Allocating COVID-19 Vaccines: Save One for the Second Dose?

Robert A. Shumsky[1], James E. Smith[1], Anne G. Hoen[2], and Michael Gilbert[3]

[1]Tuck School of Business at Dartmouth, Hanover, New Hampshire

[2]Department of Epidemiology, Geisel School of Medicine at Dartmouth, Hanover, New Hampshire

[3]Catholic Medical Center, Manchester, New Hampshire

August 10, 2021

Abstract

The majority of COVID-19 vaccines currently available are typically administered in two doses, with a prescribed number of weeks separating the two doses. Because of uncertainty in vaccine supply, many vaccination managers have saved doses in inventory to ensure on-time second doses. However, saving doses slows the administration of first doses and potentially delays completing the vaccination of the target population. In this paper, we use a mathematical model to explore the performance of policies to manage the administration of two-dose vaccines given supply uncertainty. The mathematical structure of the model suggests simple "set-aside" policies that prioritize second doses and set aside some vaccine for second doses that are due in the coming weeks; any remaining vaccine is given as first doses. We can determine good set-aside policies using a simple simulation model. In our numerical experiments, the performance of the recommended set-aside policy is close to a bound generated by the optimal allocation given perfect information about vaccine supply; thus the set-aside policies are nearly optimal given uncertainty in supply. We also find that a "lockbox" policy that ensures no second doses will be delayed appears to be overly conservative. We also show that set-aside policies perform well when vaccination supply is increasing and when two two-dose vaccines are available.

1. Introduction

The majority of COVID-19 vaccines administered worldwide require two doses, with a prescribed number of weeks between doses (CDC, 2021a). In the U.S., the CDC recommends that two doses of the Pfizer-BioNTech vaccine be administered three weeks apart (CDC, 2021b) and the Moderna vaccine doses be administered four weeks apart (CDC, 2021c). Uncertainty in vaccine supplies, however, can lead to temporary shortages and the risk of delays to second doses. Early in the U.S. vaccine rollout, public health officials had significant concerns about missed second doses (Choi and Renault, 2021; AP, 2021). On-time administration of second doses continues to be a problem as vaccines are being rolled out throughout the world, e.g., in Europe (Appuzzo et al., 2021), Brazil (Ionova, 2021), Chile (Reuters, 2021) and Nepal (Callamard, 2021).

Our interest in this problem was motivated by the experience of the fourth author (Gilbert), who is the Chief Medical Officer for the Catholic Medical Center (CMC) in Manchester, New Hampshire. The CMC was responsible for vaccinating approximately 3,000 people during Phase 1a of the vaccine rollout, when doses were administered to health workers and older adults in residential care settings. The CMC was given an initial supply of 365 doses of the Pfizer vaccine in December and 840 doses of the Moderna vaccine in January and was uncertain about future vaccine supply. Given that the vaccine supply is uncertain and on-time administration of the second dose is recommended, the CMC thought it may be good to hold some vaccine in inventory for second doses. However, holding back large amounts of vaccine slows the administration of first doses and may increase the time required to vaccinate the target population.

To get a sense of the magnitude of the variation in vaccine supplies, Figure 1(a) shows the deliveries to the CMC during Phase 1a of the vaccination process. Figure 1(b) shows vaccine deliveries to the state of Massachusetts during the early phases of the vaccination process (Mass Department of Public Health, 2021). Uncertainty in supply stems from many sources including issues in the manufacturing processes (Deo and Corbett, 2009), national and state policies changing over time, weather conditions (Harper, 2021), the availability of "low dead space" syringes to extract extra doses from vials of vaccine (O'Donnell, 2021), among many other things.

To study the tradeoff between ensuring on-time second doses and getting "shots in arms," we developed a model of a two-dose vaccination process which is managed by, for example, a hospital, a state, or a federal government. Each week an uncertain supply of vaccine is delivered and the vaccine manager must decide how many doses to allocate to first or second doses and how many to save for later use. The goal is to vaccinate the target population as quickly as possible while also minimizing delays in second doses. The model has a nice structure (specifically, $L^\natural$-convexity) and the optimal policies have a "set aside" form that hold inventory for second doses that are due in coming weeks. Simplified versions of these policies are easy to implement and can be tuned given different process parameters; the policies perform well in our numerical experiments.

In the U.S., the decision-making responsibility for allocating vaccines to first or second doses varied over time and across states. In the early stages, managers at medical centers (such as the CMC) and at the state

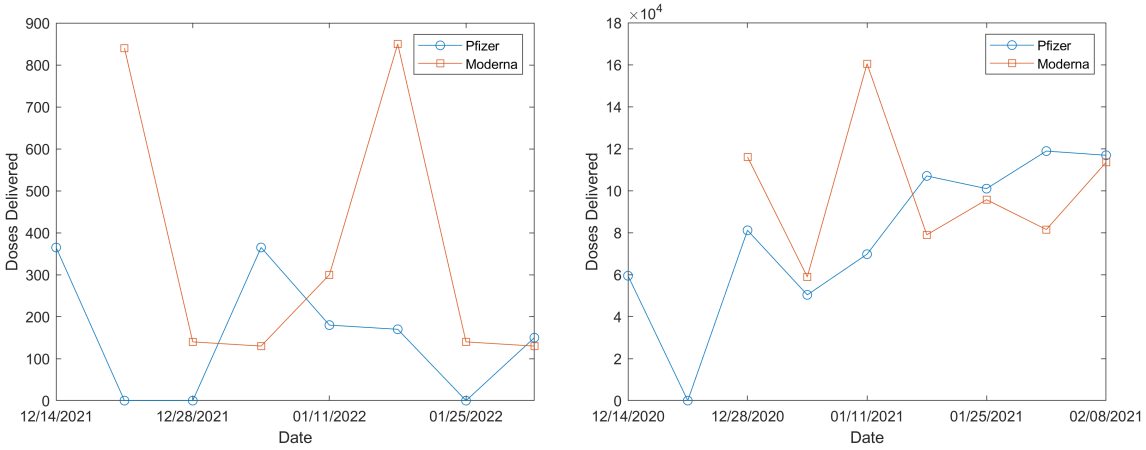

(a) Deliveries to the Catholic Medical Center

(b) Deliveries to Massachusetts

Figure 1: Vaccine deliveries during the early part of the COVID-19 vaccination process

level were concerned about missing second doses and frequently held inventory to ensure adequate future supply. In late January 2021, for example, the state of Maryland was holding 208,000 doses (Frost, 2021; Miller and Cohn, 2021). Meanwhile, the federal government was also storing inventory at the national level (Thomas, 2020; Edwards, 2021). According to Tuite et al. (2021), the initial U.S. policy was to hold one dose in inventory for every first dose given to be distributed later as a second dose; we will call this the *lockbox* policy. In early February 2021, the Biden administration assured states that second doses would always be available when needed and asked them not to hold inventory (Klein, 2021). Thus, though the locus of responsibility varied during the COVID-19 vaccine rollout, the question of whether to set aside some vaccine for second doses remained as long as the vaccination process was constrained by supply.

In §2, we review the related literature and, in §3, we describe the model and our proposed allocation policies. In §4, we describe numerical experiments that explore how these policies perform given various levels of supply uncertainty; we also test the model over a wide range of alternative scenarios. In §5, we offer a brief conclusion. To make the paper as accessible as possible to practitioners, the formal description and analysis of the model are presented in the appendix.

Though our discussion focuses on COVID-19 vaccinces, the challenge of managing multi-dose vaccines arises elsewhere. For example, unanticipated shortages of the shingles vaccine Shingrix led to concerns about second-dose delays when it was first introduced (Cimons, 2018). Of course, there may be other situations – e.g., vaccines for other viruses – where we will face the same issues in the future.

2. Literature Review

In our model, the goals are to minimize the average time to complete the two-dose vaccine regimen and to deliver second doses close to the recommended time; we show that this leads to an allocation policy that prioritizes second doses over first doses. While this is consistent with current practices in the U.S.,

there is an ongoing debate over this policy. Barnabas and Wald (2021), Bieniasz (2021), and Kadire et al. (2021) consider the public health trade-off between prioritizing inoculation of a greater proportion of the population with a single dose vs. fully vaccinating fewer individuals with the recommended two doses by intentionally extending the time between the first and second doses, as has been official policy in the UK. Paltiel et al. (2021) and Moghadas et al. (2021) formulate COVID-19 transmission and vaccination models to study whether first or second doses should be prioritized. The answer to this prioritization question depends on the degree and duration of partial protection against both severe disease and viral transmission that is afforded by the first dose and is complicated by the circulation of novel viral variants that may be suboptimally blocked by available vaccines. We do not model disease dynamics and focus on developing an inventory policy to optimize the completion of a recommended two-dose regimen. We also do not consider how to prioritize people within or among target populations; this is an important area of ongoing research that touches on ethics, epidemiology, and optimization (Enayati and Özaltın, 2020; Chen et al., 2020; Persad et al., 2020; Chen et al., 2021, Foy et al., 2021). Dai and Song (2021) describe how management science techniques may be applied to study a number of issues related to efficiency and equity of the vaccine rollout. Similarly, we do not consider the "demand" for the vaccine and issues related to vaccine hesitancy. In essence, given an ordered queue of people who are willing to be vaccinated, our problem is to work through the queue as fast as possible while minimizing delays in second doses given uncertainty in the vaccine supply.

Mak et al. (2021) and Tuite et al. (2021) combine deterministic models of vaccine supply with epidemiological models to explore the benefits and costs of holding inventory for second doses. These papers compare the lockbox policy with specific alternatives: Mak et al. (2021) examine the implications of not holding inventory at all and of stretching the time period between doses, while Tuite et al. (2021) varies the percentage of doses held in inventory according to a fixed schedule. In our analysis, we describe optimal inventory policies given uncertain (rather than deterministic) supply and suggest easy-to-implement near-optimal policies that are typically somewhere between holding no inventory and the lockbox policy.

Within the operations management literature, our vaccine inventory problem has some similarity to inventory problems with multiple demand classes and rationing (e.g., Deshpande et al., 2003; Koçağa and Şen, 2007): people requiring first and second doses can be seen as two customer classes. In the vaccine problem, however, the inventory model is complicated by the fact that individual second doses are linked to individual first doses. Our vaccine inventory problem is also similar in some ways to problems in patient scheduling (e.g., Green and Savin, 2008; Liu et al., 2010). However, our objective – rapid and effective vaccination of a population – is fundamentally different from the objectives of traditional patient scheduling models. Our model also bears some similarity to inventory models with lead times and/or perishable inventory and we analyze the model using $L^\natural$-convexity techniques as in this literature (see, e.g., Zipkin, 2008 and Chen et al., 2014). However, the structure of the vaccine model is different from those in the lost-sales or perishable goods literatures. We believe that our model represents a novel contribution to the operations management literature with interesting theoretical properties as well as immediate practical applications.

3. The Model and Proposed Policies

Our model considers vaccine supply and distribution on a weekly basis. Each week a random number of doses arrive and these doses, as well as any doses currently in inventory, may be allocated to second doses, allocated to first doses, or saved in inventory for later use. We assume that everyone who receives a first dose will return to receive a second dose and that nobody joins or leaves the target population during the vaccination process. For specificity, our discussion will focus on the Moderna vaccine whose doses are prescribed to be four weeks apart. Our simulations will initially assume that the weekly supplies are independent and identically distributed, but the model and framework can be used with a variety of supply forecast models. In §4.4.1, for example, we consider a case when the mean of the supply disitribution is increasing over time.

We formally describe the model in the Appendix, both as difference equations in Appendix §A and a stochastic dynamic program (DP) in Appendix §B. The former is useful for simulation studies and the latter is used to characterize optimal policies.

3.1 Objectives

We consider two performance metrics: (i) the average time until people receive their second dose (*average time to completion*) and (ii) the average delay in administering second doses beyond their due date, i.e., four weeks after the first dose (*average delay in second dose*). We combine these two metrics into a single objective function (*penalized average time to completion*) where each week of second-dose delay is multiplied by a penalty parameter $c \geq 0$:

$$\begin{aligned}(\text{penalized average time to completion}) = \\ (\text{average time to completion}) + c \times (\text{average delay in second dose})\end{aligned} \tag{1}$$

If $c = 0$, there is no penalty for delaying the second dose. The manager's objective is to minimize the expected value of this combined objective.

