

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

Robert A. Shumsky¹, James E. Smith¹, Anne G. Hoen², and Michael Gilbert³

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

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

³Catholic Medical Center, Manchester, New Hampshire

March 29, 2021

Abstract

The two main COVID-19 vaccines currently available in the U.S. are typically administered in two doses, with a prescribed number of weeks separating the two doses. Because of uncertainty in vaccine supply, many vaccination centers are saving doses in inventory to ensure on-time second doses. However, saving doses in inventory 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 first and second doses in the face of supply uncertainty. The structure of the model suggests simple “set-aside” policies that reserve doses for second doses that are due in the coming weeks. For example, using parameters based on data collected from a hospital-run vaccination center, we recommend that each week: (i) the vaccination center should complete all second doses due, if possible, and (ii) if doses remain, they should set aside vaccine for second doses scheduled for the following one to two weeks. Any doses beyond this should be used for first doses. A policy that ensures no second doses will be delayed appears to be overly conservative. In our 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; this set-aside policy is thus nearly optimal given uncertainty in supply. We also show that set-aside policies perform well for larger-scale vaccination centers and centers with supply parameters based on state-level data. Finally, we use the model to quantify the value of reducing the variability and uncertainty in vaccine supply.

1. Introduction

On February 9, 2021, the Associated Press reported on the logistical challenges presented by two-dose COVID-19 vaccines:

The need to give each person two doses a few weeks apart vastly complicates the country’s biggest-ever vaccination campaign. And persistent uncertainty about future vaccine supplies fuels worries that some people will not be able to get their second shots in time. . . . In some cases, local health departments and providers have said they must temporarily curb or even cancel appointments for first doses to ensure there are enough second doses for people who need them. (Choi and Renault, 2021)

The CDC recommends that two vaccines for COVID-19 be administered in two doses, with the Pfizer-BioNTech doses administered three weeks apart (CDC, 2021a) and the Moderna doses administered four weeks apart (CDC, 2021b); the CDC says “the second dose should be administered as close to the recommended interval as possible” (CDC, 2021c). Given uncertainty in future supplies, however, vaccine supply chain managers and site administrators at the national, state, and local levels have struggled to decide how many doses of vaccine to hold in inventory for future second doses (AP, 2021; Frost, 2021). On one hand, as then-Health and Human Services Secretary Alex Azar stated in January 2021, “Every vaccine dose that is sitting in a warehouse rather than going into an arm could mean one more life lost.” (Edwards, 2021). On the other hand, experts are concerned about delayed second doses, so that holding inventory is necessary when there is uncertain supply (Goodin-Smith, 2021; Miller and Cohn, 2021). Current inventory policies vary across vaccination centers, U.S. states, and countries (Choi and Renault, 2021; AP, 2021; Frost, 2021; Appuzzo et al., 2021).

To provide a sense of the uncertainty in vaccine supply, Figure 1 shows weekly supplies of vaccine to the vaccine clinic operated by the Catholic Medical Center (CMC) in Manchester, New Hampshire from the New Hampshire Division of Public Health Services. During this time, CMC was responsible for vaccinating approximately 3,000 people as part of Phase 1a where doses are administered to health workers and older adults in residential care settings (NH Bureau of Infection Disease Control, 2021).

Given that the vaccine supply is uncertain and on-time administration of the second dose is recommended, it may make sense to maintain some inventory of vaccines for second doses. However, holding back large amounts of vaccine slows the administration of first doses and may delay the overall rate of vaccination. To study this tradeoff, we develop a model of a vaccination center, e.g., a center operated by a hospital or a state. Each week the center receives a supply of vaccine and the center’s goal is to vaccinate the target population as quickly as possible. The model has a nice structure that suggests “set-aside” policies that reserve doses for second doses that are due in coming weeks. These policies are easy to implement and perform well in numerical experiments, close to a theoretical performance bound generated by the optimal

