sparks fierce debate over risk. *Financial Times* (January 8). <http://on.ft.com/3pgBRVr>
- Miller, H., & Cohn, M. (2021). Maryland facing shortage of second COVID vaccine doses, health secretary says. *Baltimore Sun* (February 1). <http://bit.ly/md2dosehb>
- Muhumuza, R., & Mutsaka, F. (2021). In poorest countries, surges worsen shortages of vaccines. *Associated Press* (June 17). <http://bit.ly/africavaxap>
- Natarajan, K. V., & Swaminathan, J. M. (2017). Multi-treatment inventory allocation in humanitarian health settings under funding constraints. *Production and Operations Management*, 26(6), 1015–1034.
- O'Donnell, C., & Spalding, R. (2021). U.S. releases millions of COVID-19 doses, urges states to broaden vaccine access. *Reuters News* (Jan 12, 2021). <http://reut.rs/3rBPJv1>
- Paltiel, D., Schwartz, J., Zheng, A., & Walensky, R. P. (2021a). Clinical outcomes of a COVID-19 vaccine: Implementation over efficacy. *Health Affairs*, 40(1), 42–52.
- Paltiel, A. D., Zheng, A., & Schwartz, J. L. (2021b). Speed versus efficacy: Quantifying potential tradeoffs in COVID-19 vaccine deployment. *Annals of Internal Medicine*, 174(4), 568–570.
- Porteus, E. L. (1990). Stochastic inventory theory. In D. P. Heyman, & M. J. Sobel (Eds.), *Handbooks in operations research and management science* (pp. 605–652). Elsevier.
- Rawlinson, K. (2020). More data needed before giving just one vaccine dose, says Covid adviser. *The Guardian*. <https://bit.ly/3aEz6tj>
- Scarf, H. (1960). The optimality of (S, s) policies in the dynamic inventory problem. In K. J. Arrow, S. Karlin, & P. Suppes (Eds.), *Mathematical methods in the social sciences* (pp. 196–202). Stanford University Press.
- Shumsky, R. A., Smith, J., Hoen, A., & Gilbert, M. (2021). *Allocating COVID-19 vaccines: Save one for the second dose?* Dartmouth College Working Paper.
- Song, J. S., van Houtum, G. J., & Van Mieghem, J. A. (2020). Capacity and inventory management: Review, trends, and projections. *Manufacturing & Service Operations Management*, 22(1), 36–46.
- Strauss, M. (2021). Belgium halts J&J COVID vaccine for under 41s after first EU death. *Reuters* (May 26). <https://reut.rs/3gsfKJl>
- Tuite, A. R., Zhu, L., Fisman, D. N., & Salomon, J. A. (2021). Alternative dose allocation strategies to increase benefits from constrained COVID-19 vaccine supply. *Annals of Internal Medicine*, 174(4), 570–572.
- Voysey, M., Costa Clemens, S. A., Madhi, S. A., Weckx, L. Y., Folegatti, P. M., Aley, P. K., Angus, B., Baillie, V. L., Barnabas, S. L., Borat, Q. E., Bibi, S., Briner, C., Cicconi, P., Clutterbuck, E. A., Collins, A. M., Cutland, C. L., Darton, T. C., Dheda, K., Dold, C., ... Pollard, A. J. (2021). Single-dose administration and the influence of the timing of the booster dose on immunogenicity and efficacy of ChAdOx1 nCoV-19 (AZD1222) vaccine: A pooled analysis of four randomised trials. *Lancet*, 397(10277), 881–891.
- Wilde, H., Mellan, T., Hawryluk, I., Dennis, J. M., Denaxas, S., Pagel, C., Duncan, A., Bhatt, S., Flaxman, S., Mateen, B. A., & Vollmer, S. J. (2021). The association between mechanical ventilator availability and mortality risk in intensive care patients with COVID-19: A national retrospective cohort study. *BMJ Medicine*, 19(1), 1–12.
- Yamin, D., & Gavius, A. (2013). Incentives' effect in influenza vaccination policy. *Management Science*, 59(12), 2667–2686.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Mak, H.-Y., Dai, T., & Tang, C. S. (2022). Managing two-dose COVID-19 vaccine rollouts with limited supply: Operations strategies for distributing time-sensitive resources. *Production and Operations Management*, 1–19. <https://doi.org/10.1111/poms.13862>