Another way to formulate the objective is to consider the "penalized average vaccination efficacy" in the target population over time. Suppose that an individual receives a benefit (i.e., in preventing infection and/or severe disease) of β_1 after one dose and an additional benefit of β_2 upon completing the second dose, with a penalty κ associated with each week of delay in the second dose. The population efficacy in a given period is then

$$\begin{aligned}(\text{population efficacy}) = \beta_1 \times (\text{number of people with first doses}) \\ + \beta_2 \times (\text{number of people with second doses}) \\ - \kappa \times (\text{number of people with delayed second doses})\end{aligned} \tag{2}$$

These population efficacies for each period are then averaged over the course of the vaccination process. This objective turns out to be equivalent to that of (1) if we normalize the coefficients in (2) so that $\beta_1 + \beta_2 = 1$ and take $c = (\kappa - \beta_1)$ (see §A.3). This alternative objective provides a way to interpret c as the net cost of delaying a second dose given as the difference between the cost of delay κ (reflecting psychological concerns or long-term loss in efficacy) less the benefit β_1 associated with partial vaccination.

Although the DP model is well defined whenever $c > -1$, in our analysis and experiments, we will focus on the case where $c \geq 0$ and show that it is optimal to prioritize second doses in this case. This priority is consistent with the policy for the Moderna and Pfizer vaccines in the U.S., though other countries (notably the U.K. and Canada) have prioritized first doses. Note, however, that a value $c < 0$ does not necessarily imply it is optimal to prioritize first doses; for such c, the optimal priority may depend on the details of the state of the system. For example, we have identified some examples where $c = -0.5$ still leads to prioritizing second doses and others where c just below zero leads to prioritizing first doses.

In practice, the appropriate penalty c may be difficult to choose with confidence or precision. In our analysis, we will consider a range of values for c and a variety of amounts to set aside. Given results for a range of policies, a decision-maker might consider the potential delays experienced with particular policies (e.g., by considering the results in Figure 4 below) and ask if they are comfortable with these tradeoffs: e.g., would it be acceptable if, say, people have 40% (or 10%) chance of experiencing a delay of one or two weeks in getting their second dose if it means vaccinating people one week earlier on average? Such concerns might guide their choice of amounts to set-aside (x in Figure 4) rather than optimizing for a given penalty c.

3.2 The Structure of Optimal Policies

Unfortunately, the DP model for this problem is difficult to solve to optimality because the state space is quite large. The cost-to-go function and optimal policy will depend on: the number of doses of vaccine supplied in the given week plus the number of vaccines in inventory; the number of people currently due (or overdue) for their second dose; the number of people who received first doses in each of the previous three weeks (who will be due for their second dose in the weeks ahead); and the number of people in the target population who have not yet received their first dose.

Though the DP is difficult to solve to optimality, the optimal policies have a nice structure that we can use to guide the choice of simple policies that are easy to evaluate and implement. In particular, we show in the appendix that when the penalty $c \geq 0$, an optimal policy will:

(i) First, administer vaccines to those who are due or overdue for their second dose, up to the available inventory (after receiving the week's supply).

(ii) Then, administer some or all of the remaining inventory as first doses to those who need it. The number of first doses given is decreasing in the number of doses due in the coming weeks and decreasingly so: that is, if we increase the number of second doses due in week j ($0 \leq j \leq 3$) by, say, 100, it is optimal

to decrease the number of first doses administered by some number n, where n is less than 100 and n is decreasing in j.

These policies can be interpreted as prioritizing second doses and setting aside some of the remaining inventory (if any) for future second doses. The optimal setaside varies with the number of doses due and the setasides are greater for doses due sooner. Moreover, the optimal setasides and number of first doses given are both (weakly) increasing in the amount of inventory on hand.

Given the complexity of finding truly optimal policies, we focus on simplified versions of these policies that we describe next. In §4, we evaluate these policies in numerical experiments and compare the performance of these policies with a theoretical performance bound provided by a policy that allocates vaccines optimally given perfect information about vaccine supply; we describe this perfect information policy in §3.4.

3.3 Set-aside Policies

The simplified set-aside policies we study mimic the features of an optimal policy for the DP just discussed. In these policies, we first administer second doses to people who are currently due or overdue, prioritizing those who are most overdue. If doses remain after administering second doses, we calculate the number of first doses administered after setting aside doses for future use. The amounts set aside are determined from set-aside fractions w_1, w_2, w_3, w_4 where w_i indicates the fraction of second doses coming due in i weeks that are reserved for future use. For example, if 60, 40, and 50 doses are coming due in the next three weeks and $(w_1, w_2, w_3, w_4) = (1, 1, 0.5, 0)$, then we reserve $1 \times 60 + 1 \times 40 + 0.5 \times 50 = 125$ doses if they are available and use the remainder, if any, as first doses.

The number of doses reserved when $w_4 > 0$ is somewhat more involved because the number of doses due in four weeks depends on the number of first doses given in the current week, which itself depends on w_4. If $w_4 = 1$, then half of all leftover doses (after accounting for set-asides for earlier weeks) are administered as first doses and the rest are reserved. More generally, if D_t is the number of doses available after administering second doses, then under the set-aside policy the number of first doses administered in period t is,

$$\frac{1}{1+w_4}\left(D_t - w_1 v_1 - w_2 v_2 - w_3 v_3\right)^+ \tag{3}$$

where v_i denotes the number of second doses due in i weeks and $(x)^+ = \max(x, 0)$.[1] If there are fewer than this many people who need their first dose, we give first doses to all those who have not yet received their first dose. Any vaccines remaining after these first and second doses are administered are held in inventory for future use. Following the structure of the optimal policy, we assume the set-aside fractions are between 0% and 100% and decreasing, $w_1 \geq w_2 \geq w_3 \geq w_4$, as it is more important to set aside vaccine for second doses that are due sooner.

[1] If f is the fraction of the doses available to be administered as first doses, we want $(1 - f) = w_4 f$ (i.e., doses remaining = $w_4 \times$doses given) which leads to $f = 1/(1 + w_4)$.

When $(w_1, w_2, w_3, w_4) = (1, 1, 1, 1)$, a second dose is held in inventory for every first dose given and no second doses will be delayed. This is equivalent to dividing every vaccine delivery in half, with one half administered as first doses and the other half "locked away" for second doses. As mentioned earlier, we call this special case of the set-aside policy the *lockbox policy*.

Note that the optimal policies need not correspond to constant set-aside fractions. For example, if 10 second doses are due next week, it may be optimal to set aside nothing (0%) but, if 100 are due next week, it may be optimal to set aside 80 doses (80%) for future use. Moreover, the amounts set aside (0 and 80 in this example) in an optimal policy may vary depending on how many doses are due in two or three weeks. The optimal policies may thus be quite complex, not only to calculate but also to implement.

To further simplify our discussion and analysis, we will focus on policies that correspond to a given number of weeks (x) of doses to set aside: $x = 0$ corresponds to set-aside fractions $(0, 0, 0, 0)$, $x = 1$ to $(1, 0, 0, 0)$, $x = 2$ to $(1, 1, 0, 0)$, and so on up to $x = 4$ which is the lockbox policy $(1, 1, 1, 1)$. We will also consider fractional policies where $x = 0.5$ corresponds to $(0.5, 0, 0, 0)$, $x = 1.5$ to $(1, 0.5, 0, 0)$, and so on. When choosing set-aside policies, we will restrict our search to the one-dimensional class of policies of this form. Specifically, we will consider half-step increments in the number of weeks set-aside, taking $x = 0, 0.5, 1, \ldots, 4$.[2]

3.4 Perfect Information Policies

To establish a lower bound on the performance of an optimal policy, we will consider a model that optimizes the allocation of vaccines given perfect information about the supply of vaccines, i.e., an "information relaxation" performance bound (Brown et al., 2010). To determine this performance bound, for each simulated scenario of week-by-week vaccine supplies, we choose optimal dose allocations for all periods under the assumption that all supplies are known with certainty. Given the information about the future doses, this inner optimization problem is deterministic and can be formulated as a linear program that is described in Appendix A.2. These perfect information policies cannot be implemented in practice (because the vaccine supply is actually uncertain) but their performance provides a bound on the performance of any feasible policy given uncertainty about the vaccine supply: the allocations chosen by the inner optimization problem could match that of an optimal policy given uncertain supplies or could deviate to improve performance. This perfect information bound also provides a bound on the performance that would be obtainable with better information about future supplies. These performance bounds likely could be improved by incorporating a penalty as discussed in Brown et al. (2010), but we will consider only the simple perfect information bound.

[2] In our experiments, we also considered several other policies including a "hold up to" policy that is based on the "order up to" policies that are common in inventory settings. In the hold-up-to policy, we first administer second doses to people who are overdue or due within the current week. Then, after administering second doses, we reserve up to a fixed number of remaining doses $\bar{q}$ for future use. Any remaining doses (i.e., above $\bar{q}$) are used for first doses. The target inventory level $\bar{q}$ was optimized in our experiments. These hold-up-to policies (and others we considered) did not perform as well as the set-aside policies and will not be discussed here.

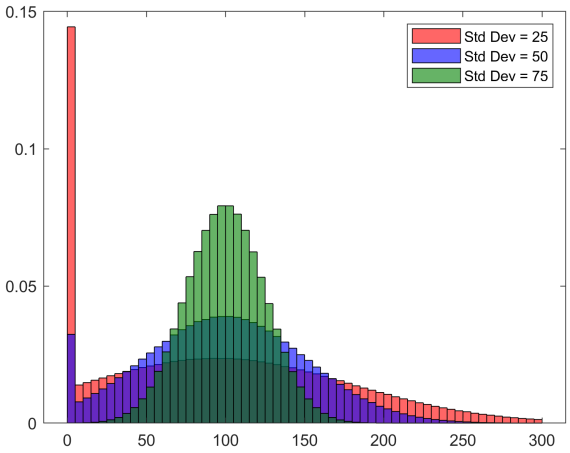

Figure 2: Supply distributions for numerical experiments

4. Numerical Experiments

In our numerical experiments, we will focus on a standardized population of 1,000 people to be vaccinated with the Moderna vaccine with four weeks recommended between doses. We will initially assume that vaccine supplies are independent from week to week and consider three different distributions for the supply, all rectified normal distributions meaning any negative observations are moved to zero (rectified). The distributions all have mean equal to 100 (after rectification) and standard deviations of 25, 50, and 75 (again after rectification), representing low, medium, and high uncertainty cases; see Figure 2. These distributions have a probabilities of zero supply equal to 0.00003, 0.0259, and 0.1313, respectively. As a point of comparison, the Pfizer and Moderna deliveries to the CMC shown in Figure 1(a) had coefficients of variation (CVs) of 0.99 and 0.93 (respectively) and the deliveries to Massachusetts in Figure 1(b) had CVs of 0.49 and 0.33. We evaluate the proposed policies using Monte Carlo simulations with 100,000 trials.

In this section, we first study the key tradeoff between reducing the average completion time and the average delay in second doses using set-aside policies. We then consider optimal set-aside policies for a range of penalties for delayed second doses and compare the performance of these policies to the perfect information bound. Next, we consider the uncertainty in these performance measures, reflecting the uncertainty in the vaccine supply. Finally, we study how the results change with alternative assumptions.

4.1 The Key Tradeoff

Figure 3 shows how the expected average time to completion changes as more vaccine doses are held back under the set-aside policy. Here we see that the lockbox policy (with setaside x=4) is quite conservative, delaying the average completion time by more than a week compared to the set-aside policies that reserve two or fewer weeks for future use. In Section 4.2 we will discuss the perhaps surprising fact that the expected average completion time, given high uncertainty, is longer when we hold no inventory ($x = 0$) than when we hold some inventory ($x = 0.5$). We also see the value of reducing the standard deviation of vaccine supplies:

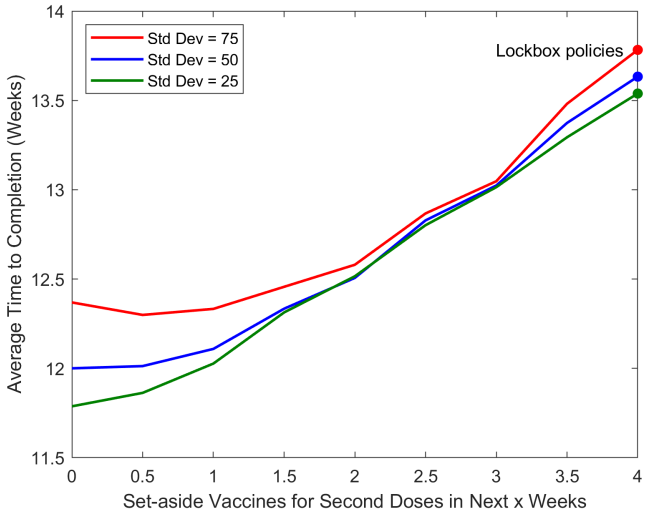

Figure 3: Average time to completion under set-aside policies

in the low uncertainty scenario, the average completion time is substantially less when no inventory is set aside. Note that the values shown in Figure 3 are the estimated expected average completion times based on the Monte Carlo simulation; we will discuss the uncertainty in these average completion times in §4.3 below. (The mean standard errors associated with these estimates due to simulation error are all less than 0.0065 weeks.)