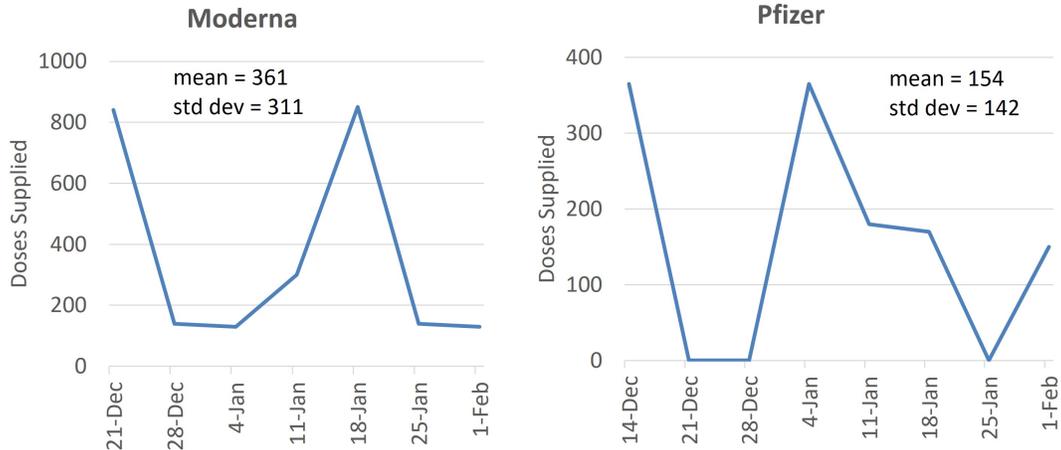


Figure 1: Vaccine supplies to the Catholic Medical Center in Manchester, NH

allocation given perfect information about vaccine supply. Therefore the set-aside policies we propose are nearly optimal with uncertainty in 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 our recommendations perform given the level of uncertainty faced by CMC and by state-level vaccination centers. We also test the model over a wide range of alternative parameters. 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.

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 at the recommended time; 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 will focus on optimizing the timely completion of the two-dose regimen that was proven safe and effective in large trials. We also do not consider how to prioritize target populations; this is an important area of ongoing research that touches on ethics, epidemiology, and optimization (Enayati and Özaltn, 2020;

Chen et al., 2020; Persad et al., 2020; Chen et al., 2021, Foy et al., 2021).

As in our model, Tuite et al. (2021) prioritize second doses and explores the benefits and costs of reserving vaccines to ensure on-time second doses. Their model describes the problem on a national level, as opposed to our focus on a particular vaccination center. In particular, their model assumes a deterministic vaccine supply with predictable changes in volumes, whereas we focus on the effects of uncertainty in the supply.

Within the operations management literature, our vaccine reservation problem has some similarity to problems in inventory theory with multiple demand classes and rationing (e.g., Deshpande et al., 2003; Koçağa and Şen, 2007), for people requiring first and second doses can be seen as two customer classes. In our case, however, the inventory model is complicated by the fact that individual second doses are linked to individual first doses. The 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), but the model structure (with second doses tied to earlier first doses) is different from those in this literature. 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 captures vaccine supply and administration on a weekly basis. First, a random number of doses arrive at the center. The uncertainty in supply stems from many sources including issues in the manufacturing processes, national and state policies that change over time, weather conditions (Harper, 2021), whether “low dead space” syringes are available to extract extra doses from vials of vaccine (Rowland, 2021), among many other things.

The arriving doses, as well as any doses currently in inventory, may then 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 from the vaccination center also receives a second dose from the same center 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.

We consider two performance metrics: (i) the average time until people receive both doses (*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 $c \geq 0$:

$$\begin{aligned} (\text{penalized average time to completion}) = & \hspace{15em} (1) \\ & (\text{average time to completion}) + c \times (\text{average delay in second dose}) \end{aligned}$$

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

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 . To settle on a particular value for c , a decision-maker might consider the potential reduction in the efficacy of the vaccine if the second dose is delayed; however, this is something that is not currently well understood based on clinical trial data (Amit et al., 2021). Alternatively, a decision-maker might consider the potential delays experienced with particular policies (e.g., by considering Figures 2, 3, and 6 below) and ask if they are comfortable with these tradeoffs: e.g., would it be acceptable if, say, 40% (or 10%) of the population experience a delay of one or two weeks if it means vaccinating people one week earlier on average?

This model can be formulated as a stochastic dynamic program (DP); we formally describe this DP in the appendix. Unfortunately, the DP is difficult (or impossible) to solve to optimality because the state space is quite large. The cost-to-go function and optimal policy for the DP will depend on: the number of doses of vaccine supplied in the given week; the number of vaccines in inventory that week; 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 (or impossible) 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 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 amount set aside varies with the number of doses due and the set-asides are greater for doses due sooner.

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 §3.2.

3.1 Set-aside Policies

The simplified set-aside policies we focus on 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, 50 and 30 doses are coming due in the next one, two, three, and four weeks, respectively, and if $(w_1, w_2, w_3, w_4) = (1, 1, 0.5, 0)$, then we reserve $1 \times 60 + 1 \times 40 + 0.5 \times 50 + 0 \times 30 = 125$ 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,

$$(1 - 0.5w_4) (D_t - w_1v_1 - w_2v_2 - w_3v_3)^+ \tag{2}$$

where v_i denotes the number of second doses due in i weeks and $(x)^+ = \max(x, 0)$. (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.

When $(w_1, w_2, w_3, w_4) = (1, 1, 1, 1)$, a second dose is held in inventory for every first dose administered 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. We will refer to this special case of the set-aside policy as 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. 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 $x = 0, 0.5, 1, \dots, 4$.¹

3.2 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 the appendix. 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.

4. Numerical Experiments

The first set of model parameters that we use in our numerical experiments are based on the CMC data described in the introduction. We assume that there are 3,000 people in the target population and that the number of vaccine doses supplied each week is independent from week to week and normally distributed rectified at 0, meaning any negative observations are reset to zero (rectified).² Specifically, we assume this rectified normal distribution has mean 361 and standard deviation 311 (both after rectification), matching the observed variability of Moderna doses supplied to CMC; this corresponds to mean 318.6 and standard

¹In 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.

²We also evaluated these policies using a Beta distribution with the same mean and variance and with a range from 130 to 850, the range observed by CMC. The results were nearly identical to results generated using the rectified normal distribution.

deviation 373.7 before rectification. These parameters imply that there is a 20% chance of receiving zero doses of vaccine in a given week.

Though weekly vaccine supplies to CMC have a coefficient of variation (CV) of $311/361 = 0.86$, we would expect that other vaccination centers face different demand patterns. Therefore we also test our allocation policies using models with supply parameters based on state-level data. Our analysis of CDC vaccine shipments to states (CDC, 2021d; CDC, 2021e) and doses received by Massachusetts (Mass Department of Public Health, 2021, slide 4) suggest that weekly deliveries to states have had CVs in the range $0.25 - 0.35$. Therefore, we test the model with a standard deviation = 100 (after rectification), for a CV of $100/361 = 0.28$. This corresponds to a normal distribution with mean 358.5 and standard deviation 100 before rectification and less than a 0.2% chance of receiving zero doses in a given week. We will refer to models with this supply distribution as having “low uncertainty” and those based on the CMC data as having “high uncertainty.” We evaluate the proposed policies in both cases using Monte Carlo simulations with 50,000 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 2 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 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 from 311 to 100: in the low uncertainty scenario, the average completion time is almost a week less when no inventory is set aside. Note that the values shown in Figure 2 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 0.004 weeks.)

Figure 3 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 one week on average in the high uncertainty case and approximately 0.2 weeks in the low uncertainty case; see Figure 3a.

Focusing on the high uncertainty case, in Figure 3b we see that when no doses are set aside ($x = 0$)