Figure 4 shows the cost of holding little inventory: potential delays for second doses. While the lockbox policies have no second-dose delays, if we set aside no inventory for future doses (i.e., take $x = 0$) second doses are expected to be delayed by approximately 0.9 weeks on average in the high uncertainty case and approximately 0.2 weeks in the low uncertainty case; see Figure 4a. Focusing on the high uncertainty case, in Figure 4b we see that when no doses are set aside ($x = 0$) approximately 45% of people experience some delay in getting their second dose: 23% experience a delay of one week; 12% experience a delay of two weeks; 6% experience a delay of 3 weeks; and 5% experience a delay of 4 weeks or more. Recall that the probability of receiving zero doses in a given week is 13% in the high uncertainty case. When holding zero inventory, weeks with zero supply will result in the delay of second doses if second doses are due that week. Following a week with zero supply, the vaccination process can get further behind if later vaccine supplies are insufficient to cover those who are overdue for their second dose, let alone those who are currently due. Increasing the amount of vaccine set-aside (x) reduces the frequency and severity of these delays.

4.2 Minimizing the Penalized Average Time to Completion

We now consider the set-aside policies that minimize the penalized average completion time (1). Figure 5 shows the penalized average completion time as the set-aside parameter x varies, with penalty $c = 1$. Here we see that in the high uncertainty case, the penalized average time to completion is minimized with 1.5 weeks set aside; in the middle and low uncertainty cases, 0.5 weeks set aside is optimal. Figure 6 shows how the

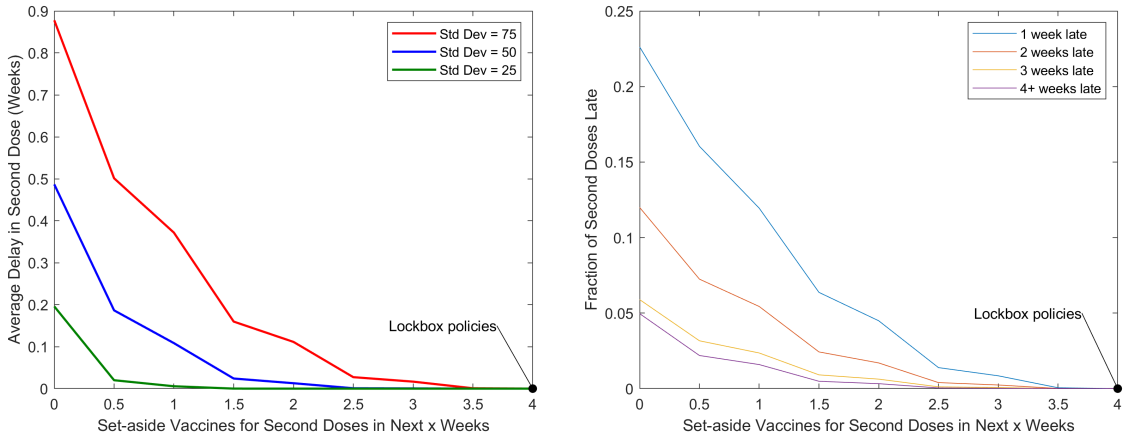

(a) Average delay in second doses

(b) Duration of delays in high uncertainty case

Figure 4: Delays in second dose for set-aside policies

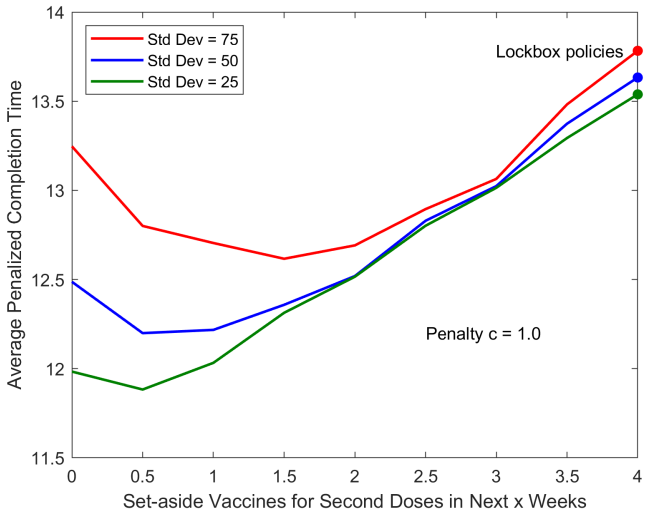

Figure 5: Average penalized completion time for varying set-aside quantities with $c = 1$

penalized time to completion changes as c varies from 0 to 3 for the three uncertainty cases; the corresponding optimal set-asides are shown in the figure. As c increases and we place more weight on delaying second doses, it is optimal to set-aside more vaccine. We also see that the lockbox policies are too conservative, even with penalties as high as $c = 3$. Moreover, the recommended set-asides appear to be robust: for example, in the high uncertainty case, set-aside policies of 1.5 to 2 weeks are optimal or near-optimal over a wide range of penalties c.

The fact that it is optimal to set aside some doses when $c = 0$ in the high uncertainty case is perhaps surprising. In the $c = 0$ case, the combined objective reduces to the average completion time shown in Figure 3. One might expect that if there is no penalty for delaying second doses, it would be best to hold no inventory and administer first doses as fast as possible. The fact that it is optimal to set aside some doses in this scenario is due to a smoothing effect of the set-aside policy: if all available doses are used for

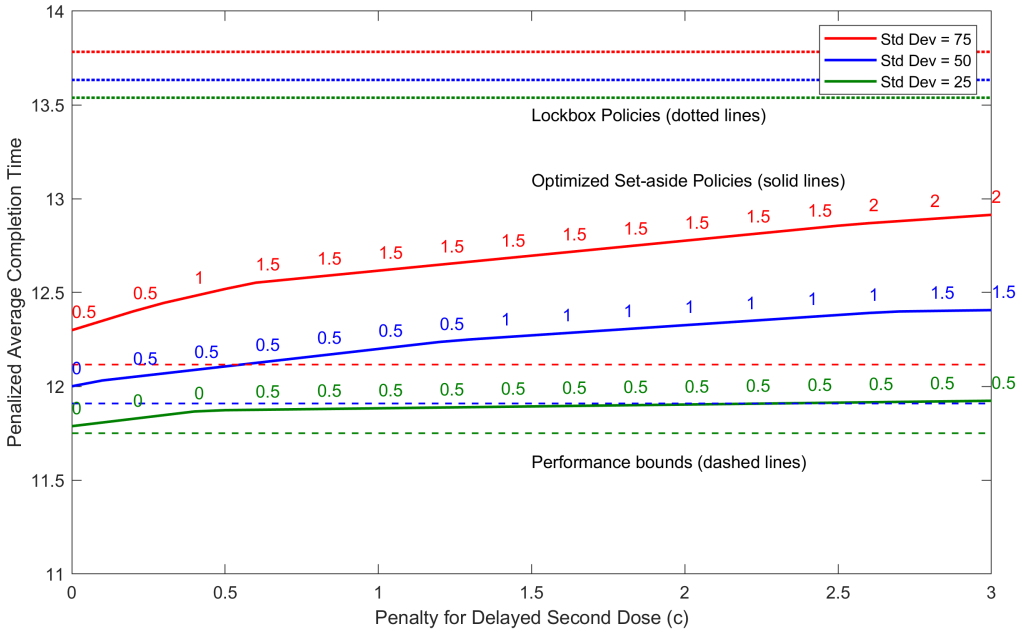

Figure 6: Penalized average completion times for optimized set-aside policies

first vaccinations and there is a week or more with zero or small supplies (a "supply gap"), second doses are delayed during the gap. Having some doses in inventory before such a gap reduces the average completion time because these doses can be used to administer second doses on time. Using these stored doses for first doses instead would not reduce the average completion time because their corresponding second doses would be delayed anyway. Such supply gaps are much rarer in the medium and low uncertainty cases and it is not optimal to set aside any inventory if the penalty for delayed second doses is small ($c = 0$ or 0.5).

Figure 6 also shows the perfect information lower bounds which, as discussed in §3.4, maintain supply variability but eliminate supply uncertainty. Here we see that eliminating uncertainty can be valuable though the set-aside policy is close to the lower bound when c is small: with $c = 0$, the optimal set-aside policy is within 0.3%, 0.8%, and 1.4% of the perfect information lower bound in the low, medium and high uncertainty cases (respectively). With $c = 1$, the optimal set-aside policy is within 1.1%, 2.4%, and 3.9% in these same cases. These bounds show the extent to which using a truly optimal policy (rather than a set-aside policy of the form considered here) and/or improving supply forecasts can help improve performance. These results suggest that there is relatively little to be gained by considering more sophisticated allocation policies when there is little uncertainty or relatively small penalties for delaying second doses.

The structure of the optimal policy with perfect information is also interesting: if there is a positive cost of second-dose delay ($c > 0$), the optimal allocations given perfect information about supplies ensure that second doses are never delayed and the perfect information bound is insensitive to c. To see why, suppose there is somebody whose second dose is delayed in period t, for the first time. We can delay this person's first dose (four periods earlier) by one period. Note that this is feasible: delaying this person's dose increases the end-of-period inventory in period $t-4$ by one, thereby creating an additional unit of inventory to provide

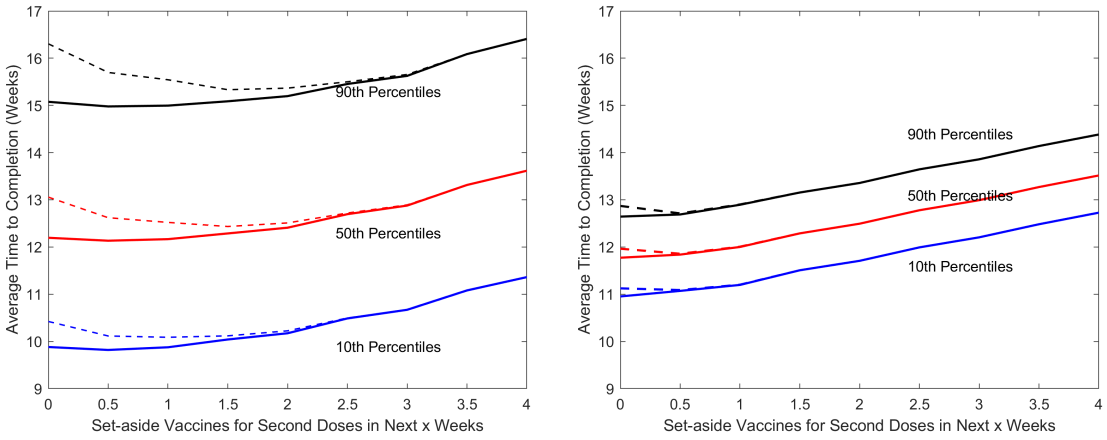

(a) High uncertainty case doses

(b) Low uncertainty case

Figure 7: Percentiles for average completion times

for that person's first dose one period later. This shift of timing does not affect the inventory or supply of doses in any later periods and thus reduces the penalty for the delayed second dose without changing the average time to completion. Such a shift of first-dose timing can be repeated for everybody with a delayed second dose and leads to a solution with the same average time to completion and no delayed second doses. This argument implies that it is not the *variation* in the vaccine supply that leads to delayed second doses, it is the *uncertainty* about the supply. Incorporating a "penalty" in the information relaxation bound could "punish the decision maker" for using information in this way and lead to tighter bounds (see, e.g., Brown et al., 2010).

4.3 Uncertainty in Average Completion Times

It is important to note that the results in Figures 3-6 focus on *expected* average completion times and expected delays, averaging across possible supply scenarios. Figure 7 describes the uncertainty in the average completion times by showing the 10th, 50th, and 90th percentiles of the completion time distributions for varying numbers of weeks set aside; the average completion times are shown with solid lines and the penalized average completion times (with $c = 1$) are shown with dashed lines.

In Figure 7, we see that uncertainty in the supply creates considerable uncertainty in the average completion time. In the high uncertainty case, the 10th and 90th percentiles range from 10 to 16 weeks, depending on the policies and objective. The ranges for the low uncertainty case are naturally much narrower. The supply uncertainty affects the average completion times for different set-asides similarly as the widths of the 10-90 ranges are approximately equal for all set-asides. However, the penalized average completion times are somewhat more uncertain, particularly for lower set-asides in the high uncertainty case. The wider uncertainty ranges reflect the additional uncertainty in delays in second doses due to the supply uncertainty.