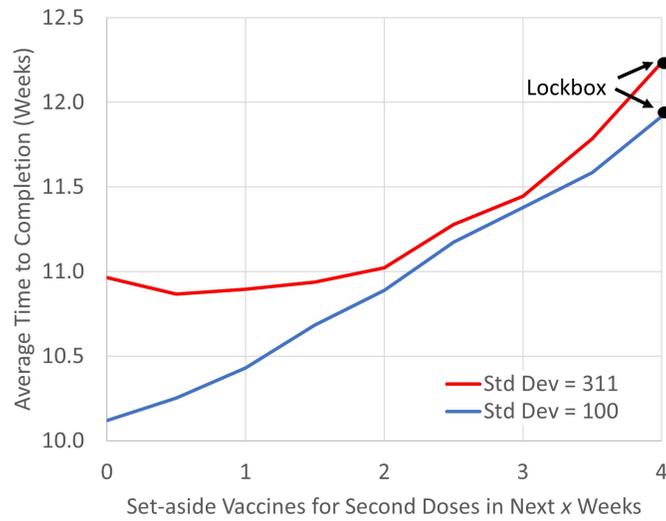


Figure 2: Average time to completion under set-aside policies

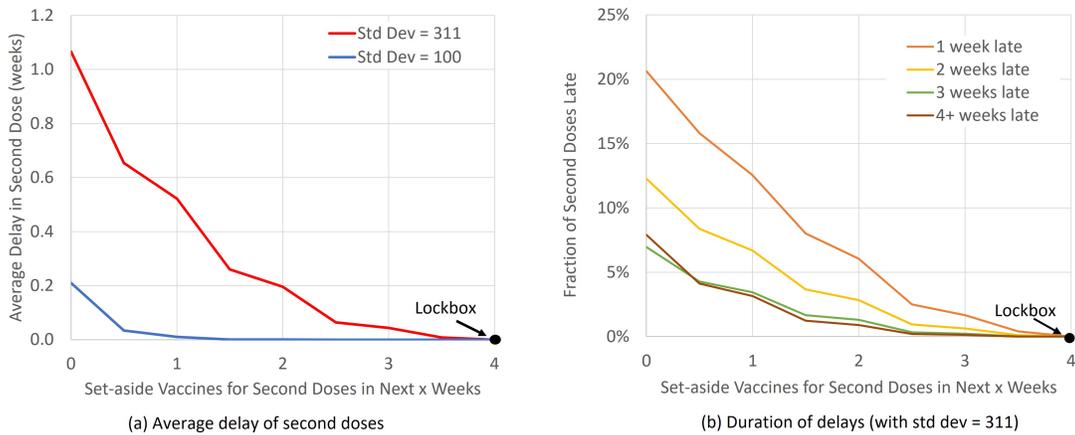


Figure 3: Delays in second dose for set-aside policies

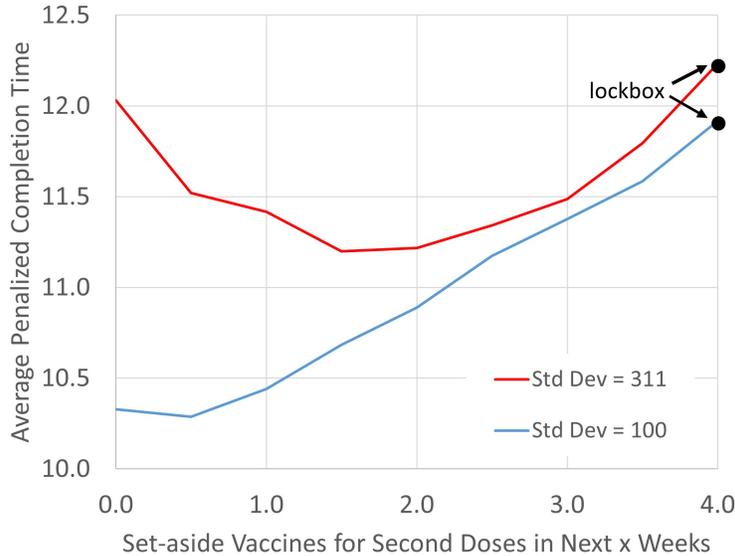


Figure 4: Average penalized completion time for varying set-aside quantities with $c = 1$

approximately 48% of people experience some delay in getting their second dose: 21% experience a delay of one week; 12% experience a delay of two weeks; 7% experience a delay of 3 weeks; and 8% experience a delay of 4 weeks or more. Recall that the probability of receiving zero doses in a given week is 20% 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. Moreover, the vaccination center can get further behind when 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 4 shows the penalized average completion time as the set-aside parameter x varies and with $c = 1$. 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 low uncertainty case, 0.5 weeks set aside is optimal. In general, as c increases and we place more weight on delaying second doses, it is optimal to set-aside more vaccine. Figure 5 shows how the penalized time to completion changes as c varies from 0 to 3 for the high and low uncertainty cases; the corresponding recommended set-asides are shown in the figure. In Figure 5, we see that the lockbox policies are too conservative, even with penalties as high as $c = 3$.

Although the CDC recommends on-time second doses, the CDC also states that a delay of up to two weeks may be acceptable (CDC, 2021c). Given high supply uncertainty, in Figure 3b we see that a 1.5-week

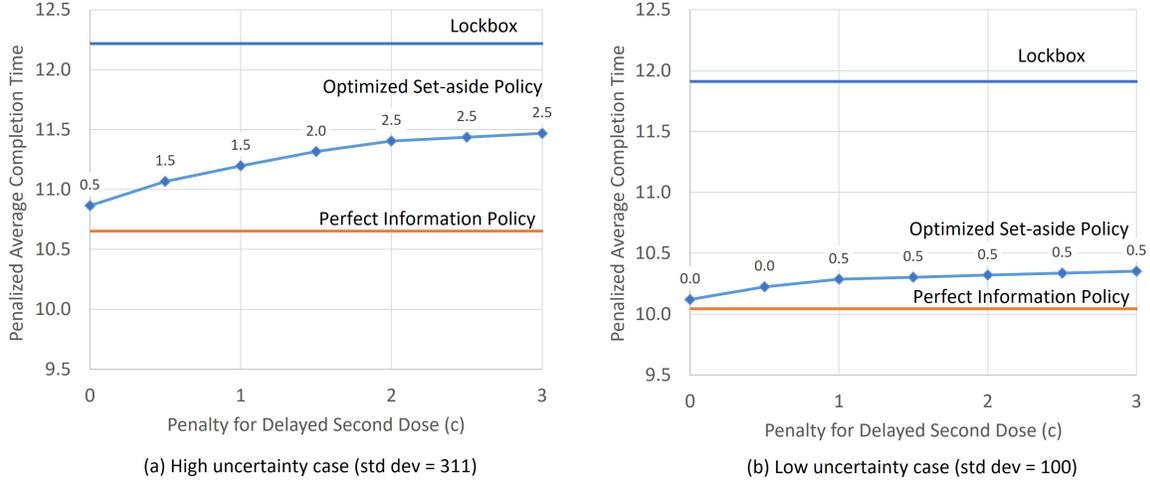


Figure 5: Penalized average completion times for optimized set-aside policies

set-aside policy has a low probability of delays greater than two weeks. In Figure 5, we see that set-aside policies of 1.5 to 2 weeks are optimal or near-optimal over a wide range of penalties c . Therefore, we recommend set-asides between one and two weeks when there is high uncertainty. With low uncertainty, a set-aside of half a week is optimal over a broad range.

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 2. 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 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 low uncertainty case 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 5 also shows the perfect information lower bound which maintains supply variability but eliminates supply uncertainty. Here we see that eliminating uncertainty can be valuable, although the set-aside policy is close to the lower bound when c is small: with $c = 0$, the optimal set-aside policy (setting aside 0.5 weeks) is within 2.0% of the perfect information lower bound in the high uncertainty case and within 0.7% in the low uncertainty case. With $c = 1$, the optimal set-aside policy is 5.1% higher than the lower bound in the high uncertainty case and 2.4% higher in the low standard deviation case. 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

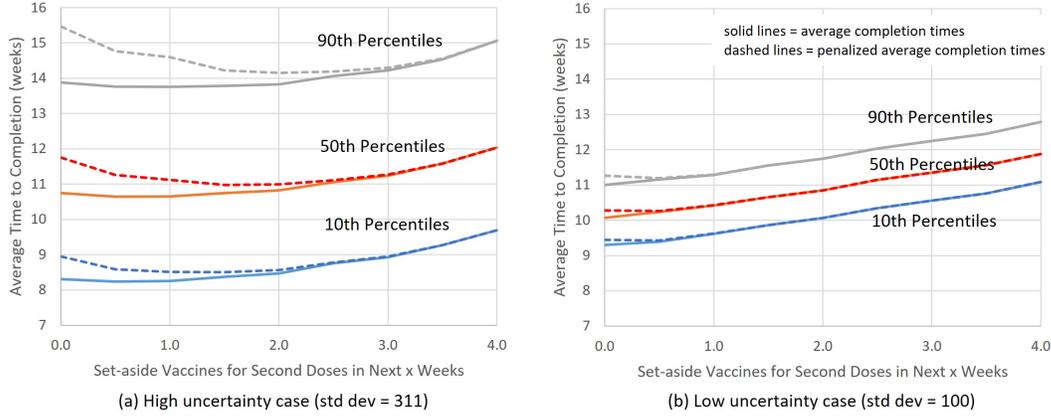


Figure 6: Percentiles for the average completion times

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 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.