4.4 Alternative Assumptions

In the numerical experiments above, the population size is 1,000 people and the supply distributions have a rectified normal distributions with means 100 and standard deviations of 25, 50 and 75. The results and recommendations of our analysis will be unchanged if we scale the size of the target population and the mean and standard deviation of the supply in proportion. In practice, the vaccine allocations tend to be scaled with the size of the target population, so our results have relevance over a broad range of population sizes. Moreover, the recommendations appear to be not very sensitive to changes in the population size. For example, if we double the population size to 2,000 without changing the supply distributions, the average completion times roughly double (as one might expect) but the recommended number of weeks to set aside shown in Figure 6 are unchanged.

4.4.1 Supply Growth

We also ran simulations that assume the vaccine supply is increasing over time. Specifically, we considered the medium uncertainty scenario with mean 100 doses per week and standard deviation 50 and two growth scenarios, with the (before rectification) mean growing at a rate of 5 or 10 doses per week. The results are summarized in Figure 8. Here we see that the penalized average completion times decrease with increasing supply, as one would expect. The optimal setasides (shown in text in the figure) are decreasing with growth in supply because one can use growth in supply to help ensure the availability of future doses. Although we could experiment with setasides that are changing over time in response to changes in the supply distribution, the simple set-aside policies with growth are very close to the perfect information bounds indicating that there is little to be gained by considering more sophisticated policies.

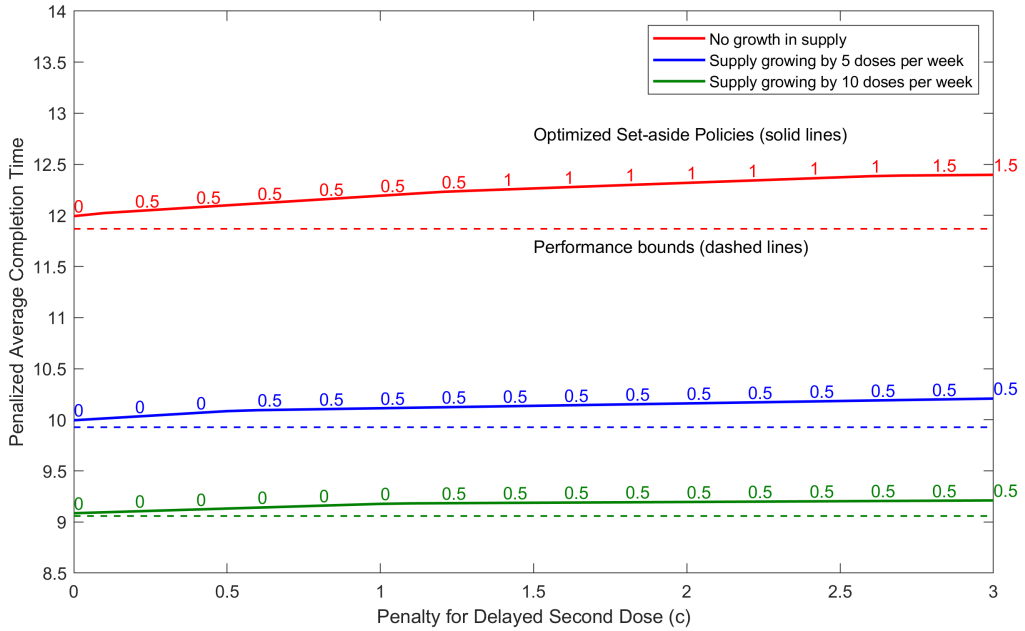

Figure 8: Penalized average completion times for optimized set-aside policies with growing supplies

4.4.2 Two Vaccines

We next consider a scenario where there are two two-dose vaccines available, one with 4 weeks between doses – like the Moderna vaccine – and the other with 3 weeks between doses – like the Pfizer vaccine – both with uncertain supplies. We assume, as per CDC guidance, that the second dose must be the same vaccine as the first dose. Though we will not describe or analyze a formal model for this case, we would expect the optimal policies to have a similar structure as in the single vaccine case. Assuming a penalty $c \geq 0$, for each vaccine, we would prioritize second doses that are due, set aside some for future vaccinations that are due, and use any remaining vaccine for first doses; those receiving first doses are drawn from a common pool of unvaccinated people. Intuitively, the motivation for setting aside doses for second doses due in the future is related to flows for a single vaccine. For most of the vaccination process, we would expect the optimal setasides for one vaccine to be relatively insensitive to the number of doses due for the other vaccine. Only when nearly the entire population has been vaccinated might the vaccine manager prefer to administer the vaccine with the shorter time between doses or more doses in inventory in order to complete the vaccination process sooner. Therefore, we would expect reasonably good performance from using fixed set-aside quantities for each vaccine.

To test these policies, we ran a simulation where the Moderna and Pfizer vaccine supplies are allocated according to set-aside policies, but to a shared target population. We consider high, medium and low uncertainty supply distributions as described above, assuming both vaccines have the same distribution and the draws are independent. We double the population to 2000 people to make the two-vaccine results comparable to the earlier one-vaccine results, e.g., those in Figure 6. The results are summarized in Figure 9. Here, for each supply distribution scenario, we found the optimal number of weeks to set aside for each vaccine by jointly optimizing in the simulation; these optimal setasides (Pfizer/Moderna) are shown in the figure. (There is little lost if we optimize independently.) Also shown are perfect information performance bounds for the two-vaccine scenarios.

Qualitatively the results for the two-vaccine case in Figure 9 are quite similar to those for the one-vaccine case in Figure 6, though the (penalized) average completion times are somewhat shorter in the two-vaccine case because the time between doses for the Pfizer vaccine is less than that for the Moderna vaccine that was considered in the one-vaccine case. As in the one-vaccine results, we see that when applying set-aside policies with two vaccines, the optimal policies appear to be robust and nearly optimal (particularly with small c or low uncertainty) whereas lockbox policies are clearly suboptimal.

4.4.3 Incorporating Second Dose "No Shows"

Finally, we also considered a simulation model that incorporated "no shows" for second doses; according to news reports approximately 10% of people in the U.S. who receive a first Covid-19 vaccine dose fail to show up for a second dose (McPhillips, 2021). In these simulations, we found that the set-aside policies perform well in this setting and the qualitative conclusions remain unchanged. We also found that it is effective to

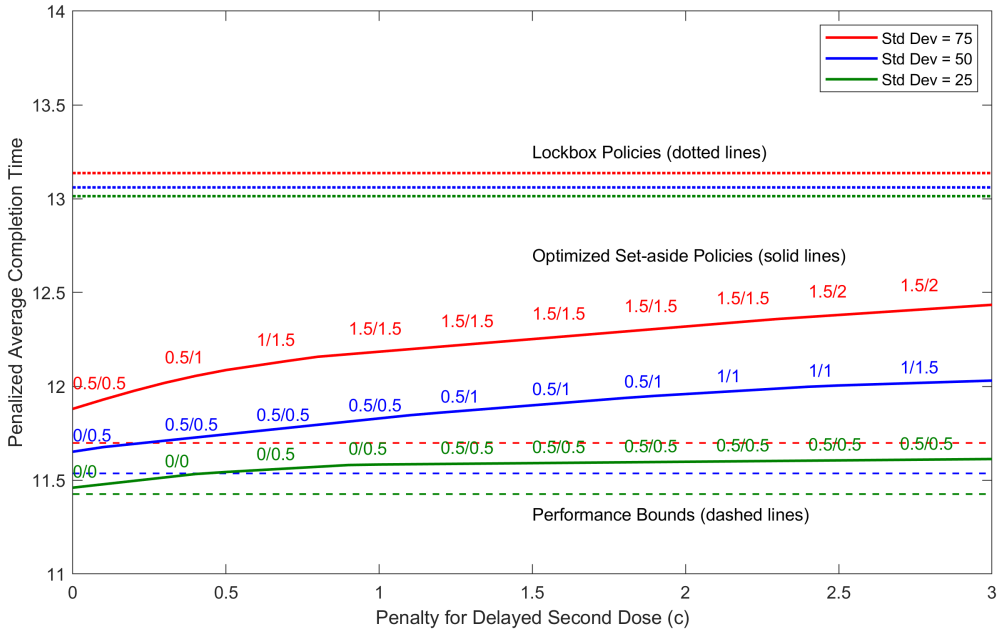

Figure 9: Penalized average completion times for optimized set-aside policies with two vaccines

simply scale the set-aside quantities to take no-shows into account. For example, if 90% of those receiving a first dose return for their second dose, it is reasonable to set aside 90% of the quantity one would if we were to assume that 100% returned.

5. Conclusion

Our model of the two-dose vaccine rollout demonstrates how the optimal number of doses to hold in inventory depends on the cost of delayed second doses and the uncertainty in vaccine supplies, as well as the forecasts for future supplies. The formal analysis of the model shows that the optimal number of doses to set aside increases with the number of second doses due in the coming weeks, with second doses due sooner having a larger impacts on the optimal setaside than doses due later. We also prove that both the optimal setasides and number of first doses to deliver within a period increase with the quantity of inventory available. We use these insights to construct simple set-aside policies that can be calibrated using a simulation model. These heuristic set-aside policies are very simple, requiring the manager to set aside vaccine for second doses due in the coming weeks, without considering complex supply forecasts. Numerical experiments demonstrate that these simple policies are near-optimal for a wide range of supply uncertainties, delay costs, and supply forecasts. A lockbox policy that ensures no second doses will be delayed - a policy that was widely implemented for the U.S. vaccine rollout - appears to be overly conservative.

While in the U.S. the primary limitation on COVID-19 vaccination rollout is now the lack of demand for vaccines (due in part to vaccine hesitancy), vaccine supplies remain constrained and uncertain in many other countries. The two-dose vaccine model developed here thus continues to be relevant for the current

pandemic and may be relevant again in the future as we face other novel viruses or other diseases with new multi-dose treatments and limited, unreliable supplies.

A. Appendix: Model Details

A.1 Model and Policy Definitions

Variable Definitions: To formally describe the model, we define the following variables:

P = number of people in the target population to be vaccinated

L = recommended number of weeks between first and second doses

T = model horizon (weeks)

t = week index ($0 \leq t \leq T$)

$\tilde{a}_t$ = random number of vaccine doses arriving in week t

I_t = inventory of available doses available in week t, after vaccine arrival

U_{it} = number of people who have not received dose $i, i = 1, 2$ at start of week t

d_t = number of second doses due at start of week t. This includes those who received first doses L periods before, as well as overdue second doses.

v_{it} = number of doses of dose i ($i = 1, 2$) administered at time t

All variables are nonnegative and we can take them to be either real numbers or restricted to integers; we will assume they are real valued in our analysis and our numerical examples. We can assume any nonnegative probability distribution for the weekly supplies $\tilde{a}_t$ and these distribution can vary from week to week and be dependent.

Difference Equations: To describe the dynamics of the model for use in simulation (and elsewhere), we define difference equations for $t \geq 1$. First define $I_0 = v_{2,0} = 0$, $U_{1,0} = U_{2,0} = P$, and $v_{1,t} = 0$ for $t \leq 0$. For $1 \leq t \leq T$:

$$
\begin{aligned}
I_t &= I_{t-1} + \tilde{a}_t - v_{1,t-1} - v_{2,t-1} \\
U_{1t} &= U_{1,t-1} - v_{1,t-1} \\
U_{2t} &= U_{2,t-1} - v_{2,t-1} \\
d_t &= \begin{cases} 0 & \text{for } t = 0, \ldots, L-1 \\ d_{t-1} + v_{1,t-L} - v_{2,t-1} & \text{otherwise} \end{cases}
\end{aligned}
$$

Objective: The objective is to minimize a weighted sum of two quantities, (i) the number of weeks spent by

people without both doses (*the time to completion*), and (ii) the number of weeks of delayed second doses. The second-dose delay is weighted by penalty $c > -1$:

$$\sum_{t=1}^{T} U_{2t} + c \sum_{t=1}^{T} (d_t - v_{2t}) \ .$$

In the body of the paper (e.g., in the numerical experiments), we normalize this objective by dividing by the target population size P.