4.3 Uncertainty in Average Completion Times

It is important to note that the results in Figures 2-5 focus on *expected* average completion times and expected delays, averaging across possible supply scenarios. Figure 6 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 6, 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 8 to 16 weeks, depending on the policies and objective. The ranges for the low uncertainty case are naturally much

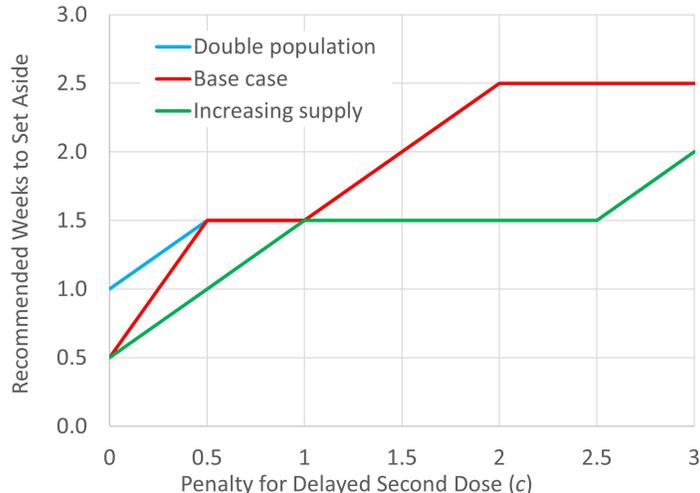


Figure 7: Optimal set-asides for varying penalties (c) for alternative assumptions

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 3,000 people and the supply distribution has a rectified normal distribution with mean 361 and standard deviations of 311 or 100 (both after rectification), matching the observed variability of doses supplied to CMC and to U.S. states. 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 supply allocations to vaccine centers tend to be scaled with the size of the target population, so our results have relevance over a broad range of population sizes.

We also experimented with alternative assumptions that go beyond simply changing the scale of the problem. First, we ran simulations with 6,000 people (instead of 3,000 people) holding the supply distribution constant. Although the average time to completion approximately doubles (as one would expect), the tradeoffs and recommendations are similar to the base case; see Figure 7. The only difference in recommendations is when $c = 0$ where the recommended set-aside is 1.0 week with the larger population rather than 0.5 weeks in the base case. As shown in Figure 2, the average completion times are quite similar for these two set-asides in the base case, so the change in recommended policy does not correspond to a substantial change in performance.

We also ran simulations assuming that the vaccine supply is increasing over time. In this case, we assumed the average supply grows by 10 doses each week, starting from our base case mean of 361; the standard deviation of the underlying normal distribution (before rectification) was held constant. A typical maximum time to completion in our baseline model is about 20 weeks, so this assumption corresponds to growth of $20 \times 10/361 = 55\%$ in the vaccine supply over this time period. As shown in Figure 7, with $c = 1$, setting aside 1.5 weeks is again optimal, though somewhat smaller set-asides are optimal for other penalties. This makes sense because the growing supply can provide second doses with less inventory set aside. One could also experiment with set-aside fractions that are changing over time in response to changes in the supply distribution.

5. Conclusion

Our model demonstrates how the optimal number of doses to hold in inventory depends on the relative cost of delayed second doses and the uncertainty in vaccine supplies, as well as the trajectory of the expected supplies over time. For CMC and similar vaccination centers (e.g., differing in scale), we recommend calculating the number of second doses due in the next one to two weeks and holding that number in inventory. With less uncertainty in the supply, as seen in the state-level data, a set-aside of zero to one week may be preferred. In both cases, a lockbox policy that ensures no second doses will be delayed appears to be overly conservative. The model can easily be adjusted to generate recommendations given other parameters. While vaccine production is increasing, we expect variability and uncertainty will remain and the tradeoffs and issues studied here will remain important as we fight the COVID-19 pandemic. Similar optimization problems may also arise in other settings requiring multiple doses of vaccines (or courses of treatment) with limited and 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 ($0 \leq t \leq T$)

A_t = 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.