Set-aside Policy: Given set-aside fractions $(w_1, \ldots, w_L)$, the second and first doses are given by:

$$\begin{aligned} v_{2t} &= \min\left(d_t, I_t\right) \\ v_{1t} &= \min\left(U_{1t},\ \frac{1}{1+w_L}\left(I_t - v_{2t} - \sum\nolimits_{j=1}^{L-1} w_j v_{1,t-(L-j)}\right)^+\right). \end{aligned}$$

A.2 Optimization with Perfect Information

Given perfect information about vaccine supplies, the vaccine allocation problem can be formulated as a linear program (LP) with decision variables v_{1t}, v_{2t} $(0 \le t \le T)$ representing the number of first and second doses in each week. The objective is to minimize the *penalized time to completion* given as

$$\begin{aligned} &\min\Big\{(\text{time to completion}) + c \times (\text{delay in second dose})\Big\} \\ &\quad = \min\left\{\sum_{t=0}^{T-1} t v_{2t} + c \sum_{t=0}^{T-1} d_t\right\} \end{aligned} \tag{A.1}$$

where d_t is the number of delayed dosed in period t. We can write the objective directly in terms of the doses given. First, note that the delayed doses in period t are equal to the cumulative number of first doses given up to L periods previously, less the cumulative number of second doses given

$$d_t = \sum_{\tau=0}^{t-L} v_{1,t} - \sum_{\tau=0}^{t} v_{2t} \ . \tag{A.2}$$

Then the objective (A.1) can be rewritten as

$$\min\left\{\sum_{t=0}^{T-1} t v_{2t} + c \sum_{t=0}^{T-1} \left(\sum_{\tau=0}^{t-L} v_{1,t} - \sum_{\tau=0}^{t} v_{2t}\right)\right\} \tag{A.3}$$

Let $\tilde{A}_t = \sum_{\tau=0}^{t-1} \tilde{a}_\tau$ denote the cumulative supply up to period t. The constraints for the LP can then be written

$$\begin{aligned}
&\sum_{\tau=0}^{t} v_{1\tau} + \sum_{\tau=0}^{t} v_{2\tau} \le \tilde{A}_t && \text{for } t = 0, \ldots, T-1 \\
&\sum_{\tau=0}^{t} v_{2t} - \sum_{\tau=0}^{t-L} v_{1,t} \le 0 && \text{for } t = 0, \ldots, T-1 \\
&\sum_{t=0}^{T-1} v_{2,t} = P \\
&v_{1t}, v_{2t} \ge 0 \ .
\end{aligned}$$

The first constraint requires cumulative doses administered to be less than or equal to cumulative supply. The second requires cumulative second doses given to be less than or equal to the cumulative number of first doses given up to L periods earlier or, equivalently using (A.2), $d_t \ge 0$. The third constraint requires everyone to receive a second dose.

A.3 Equivalence of Objectives

Here we show that the original objective (A.1) is equivalent to the alternative objective discussed in §3.1. This alternative objective is to maximize the *average population efficacy* over the course of the vaccination period. Let $V_{it} = \sum_{\tau=0}^{t} v_{i\tau}$ $(i = 1, 2)$ denote the cumulative number of first and second doses given by period t. Working with the total over time rather than the average, this alternative objective can be written as

$$\begin{aligned}
&\max \Bigg\{ \sum_{t=0}^{T-1} \Big(\beta_1 \times (\text{number of people with first doses in period } t) \\
&\qquad + \beta_2 \times (\text{number of people with second doses in period } t) \\
&\qquad - \kappa \times (\text{number of people with delayed second doses in period } t) \Big) \Bigg\} \\
&= \max \left\{ \beta_1 \sum_{t=0}^{T-1} V_{1t} + \beta_2 \sum_{t=0}^{T-1} V_{2t} - \kappa \sum_{t=0}^{T-1} d_t \right\} \\
&= \beta_1 PL + \max \left\{ \beta_1 \sum_{t=0}^{T-1} V_{1,t-L} + \beta_2 \sum_{t=0}^{T-1} V_{2t} - \kappa \sum_{t=0}^{T-1} d_t \right\} \\
&= \beta_1 PL + \max \left\{ (\beta_1 + \beta_2) \sum_{t=0}^{T-1} V_{2t} + \beta_1 \sum_{t=0}^{T-1} (V_{1,t-L} - V_{2t}) - \kappa \sum_{t=0}^{T-1} d_t \right\} \\
&= \beta_1 PL + PT - \min \left\{ (\beta_1 + \beta_2) \sum_{t=0}^{T-1} t v_{2t} + (\kappa - \beta_1) \sum_{t=0}^{T-1} d_t \right\}. \qquad \text{(A.4)}
\end{aligned}$$

The second equality follows from noting that $V_{1t} = 0$ for $t < 0$ (by definition) and $V_{1t} = P$ for $t > T - L$ since everyone is fully vaccinated by T and therefore everyone must have received the first dose by $T - L$.

The third and fourth equalities follow from rearranging and using the identity (A.2) (which can be written in terms of cumulative doses as $d_t = V_{1,t-L} - V_{2t}$) and observing

$$\sum_{t=0}^{T-1} V_{2t} = PT - \sum_{t=0}^{T-1} tv_{2t} \ ,$$

which applies "summation by parts," using the fact that everyone has been vaccinated by period T-1. Comparing (A.4) to (A.1), we see that the alternative objective (A.4) differs by a constant from the original objective (A.1) if we take $\beta_1 + \beta_2 = 1$ and $c = \kappa - \beta_1$. Thus the alternative objective is equivalent to the original objective (i.e., the two objectives paired with the same constraints yield the same optimal solutions) if we normalize β_1 and β_2 and define c this way.

B. Stochastic Dynamic Programming Formulation

Because the stochastic dynamic program (DP) is not dependent on time, we use somewhat different notation here than in the previous section:

ℓ : weeks between doses less one (i.e., $L-1$ which is 3 for the Moderna vaccine)

d_0 : second doses currently due or overdue (beginning of week)

d_i : second doses due in i weeks, $i = 1, \ldots, \ell$ (beginning of week)

U_1 : number who have not yet received any doses of vaccine (beginning of week)

I : inventory of vaccine on hand (beginning of week)

$\tilde{a}$: uncertain doses arriving at the end of week

v_1 : first doses given in current week

v_2 : second doses given in current week

The cost-to-go function $\hat{f}$ can be written recursively as

$$\begin{aligned} \hat{f}(d_0, \ldots, d_\ell, U_1, I) = \min_{(v_1, v_2) \in \mathbb{V}} \{ & (d_0 + \ldots + d_\ell + U_1 - v_2) + c(d_0 - v_2) \\ & + \mathbb{E}\left[\hat{f}(d_0 - v_2 + d_1, \ldots, d_\ell, v_1, U_1 - v_1, I - v_1 - v_2 + \tilde{a})\right] \Big\} \end{aligned} \tag{B.1}$$

where

$$\mathbb{V} = \{(v_1, v_2) : v_1, v_2 \geq 0, v_2 \leq d_0, v_1 \leq U_1, v_1 + v_2 \leq I\} \ .$$

The terminal case is when there is nobody left to vaccinate $\hat{f}(0, \ldots, 0, 0, I) = 0$. As long as there is a positive probability of receiving vaccine doses in every period, we will eventually reach the terminal state where there

is nobody left to vaccinate and the DP recursion is well defined. (More formally, we can take the infinite horizon limit of models with finite horizons, as in the proof of Theorem B.1 below.)

We can interpret (B.1) as follows. The period cost function is the total number of people unvaccinated $(d_0 + \ldots + d_\ell + U_1 - v_2)$ at the end of the week plus the penalty c times the number of people who are overdue for their second dose $(d_0 - v_2)$ at the end of the week. In the state transitions, people with doses due in i weeks (d_i) become due in $i-1$ weeks; those receiving their first dose (v_1) become due in ℓ weeks; the second doses (v_2) reduce the pool of those due for their second dose (d_0); the first doses (v_1) reduce the pool of those who have not yet received a first dose (U_1); both doses deplete the inventory of vaccines (I) and the random supply $(\tilde{a})$ replenishes inventory. The feasible set $\mathbb{V}$ requires second doses administered (v_2) to be less than the number who are due for their second dose (d_0), the number of first doses (v_1) to be less than the number of people who have not yet received a first dose (U_1), and the sum of first and second doses $(v_1 + v_2)$ to be less than or equal to the vaccine available (I).

Our formulation of the DP (B.1) assumes the arriving supplies $\tilde{a}$ are independent and identically distributed over time. This is for ease of notation. We could consider distributions that vary over time and are dependent on previous arrivals: this would require considering value functions $\hat{f}_t$ that depend on time (to capture time dependence) and include past observations of supplies $\tilde{a}_\tau$ $(\tau < t)$ (or a summary statistic based on these observations) as an additional state variable to capture dependence on past supplies. Including time- and/or past-supply dependence would not change any of the results or proofs below. However, including time- and/or past-supply dependence would introduce the opportunities for additional comparative statics analysis: e.g., how do the amounts set aside or the number of first doses given vary with changes in the supply forecasts? We hope to address these kinds of questions in future work.

B.1 $L^\natural$-Convex Functions and Sets

We will use properties of $L^\natural$-convex functions and sets to derive structural properties for this stochastic DP. The concept of $L^\natural$-convexity was introduced in Murota (1998) and, as mentioned in the introduction, related results have been used in the study of inventory models with lead times and/or perishable inventory models (e.g., Zipkin, 2008, Chen et al., 2014). There are several equivalent definitions of $L^\natural$-convexity; we will use the following, adapted from Chen et al. (2018).

Definition B.1 *A function $f : \mathbb{R}^n \to \bar{\mathbb{R}}$ is $L^\natural$-convex if and only if, for all $x', x'' \in \mathbb{R}^n$ and $\omega \geq 0$,*

$$f(x') + f(x'') \geq f((x' + \omega \boldsymbol{e}) \wedge x'') + f(x' \vee (x'' - \omega \boldsymbol{e})) \tag{B.2}$$

where $\boldsymbol{e}$ is an n-vector of all ones. A set is $L^\natural$-convex if its indicator function is $L^\natural$-convex.

Here $\wedge$ and $\vee$ denote componentwise minimization and maximization and $\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$. $L^\natural$-convexity implies ordinary convexity (see, e.g., Murota, 1998) and submodularity which requires (B.2) to hold when $\omega = 0$. More generally, $L^\natural$-convexity is an example of a C5 (closed convex cone containing constants)

property, a class of properties of functions that are often well-behaved in stochastic DP models; see Smith and McCardle (2002).

The following proposition is adapted from Proposition 1 in Chen et al. (2018) and summarizes many useful properties of $L^\natural$-convex functions and sets.

Proposition B.1 (Properties of $L^\natural$-Convexity)

(a) *Any nonnegative linear combination of $L^\natural$-convex functions is $L^\natural$-convex.*

(b) *If f_k is $L^\natural$-convex for $k = 1, 2, \ldots$ and $\lim_{k\to\infty} f_k = f$, then f is $L^\natural$-convex.*

(c) *Assume $f(\cdot, \cdot)$ is defined on $\mathbb{R}^n \times \mathbb{R}^m$ and $f(\cdot, \varepsilon)$ is $L^\natural$-convex for any $\varepsilon \in \mathbb{R}^m$. Then for a random vector $\tilde{\varepsilon}$ defined on $\mathbb{R}^m$, the expectation $\mathbb{E}[f(x, \tilde{\varepsilon})]$ is $L^\natural$-convex, provided it is well defined.*

(d) *If $f : \mathbb{R}^n \to \bar{\mathbb{R}}$ is $L^\natural$-convex, then $g : \mathbb{R}^n \times \mathbb{R} \to \bar{\mathbb{R}}$ defined as $g(x, \lambda) = f(x - \lambda \boldsymbol{e})$ is also $L^\natural$-convex.*

(e) *Assume that $\mathcal{A}$ is a $L^\natural$-convex subset of $\mathbb{R}^n \times \mathbb{R}^m$ and $f(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^m \to \bar{\mathbb{R}}$ is a $L^\natural$-convex function. Let $\boldsymbol{e}$ and $\hat{\boldsymbol{e}}$ denote n- and m-vectors of ones. Then*

(i) *$g(x) = \inf_{\{y:(x,y)\in\mathcal{A}\}} f(x, y)$ is $L^\natural$-convex on $\mathbb{R}^n$ if $g(x) \neq -\infty$ for any $x \in \mathbb{R}^n$.*

(ii) *$\arg\min_{\{y:(x,y)\in\mathcal{A}\}} f(x, y)$ is increasing in x.*

(iii) *For $\omega \geq 0$,*

$$\operatorname*{arg\,min}_{\{y:(x,y)\in\mathcal{A}\}} f(x + \omega \boldsymbol{e}, y) \leq \omega \hat{\boldsymbol{e}} + \operatorname*{arg\,min}_{\{y:(x,y)\in\mathcal{A}\}} f(x, y) \ .$$

In (ii) and (iii) increasing and $\leq$ are defined in terms of the set ordering where $Y' \leq Y''$ if $y' \in Y'$ and $y'' \in Y''$ implies $y' \wedge y'' \in Y'$ and $y' \vee y'' \in Y''$.

(f) *Let x_i denote a component of $x \in \mathbb{R}^n$. Any set of the form $\{x \in \mathbb{R}^n : l \leq x \leq u, x_i - x_j \leq v_{ij}, \forall i \neq j\}$ is $L^\natural$-convex in $\mathbb{R}^n$, where $l, u \in \mathbb{R}^n$ and $v_{ij} \in \mathbb{R}$.*

B.2 Structural Properties of the DP

To study properties of the DP (B.1), it is useful to transform some of the state variables into cumulative form as follows:

$$\begin{aligned}
s_0 &= d_0 && \text{(second doses currently due or overdue)} \\
s_i &= s_{i-1} + d_i \ , \text{ for } i = 1, \ldots, \ell && \text{(second doses due in next } i \text{ weeks)} \\
U_2 &= U_1 + s_\ell && \text{(people who have not yet received their second dose)}
\end{aligned}$$

Nonnegativity of the original variables implies that these transformed variables satisfy $s_0 \leq \ldots \leq s_\ell \leq U_2$.

With these transformed variables, the cost-to-go function may be decomposed into nested optimization problems and rewritten as

$$f(s_0, \ldots, s_\ell, U_2, I) = \min_{\{v_1 \,:\, 0 \le v_1 \le I, v_1 \le U_2 - s_\ell\}} \phi(s_0, \ldots, s_\ell, s_\ell + v_1, U_2, I - v_1) \tag{B.3}$$

where

$$\begin{aligned} \phi(s_0, \ldots, s_\ell, s_{\ell+1}, U_2, I) = \min_{\{v_2 \,:\, 0 \le v_2 \le I, v_2 \le s_0\}} \{(U_2 - v_2) + c(s_0 - v_2) + \\ + \mathbb{E}[f(s_1 - v_2, \ldots, s_\ell - v_2, s_{\ell+1} - v_2, U_2 - v_2, I - v_2 + \tilde{a})]\} \ . \end{aligned} \tag{B.4}$$

Here we choose the allocation to first doses v_1 in the first problem and use this choice in the nested subproblem where the number of second doses v_2 is selected; note that the nested objective function ϕ has one more argument ($s_{\ell+1}$) than f which represents the number of second doses due in the next $\ell + 1$ weeks and is given by $s_\ell + v_1$. The terminal case is $f(0, \ldots, 0, 0, I) = 0$, as before.

Let $v_1^*(s_0, \ldots, s_\ell, U_2, I)$ and $v_2^*(s_0, \ldots, s_{\ell+1}, U_2, I)$ denote optimal first and second doses for (B.3) and (B.4): if there are multiple optimal solutions, we take the smallest optimal number of first doses and largest optimal number of second doses. We define $R^*(s_0, \ldots, s_\ell, U_2, I) = I - v_1^*(s_0, \ldots, s_\ell, U_2, I)$ to represent the optimal amount of vaccine reserved for second doses in the current period or set aside for future periods.

Our main result is the following.

Theorem B.1 (Structural Properties)

(a) *$f(s_0, \ldots, s_\ell, U_2, I)$ and $\phi(s_0, \ldots, s_\ell, s_{\ell+1}, U_2, I)$ are decreasing in I and $L^\natural$-convex.*

(b) *If $c \ge 0$, for $\omega \ge 0$,*

$$f(s_0 + \omega, \ldots, s_\ell + \omega, U_2 + \omega, I + \omega) = f(s_0, \ldots, s_\ell, U_2, I) \tag{B.5}$$

(c) (i) *If $c \ge 0$, $v_2^*(s_0, \ldots, , s_\ell, U_2, I) = s_0 \wedge I$*

(ii) *$R^*(s_0, \ldots, s_\ell, U_2, I)$ is increasing in its arguments and, for any $\omega \ge 0$,*

$$R^*(s_0 + \omega, \ldots, s_\ell + \omega, U_2 + \omega, I + \omega) \le \omega + R^*(s_0, \ldots, s_\ell, U_2, I) \ . \tag{B.6}$$

Parts (a) and (b) of the theorem are of interest here primarily because we use these properties to characterize the optimal policies in part (c). In particular, part (b) implies that, if the penalty $c \ge 0$, it is better to give second doses (thereby removing people from the pipeline of those awaiting vaccination) than to give first doses and move people forward in the pipeline. A proof of the theorem is provided in the online appendix; the proof relies on properties of $L^\natural$-convex functions, as described in Proposition B.1.

We can translate the policy results of part (c) of Theorem B.1 back to the original model and interpret the optimal policy as a set-aside policy of the form considered in §3. Let $\hat{v}_1^*(d_0, \ldots, d_\ell, U_1, I)$ and $\hat{v}_2^*(d_0, \ldots, d_\ell, U_1, I)$ denote the optimal first and second doses for the original DP (B.1) and let

$$S^*(d_0, \ldots, d_\ell, U_1, I) = I - \hat{v}_2^*(d_0, \ldots, d_\ell, U_1, I) - \hat{v}_1^*(d_0, \ldots, d_\ell, U_1, I)$$

denote the optimal quantity to set aside in a given week after administering first and second doses.

Corollary B.2 (Set-aside Policies) *If $c \geq 0$, then*

(a) $\hat{v}_2^*(d_0, \ldots, d_\ell, U_1, I) = d_0 \wedge I$.

(b) *For any $\omega \geq 0$ and weeks until due i, $1 \leq i \leq \ell$, the optimal amount set aside $S^*(d_0, \ldots, d_\ell, U_1, I)$ satisfies*

(i) $0 \leq S^*(d_0, \ldots, d_i + \omega, \ldots, d_\ell, U_1, I) - S^*(d_0, \ldots, d_i, \ldots, d_\ell, U_1, I) \leq \omega$,

(ii) $S^*(d_0, \ldots, d_i + \omega, \ldots, d_\ell, U_1, I) - S^*(d_0, \ldots, d_i, \ldots, d_\ell, U_1, I)$ *is decreasing in i, and*

(iii) $0 \leq S^*(d_0, \ldots, d_\ell, U_1, I + \omega) - S^*(d_0, \ldots, d_\ell, U_1, I) \leq \omega - (\omega \wedge (d_0 - I))^+$.

The interpretation of this result is as discussed in §3. (a) If the penalty $c \geq 0$, it is optimal to prioritize second doses. (b) The optimal amount to set aside (after administering first and second doses) is increasing in the number of second doses due in the coming weeks (d_i) with the increase in setaside bounded by the increase in the amount due (ω) and decreasing in the time until due (i). The optimal setaside is also boundedly increasing in the inventory level, with the bound in part (iii) taking into account the increase (if any) in the second doses enabled by the increase in inventory. The fact that the optimal amount to set aside is *boundedly* increasing in the inventory level implies that the number of first doses given is also (weakly) increasing in the inventory level. This corollary provides the justification for focusing on heuristic policies that have a similar set-aside structure as discussed in §3.3.

C. Proofs for Structural Properties (Online Supplement)

Proof of Theorem B.1. (a) The proof is by backward induction with a given finite time horizon T. To accommodate these arguments, we make the cost-to-go functions in (B.3)-(B.4) depend on the period, written as f_t and ϕ_t, with the terminal value $f_T(\cdot) = 0$; these functions are implicitly dependent on the horizon T as well. We will show that f_t and ϕ_t are decreasing in I and $L^\natural$-convex for all t with a fixed horizon T. We can then appeal to Proposition B.1(b) to conclude that the limit as $T \to \infty$ of f_t and ϕ_t approach stationary cost-to-go-functions f and ϕ which are also decreasing in I and $L^\natural$-convex. Convergence of the value functions as $T \to \infty$ follows if there is a positive probability of receiving vaccine doses (in an amount bounded away from zero) in each period: f_t and ϕ_t are monotonically increasing in the horizon (because the rewards in each period are nonnegative, since $c \geq -1$ and $U_2 \geq s_0$) and bounded above and hence convergent.[3]

Fix the horizon T. The terminal case, $f_T(\cdot) = 0$, is trivially decreasing in I and $L^\natural$-convex. Now assume ϕ_{t+1} and f_{t+1} are decreasing in I and $L^\natural$-convex; we will prove ϕ_t and f_t are decreasing in I and $L^\natural$-convex.

Given the induction hypothesis, it is easy to see that ϕ_t and f_t are decreasing in I: I enters positively in the expressions for ϕ_{t+1} and f_{t+1} in (B.3) and (B.4) (thus larger Is lead to lower continuation values) and increasing I increases the size of the feasible sets in (B.3) and (B.4) (thus larger Is lead to smaller minimum values).

We now show that $\phi_t(s_0, \ldots, s_{\ell+1}, U_2, I)$ is $L^\natural$-convex. The reward function for ϕ_t is linear in the state variables and hence $L^\natural$-convex. The expected continuation value for ϕ_t,

$$\mathbb{E}[f_{t+1}(s_1 - v_2, \ldots, s_{\ell+1} - v_2, U_2 - v_2, I - v_2 + \tilde{a})] \ ,$$

is $L^\natural$-convex in $(s_0, \ldots, s_{\ell+1}, U_2, I, v_2)$ by the induction hypothesis and Proposition B.1(c) and (d). By Proposition B.1(a), the sum of the reward and expected continuation values is $L^\natural$-convex in $(s_0, \ldots, s_{\ell+1}, U_2, I, v_2)$. Note that the feasible set for the minimization problem in (B.4) is $L^\natural$-convex by Proposition B.1(g). $L^\natural$-convexity of $\phi_t(s_0, \ldots, s_{\ell+1}, U_2, I)$ then follows from Proposition B.1(e)(i).

We now show that $f_t(s_0, \ldots, s_{\ell+1}, U_2, I)$ is $L^\natural$-convex. Let

$$(s'_0, \ldots, s'_\ell, U'_2, I') \ \text{ and } \ (s''_0, \ldots, s''_\ell, U''_2, I'') \tag{C.1}$$

denote two sets of state variables and let

$$v'_1 = v^*_1(s'_0, \ldots, s'_\ell, U'_2, I') \ \text{ and } \ v''_1 = v^*_1(s''_0, \ldots, s''_\ell, U''_2, I'') \tag{C.2}$$

denote optimal first doses given these state variables. Fix ω and denote the minimum ($\wedge$) and maximum ($\vee$) values of these pairs of state variables (as appearing in the definition of $L^\natural$-convexity) as follows

$$\begin{array}{rclcrcl} s_i^\wedge & = & (s'_i + \omega) \wedge s''_i & & s_i^\vee & = & s'_i \vee (s''_i - \omega) \\ U^\wedge & = & (U' + \omega) \wedge U'' & & U^\vee & = & U' \vee (U'' - \omega) \end{array} \tag{C.3}$$

and the corresponding first doses and inventory levels as

$$\begin{array}{rclcrcl} v_1^\wedge & = & (s'_\ell + v'_1 + \omega) \wedge (s''_\ell + v''_1) - s_\ell^\wedge & & v_1^\vee & = & (s'_\ell + v'_1) \vee (s''_\ell + v''_1 - \omega) - s_\ell^\vee \\ I^\wedge & = & (I' - v'_1 + \omega) \wedge (I'' - v''_1) + v_1^\wedge & & I^\vee & = & (I' - v'_1) \vee (I'' - v''_1 - \omega) + v_1^\vee \ . \end{array} \tag{C.4}$$

Note that $v_1^\wedge$ ($v_1^\vee$) and $I^\wedge$ ($I^\vee$) need not be the min (max) value of $\{v'_1, v''_1\}$ and $\{I', I''\}$, but are chosen to ensure that $(s_\ell^\wedge + v_1^\wedge) = (s'_\ell + v'_1 + \omega) \wedge (s''_\ell + v''_1)$, $(I^\wedge - v_1^\wedge) = (I' - v'_1 + \omega) \wedge (I'' - v''_1)$ and analogously for

[3] For an upper bound, we can take the value of a (nonoptimal) policy that administers no vaccines until inventory reaches $s_\ell + 2(U_2 - s_\ell)$, the number of doses required to complete vaccination, and then administers $(U_2 - s_\ell)$ first doses and U_2 second doses $\ell + 1$ periods later.

the joins ($\vee$) as required to apply the $L^\natural$-convexity assumption below. Then we have

$$\begin{aligned}
&f_t(s'_0,\ldots,s'_\ell,U'_2,I') + f_t(s''_0,\ldots,s''_\ell,U'_2,I') \\
&\quad= \phi_t(s'_0,\ldots,s'_\ell,s'_\ell+v'_1,U'_2,I'-v'_1) + \phi_t(s''_0,\ldots,s''_\ell,s''_\ell+v''_1,U''_2,I''-v''_1) \\
&\quad\geq \phi_t(s^\wedge_0,\ldots,s^\wedge_\ell,s^\wedge_\ell+v^\wedge_1,U^\wedge_2,I^\wedge-v^\wedge_1) + \phi_t(s^\vee_0,\ldots,s^\vee_\ell,s^\vee_\ell+v^\vee_1,U^\vee_2,I^\vee-v^\vee_1) \\
&\quad\geq f_t(s^\wedge_0,\ldots,s^\wedge_\ell,U^\wedge_2,I^\wedge) + f_t(s^\vee_0,\ldots,s^\vee_\ell,U^\vee_2,I^\vee)\ .
\end{aligned}$$

The first equality follows from the definition of v'_1 and v''_2 as optimal values in (B.3) for the given states. The next inequality follows from the $L^\natural$-convexity of ϕ_t. The final inequality follows from the observation that the first dose choices $v^\wedge_1$ and $v^\vee_1$ are feasible but not necessarily optimal for the minimal and maximal levels of the state variables appearing in the third line above. Thus f_t is $L^\natural$-convex, thereby completing the proof of part (a).

(b) We use an inductive proof as for the result of part (a). The terminal case is trivial. We assume the result of the proposition holds for period $t+1$ and will show that the result for period t as well.

Before proceeding to the result of part (b), we first observe that for $\omega \leq I$,

$$f_t(s_0,\ldots,s_\ell+\omega,U_2,I-\omega) - f_t(s_0,\ldots,s_\ell,U_2,I) \geq -\omega \tag{C.5}$$

Intuitively, administering ω first doses (moving them forward to being due in ℓ weeks and reducing inventory) may improve or hurt the objective, but the improvement can never be more than reducing the time to completion by ω weeks and incurring no additional penalties.

To establish the result of part (b), we first show

$$f_t(s_0,\ldots,,s_\ell,U_2,I) \geq f_t(s_0+\omega,\ldots,s_\ell+\omega,U_2+\omega,I+\omega)$$

and then establish the reverse inequality. Let v'_1 and v'_2 be optimal solutions for (B.3) and (B.4) with state $(s_0,\ldots,s_\ell,U_2,I)$ and let $v''_2 = v'_2+\omega$. Then

$$\begin{aligned}
&f_t(s_0,\ldots,s_\ell,U_2,I) \\
&= (U_2-v'_2) + c(s_0-v'_2) + \\
&\qquad \mathbb{E}[f_{t+1}(s_1-v'_2,\ldots,s_\ell-v'_2,s_\ell+v'_1-v'_2,U_2-v'_2,I-v'_1-v'_2+\tilde{a})] \\
&= (U_2+\omega-v''_2) + c(s_0+\omega-v''_2) + \\
&\qquad \mathbb{E}[f_{t+1}(s_1+\omega-v''_2,\ldots,s_\ell+\omega+v''_1-v'_2,U_2+\omega-v''_2,I+\omega-v'_1-v''_2+\tilde{a})] \\
&\geq f_t(s_0+\omega,\ldots,s_\ell+\omega,U_2+\omega,I+\omega)
\end{aligned}$$

The first equality follows from the definition of v'_1 and v'_2 as optimal solutions. The second equality follows from substituting $v'_2 = v''_2-\omega$. The inequality follows from the fact that (v'_1,v''_2) is feasible given $(s_0+\omega,\ldots,s_\ell+\omega,U_2+\omega,I+\omega)$ but not necessarily optimal.

The reverse inequality by considering two cases. Let v''_1 and v''_2 be optimal given $(s_0+\omega,\ldots,s_\ell+\omega,U_2+\omega,I+\omega)$. First consider the case where $v''_2 \geq \omega$ and let $v'_2 = (v''_2-\omega)$. Then

$$\begin{aligned}
&f_t(s_0+\omega,\ldots,s_\ell+\omega,U_2+\omega,I+\omega) \\
&= (U_2+\omega-v''_2) + c(s_0+\omega-v''_2) + \\
&\qquad \mathbb{E}[f_{t+1}(s_1+\omega-v''_2,\ldots,s_\ell+\omega+v''_1-v''_2,U_2+\omega-v''_2,I-v'_1-v''_2+\omega+\tilde{a})] \\
&= (U_2-v'_2) + c(s_0-v'_2) + \\
&\qquad \mathbb{E}[f_{t+1}(s_1-v'_2,\ldots,s_\ell+v''_1-v'_2,U_2-v'_2,I-v''_1-v'_2+\tilde{a})] \\
&\geq f_t(s_0,\ldots,s_\ell,U_2,I)
\end{aligned}$$

The first equality follows from (v''_1,v''_2) being an optimal solution. The second equality follows from substituting $v'_2 = v''_2-\omega$. The inequality follows from the fact that (v''_1,v'_2) is feasible given $(s_0,\ldots,s_\ell,U_2,I)$ (since $\omega \leq v''_2 \leq s_0+\omega$, we have $0 \leq v'_2 \leq s_0$) but not necessarily optimal.

Now consider the case where $v_2'' \le \omega$. Let $\Delta = \omega - v_2'' \ge 0$, $v_1' = (v_1'' - \Delta)^+$, and $\Delta' = v_1'' - v_1'$. Note that $0 \le \Delta' \le \Delta$. Then

$$\begin{aligned}
&f_t(s_0+\omega,\ldots,s_\ell+\omega,U_2+\omega,I+\omega)\\
&=(U_2+\omega-v_2'')+c(s_0+\omega-v_2'')+\\
&\qquad\mathbb{E}[f_{t+1}(s_1+\omega-v_2'',\ldots,s_\ell+\omega+v_1''-v_2'',U_2+\omega-v_2'',I-v_1''-v_2''+\omega+\tilde{a})]\\
&=(1+c)\Delta+U_2+cs_0+\mathbb{E}[f_{t+1}(s_1+\Delta,\ldots,s_\ell+\Delta+v_1'+\Delta',U_2+\Delta,I+\Delta-v_1'-\Delta'+\tilde{a})]\\
&=(1+c)\Delta+U_2+cs_0+\mathbb{E}[f_{t+1}(s_1,\ldots,s_\ell+v_1'+\Delta',U_2,I-v_1'-\Delta'+\tilde{a})]\\
&\ge(1+c)\Delta-\Delta'+U_2+cs_0+\mathbb{E}[f_{t+1}(s_1,\ldots,s_\ell+v_1',U_2,I-v_1'+\tilde{a})]\\
&\ge(1+c)\Delta-\Delta'+f_t(s_0,\ldots,s_\ell,U_2,I)\\
&\ge f_t(s_0,\ldots,s_\ell,U_2,I)
\end{aligned}$$

The first equality follows from the definition of $(v_1'',\, v_2'')$ as an optimal solution for the given state. The next two equalities follow from the definition of Δ and the induction hypothesis. The first inequality follows from the observation (C.5) and the second from the fact that $(v_1', 0)$ is feasible but not necessarily optimal given $(s_0,\ldots,s_\ell,U_2,I)$. The final inequality follows because $(1+c)\Delta - \Delta' \ge 0$ when $c \ge 0$. Thus we have

$$f_t(s_0+\omega,\ldots,s_\ell+\omega,U_2+\omega,I+\omega) \ge f_t(s_0,\ldots,s_\ell,U_2,I)$$

thereby completing the proof of part (b).

(c)(i) follows from part (b). Given $c \ge 0$, If $s_0 \le I$, then (b) implies

$$f(s_0,\ldots,s_\ell,U_2,I) = f(0,s_1-s_0,\ldots,s_\ell-s_0,U_2-s_0,I-s_0)$$

that is, it is optimal to allocate s_0 second doses (the latter expression is equivalent to the value when committing to s_0 second doses). If $s_0 \ge I$, then (b) implies

$$f(s_0,\ldots,s_\ell,U_2,I) = f(s_0-I,\ldots,s_\ell-I,U_2-I,0)$$

that is, it is optimal to allocate all of the inventory I to second doses.

For (c)(ii) consider a variation of (B.3) where given $(s_0,\ldots,s_\ell,s_{\ell+1},U_2,I)$ the decision maker selects the number of doses R to reserve for second doses this period or later use:

$$r^*(s_0,\ldots,s_\ell,s_{\ell+1},U_2,I) = \operatorname*{arg\,min}_{\{R\,:\,0\le R\le I\}} \phi(s_0,\ldots,s_\ell,s_{\ell+1},U_2,R)\ . \tag{C.6}$$

Here the decision maker takes the number of second doses due in $\ell+1$ periods to be exogenously given rather than depending on the number of first doses administered in the current period (i.e., in the original model this would be $s_{\ell+1} = s_\ell + I - R$). Note that this exogeneity assumption only affects the current period: in all future periods, first doses are allocated optimally as captured in ϕ. (If there are multiple optimal solutions, we take the largest optimal value.) Applying Proposition B.1(c), we have the following lemma.

Lemma 1 *$r^*(s_0,\ldots,s_\ell,s_{\ell+1},U_2,I)$ is increasing in its arguments and, for all $\omega > 0$,*

$$r^*(s_0+\omega,\ldots,s_\ell+\omega,s_{\ell+1}+\omega,U_2+\omega,I+\omega) \le \omega + r^*(s_0,\ldots,s_\ell,s_{\ell+1},U_2,I)\ . \tag{C.7}$$

Proof of Lemma 1. $L^\natural$-convexity of ϕ was established in Theorem B.1(a) and the feasible set in (C.6) is $L^\natural$-convex by Proposition B.1(f). The result then follows from Proposition B.1(e)(i) and (ii). ■

Now consider ordered state variables

$$(s_0',\ldots,s_\ell',s_{\ell+1}',U_2',I') \le (s_0'',\ldots,s_\ell'',s_{\ell+1}'',U_2'',I'')$$

and define

$$r'(\Delta) = r^*(s'_0, \ldots, s'_\ell, s'_\ell + \Delta, U'_2, I') \text{ and } r''(\Delta) = r^*(s''_0, \ldots, s''_\ell, s''_\ell + \Delta, U''_2, I'') \ .$$

From the "increasing" part of Lemma 1, we know that $r'(\Delta)$ and $r''(\Delta)$ are increasing in Δ and that $r'(\Delta) \leq r''(\Delta)$, as illustrated in Figure 10.

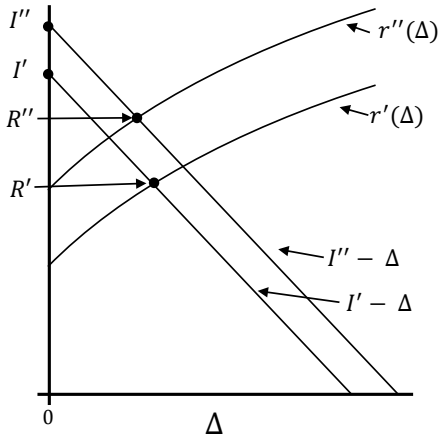

Figure 10: Illustration for proof of Theorem A.1(c)

A feasible policy when taking $s_{\ell+1}$ to be endogenously defined must choose a reserve level $r(\Delta) = I - \Delta$, where Δ is the number of first doses given. Thus, the true optimal reserve in the low scenario, $R' = R^*(s'_0, \ldots, s'_\ell, U'_2, I')$, is the value where $r'(\Delta)$ and $I' - \Delta$ cross, as shown in the figure: at this reserve level, the exogenously specified $s_{\ell+1} = s_\ell + \Delta$ and chosen $s_{\ell+1} = s_\ell + I' - r'(\Delta)$ coincide and hence R' is optimal with endogenously chosen $s_{\ell+1}$. R'' is defined analogously as the crossing value of $r''(\Delta)$ and $I'' - \Delta$.

We claim that $R'' \geq R'$ (as in Figure 10) which would establish the "increasing" part of Theorem (c)(ii). Let Δ' and Δ'' be the values of Δ corresponding to the points of intersection defining R' and R''. If $\Delta' \geq \Delta''$ (as in the figure), then

$$R'' - R' = (I'' - \Delta'') - (I' - \Delta') = (I'' - I') + (\Delta' - \Delta'') \geq 0 \ .$$

If $\Delta' \leq \Delta''$, then

$$R'' - R' = r''(\Delta'') - r'(\Delta') \geq r''(\Delta'') - r'(\Delta'') \geq 0 \ ,$$

where the first inequality comes from the fact that $r'(\Delta)$ is increasing and the second from $r'(\Delta) \leq r''(\Delta)$. This completes the proof of the "increasing" part of Theorem (c)(ii).[4]

Finally, to establish (B.6) of Theorem (c)(ii) define

$$r'(\Delta) = r^*(s_0, \ldots, s_\ell, s_\ell + \Delta, U_2, I) \text{ and } r''(\Delta) = r^*(s_0 + \omega, \ldots, s_\ell + \omega, s_\ell + \omega + \Delta, U_2 + \omega, I + \omega)$$

and let $I' = I$ and $I'' = I + \omega$. By Lemma 1, $r'(\Delta)$ and $r''(\Delta)$ are increasing and $r'(\Delta) \leq r''(\Delta)$, again as in Figure 10. The lemma also says that $r''(\Delta) \leq r'(\Delta) + \omega$. The optimal values $R' = R^*(s_0, \ldots, s_\ell, U_2, I)$ and $R'' = R^*(s_0 + \omega, \ldots, s_\ell + \omega, U_2 + \omega, I + \omega)$ are defined as the crossing values with $I' - \Delta$ and $I'' - \Delta$, as above.

We can complete the proof by showing that $R'' - R' \leq \omega$. Let Δ' and Δ'' be the values of Δ corresponding to the points of intersection defining R' and R''. If $\Delta' \leq \Delta''$, then

$$R'' - R' = (I'' - \Delta'') - (I' - \Delta') = \omega + (\Delta' - \Delta'') \leq \omega \ .$$

[4]Note that this argument assumes that the constraint $v_1 \leq U_2 - s_\ell$ (requiring the number of first doses given to be less than or equal to the number of people who have not received a first dose) is not binding. However, the argument can easily be adapted to handle the case where this constraint is binding.

If $\Delta' \geq \Delta''$, then

$$R'' - R' = r''(\Delta'') - r'(\Delta') \leq r'(\Delta'') + \omega - r'(\Delta') \leq r'(\Delta'') + \omega - r'(\Delta'') = \omega \ ,$$

where the first inequality comes from the fact that $r''(\Delta) \leq r'(\Delta) + \omega$ (from the lemma) and the second from the fact that $r'(\Delta)$ is increasing (also from the lemma). Thus $R'' - R' \leq \omega$, completing our proof of the theorem. ■

Proof of Corollary B.2. (a) follows immediately from Theorem B.1(c)(i) given that $s_0 = d_0$.

(b) follows from (B.6). For (i) and (ii), first note that the changes contemplated do not affect the number of second doses given, so the properties of differences in S^* reflect the difference in R^*. The result of (B.6) says that increasing all arguments of R^* by ω (as on the left side of (B.6)) increases R^* by ω or less. Since R^* is increasing, this implies that an increase of any subset of arguments by ω also increases R^* by ω or less. For (i), note that increasing d_i to $d_i + \omega$ increases the transformed variables $(s_i, \ldots, s_\ell, U_2)$ to $(s_i + \omega, \ldots, s_\ell + \omega, U_2 + \omega)$, leaving the other variables unchanged. For (ii), note that a larger i increases fewer arguments and results in (weakly) smaller increases.

For (iii), note that increasing the inventory level may increase the number of second doses given if it was previously constrained by the inventory level (i.e., with $v_2 = I$): in this case, increasing the inventory level will result in up to $\omega \wedge (d_0 - I)$ additional second doses being allocated, which reduces the potential increase in S^* accordingly. The bounded increase in S^* follows from the same argument (increasing a single parameter) as above. ■

References

AP. 2021. U.S. asking states to speed vaccine, not hold back 2nd dose. *Modern Healthcare* URL https://www.modernhealthcare.com/policy/us-asking-states-speed-vaccine-not-hold-back-2nd-dose.

Appuzzo, M., S. Gebrekidan, M. Pronczuk. 2021. Where europe went wrong in its vaccine rollout, and why. *The New York Times* URL https://www.nytimes.com/2021/03/20/world/europe/europe-vaccine-rollout-astrazeneca.html.

Barnabas, R.V., A. Wald. 2021. A public health covid-19 vaccination strategy to maximize the health gains for every single vaccine dose. *Annals of internal medicine* URL https://doi.org/10.7326/M20-8060.

Bieniasz, P. 2021. The case against delaying SARS-CoV-2 mRNA vaccine boosting doses. *Clinical Infectious Diseases* URL https://doi.org/10.1093/cid/ciab070.

Brown, D.B., J.E. Smith, P. Sun. 2010. Information relaxations and duality in stochastic dynamic programs. *Operations Research* **58**(4-part-1) 785–801. doi:10.1287/opre.1090.0796. URL https://doi.org/10.1287/opre.1090.0796.

Callamard, A. 2021. Covid vaccine in short supply in nepal: A week to help pull nation back from the brink – what britain must do. *Amnesty International News* URL https://www.amnesty.org/en/latest/news/2021/06/covid-vaccine-in-short-supply-in-nepal-a-week-to-help-pull-nation-back-from-the-brink-what-britain-must-do.

Chen, J., S. Hoops, Marathe A., H. Mortveit, B. Lewis, S. Venkatramanan, et al. 2021. Prioritizing allocation of COVID-19 vaccines based on social contacts increases vaccination effectiveness. *medRxiv* URL https://doi.org/10.1101/2021.02.04.21251012.

Chen, X., X. Gao, Z. Pang. 2018. Preservation of structural properties in optimization with decisions truncated by random variables and its applications. *Operations Research* **66**(2) 340–357.

Chen, X., M. Li, D. Simchi-Levi, T. Zhao. 2020. Allocation of COVID-19 vaccines under limited supply. *Available at SSRN 3678986* .

Chen, X., Z. Pang, L. Pan. 2014. Coordinating inventory control and pricing strategies for perishable products. *Operations Research* **62**(2) 284–300. doi:10.1287/opre.2014.1261.

Choi, C., M. Renault. 2021. U.S. vaccine drive complicated by 1st and 2nd dose juggling act. *Seattle Times* URL https://www.seattletimes.com/seattle-news/health/giving-1st-and-2nd-doses-at-once-complicates-vaccine-effort/.

Cimons, M. 2018. Don't panic if you're put on a a waiting list for the new shingles vaccine. *The Washington Post* URL https://www.washingtonpost.com/national/health-science/dont-panic-if-youre-

put-on-a-waiting-list-for-the-new-shingles-vaccine/2018/07/24/248d972e-8e8d-11e8-8322-b5482bf5e0f5_story.html.

Dai, T., J-S Song. 2021. Transforming covid-19 vaccines into vaccination: Challenges and opportunities for management scientists. *Health Care Management Science* .

Deo, S., C.J. Corbett. 2009. Cournot competition under yield uncertainty: The case of the us influenza vaccine market. *Manufacturing & Service Operations Management* **11.4** 563–576.

Deshpande, V., M. A. Cohen, K. Donohue. 2003. A threshold inventory rationing policy for service-differentiated demand classes. *Management Science* **49.6** 683—703.

Edwards, E. 2021. In push to get more vaccines into arms, officials recommend states give to anyone 65 and up. *NBC News* URL https://www.nbcnews.com/health/health-news/cdc-recommend-states-give-covid-vaccine-anyone-65-older-n1253887.

Enayati, S., O. Y Özaltın. 2020. Optimal influenza vaccine distribution with equity. *European Journal of Operational Research* **283**(2) 714–725.

Foy, B. H., B. Wahl, K. Mehta, A. Shet, G. I. Menon, C. Britto. 2021. Comparing COVID-19 vaccine allocation strategies in india: a mathematical modelling study. *International Journal of Infectious Diseases* **103** 431—438.

Frost, M. 2021. Gov. Hogan: Maryland is not holding back second dose of COVID vaccine. *Fox 45 News* URL https://foxbaltimore.com/news/local/gov-hogan-maryland-is-not-holding-back-second-dose-of-covid-vaccine.

Green, L. V., S. Savin. 2008. Reducing delays for medical appointments: A queuing approach. *Operations Research* **56.6** 1526—1538.

Harper, K.B. 2021. Hundreds of thousands of COVID-19 vaccine deliveries, injections delayed by winter storm. *The Texas Tribune* URL https://www.dallasnews.com/news/weather/2021/02/18/hundreds-of-thousands-of-covid-19-vaccine-deliveries-injections-delayed-by-winter-storm/.

Ionova, A. 2021. Millions in brazil are missing their second vaccine dose, adding to the burden of a hard-hit nation. *The New York Times* URL https://www.nytimes.com/2021/06/27/world/brazil-covid-vaccine.html.

Kadire, S. R., R. M. Wachter, N. Lurie. 2021. Delayed second dose versus standard regimen for COVID-19 vaccination. *New England Journal of Medicine* URL https://doi.org/10.1056/NEJMclde2101987.

Klein, B. 2021. White house cautions providers against holding back vaccine doses. *CNN politics* URL http://www.cnn.com/2021/02/01/politics/vaccines-doses-white-house-covid-19/index.html.

Koçağa, Y. L., A. Şen. 2007. Spare parts inventory management with demand lead times and rationing. *IIE Transactions* **39.9** 879—898.

Liu, N., S. Ziya, V. G. Kulkarni. 2010. Dynamic scheduling of outpatient appointments under patient no-shows and cancellations. *Manufacturing & Service Operations Management* **12.2** 347—364.

Mak, H-Y, T. Dai, C.S. Tang. 2021. Managing two-dose COVID-19 vaccine rollouts with limited supply. *Available at SSRN 3790836* .

Mass Department of Public Health. 2021. Weekly COVID-19 vaccination report URL `https://www.mass.gov/doc/weekly-covid-19-vaccination-report-february-18-2021/download`.

McPhillips, D. 2021. More than 1 in 10 people have missed their second dose of Covid-19 vaccine. *CNN health* URL `https://www.cnn.com/2021/06/24/health/missed-second-doses-delta/index.html`.

Miller, H., M. Cohn. 2021. Maryland facing shortage of second COVID vaccine doses, health secretary says. *Baltimore Sun* URL `https://www.baltimoresun.com/coronavirus/bs-md-vaccine-second-dose-shortages-20210201-lsx5bh6rjfh7ra3rhpjh247kxi-story.html`.

Moghadas, S. M., T. N. Vilches, K. Zhang, S. Nourbakhsh, P. Sah, M. C. Fitzpatrick, A. P. Galvani. 2021. Evaluation of COVID-19 vaccination strategies with a delayed second dose. *medRxiv* URL `https://doi.org/10.1101/2021.01.27.21250619`.

Murota, K. 1998. Discrete convex analysis. *Mathematical Programming* **83**(1) 313–371.

O'Donnell, C. 2021. Scarce niche syringes complicate U.S. plan to squeeze more COVID shots from Pfizer vials. *Reuters* URL `https://www.reuters.com/business/healthcare-pharmaceuticals/exclusive-us-goal-squeeze-more-covid-shots-pfizer-vials-hampered-by-syringe-2021-01-25/`.

Paltiel, A. D., A. Zheng, J. L. Schwartz. 2021. Speed versus efficacy: quantifying potential tradeoffs in COVID-19 vaccine deployment. *Annals of Internal Medicine* URL `https://doi.org/10.7326/M20-7866`.

Persad, G., M. E. Peek, E. J. Emanuel. 2020. Fairly prioritizing groups for access to COVID-19 vaccines. *JAMA* **324(16)** 1601—1602.

Reuters. 2021. Chile shifts covid vaccine drive to second doses amid concerns over first shot's protection. *Reuters* URL `https://www.nbcnews.com/news/latino/chile-shifts-covid-vaccine-drive-second-doses-concerns-first-shots-pro-rcna730`.

Smith, J. E., K.F. McCardle. 2002. Structural properties of stochastic dynamic programs. *Operations Research* **50**(5) 796–809.

Thomas, K. 2020. Hospitals discover a surprise in their vaccine deliveries: Extra doses. *The New York Times* URL `https://www.nytimes.com/2020/12/16/health/Covid-Pfizer-vaccine-extra-doses.html`.

Tuite, A. R., L. Zhu, D. N. Fisman, J. A. Salomon. 2021. Alternative dose allocation strategies to increase benefits from constrained COVID-19 vaccine supply. *Annals of Internal Medicine* URL `https://doi.org/10.7326/M20-8137`.

Zipkin, P. 2008. On the structure of lost-sales inventory models. *Operations Research* **56**(4) 937–944. doi:10.1287/opre.1070.0482.