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 $1 \leq t \leq T$:

$$I_t = I_{t-1} + A_t - v_{1,t-1} - v_{2,t-1}$$

$$U_{1t} = U_{1,t-1} - v_{1,t-1}$$

$$U_{2t} = U_{2,t-1} - v_{2,t-1}$$

$$d_t = \begin{cases} 0 & \text{for } t = 0, \dots, L-1 \\ d_{t-1} + v_{1,t-L} - v_{2,t-1} & \text{otherwise} \end{cases}$$

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 $c \geq 0$:

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

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, \dots, w_L) , the second and first doses are given by:

$$v_{2t} = \min(d_t, I_t)$$

$$v_{1t} = \min\left(U_{1t}, (1 - 0.5w_L) \left(I_t - v_{2t} - \sum_{j=1}^{L-1} w_j v_{1,t-(L-j)}\right)^+\right).$$

Optimization with Perfect Information: Given perfect information about vaccine supplies A_t , using the definitions above, the vaccine allocation problem can be formulated as a linear program with the objective given above and decision variables v_{1t}, v_{2t} ($0 \leq t \leq T$) representing the number of first and second doses in each week. The constraints are, for each week:

$$\begin{aligned} v_{1t} &\leq U_{1t} && \text{(first doses are limited by unvaccinated people)} \\ v_{2t} &\leq d_t && \text{(second doses are limited by second doses due)} \\ v_{1t} + v_{2t} &\leq I_t && \text{(total doses administered are limited by available inventory)} \\ v_{1t}, v_{2t} &\geq 0 && \end{aligned}$$

A.2 Stochastic Dynamic Program Formulation

Because the stochastic 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, $i = 1, \dots, \ell$ (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 (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

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

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 $\hat{f}(0, \dots, 0, 0, I) = 0$. 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 A.1 below.)

We can interpret (A.1) as follows. The period cost function is the total number of people unvaccinated ($d_0 + \dots + d_\ell + U_1 - v_2$) 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).

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

$$f(x') + f(x'') \geq f((x' + \omega \mathbf{e}) \wedge x'') + f(x' \vee (x'' - \omega \mathbf{e})) . \quad (\text{A.2})$$

where \mathbf{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 (A.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 A.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 $k = 1, 2, \dots$ and $\lim_{k \rightarrow \infty} f_k = f$, 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 $f : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}$ is L^\natural -convex, then $g : \mathbb{R}^n \times \mathbb{R} \rightarrow \bar{\mathbb{R}}$ defined as $g(x, \lambda) = f(x - \lambda e)$ 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 $f(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \bar{\mathbb{R}}$ is a L^\natural -convex function. Let e and \hat{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 $\omega \geq 0$,

$$\arg \min_{\{y:(x,y) \in \mathcal{A}\}} f(x + \omega e, y) \leq \omega \hat{e} + \arg \min_{\{y:(x,y) \in \mathcal{A}\}} f(x, y) .$$

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 $v_{ij} \in \mathbb{R}$.

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

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

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

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

$$f(s_0, \dots, s_\ell, U_2, I) = \min_{\{v_1 : 0 \leq v_1 \leq I, v_1 \leq U_2 - s_\ell\}} \phi(s_0, \dots, s_\ell, s_\ell + v_1, U_2, I - v_1) \quad (\text{A.3})$$

where

$$\begin{aligned} \phi(s_0, \dots, s_\ell, s_{\ell+1}, U_2, I) = & \min_{\{v_2 : 0 \leq v_2 \leq I, v_2 \leq s_0\}} \left\{ (U_2 - v_2) + c(s_0 - v_2) + \right. \\ & \left. + \mathbb{E} \left[f(s_1 - v_2, \dots, s_\ell - v_2, s_{\ell+1} - v_2, U_2 - v_2, I - v_2 + \tilde{A}) \right] \right\}. \end{aligned} \quad (\text{A.4})$$

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 $f(0, \dots, 0, 0, I) = 0$, as before.

Let $v_1^*(s_0, \dots, s_\ell, U_2, I)$ and $v_2^*(s_0, \dots, s_{\ell+1}, U_2, I)$ denote optimal first and second doses for (A.3) and (A.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 $R^*(s_0, \dots, s_\ell, U_2, I) = I - v_1^*(s_0, \dots, s_\ell, U_2, I)$ 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 A.1 (Structural Properties)

- (a) $f(s_0, \dots, s_\ell, U_2, I)$ and $\phi(s_0, \dots, s_\ell, s_{\ell+1}, U_2, I)$ are decreasing in I and L^\natural -convex.
- (b) For any $\omega \geq 0$, $f(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega) = f(s_0, \dots, s_\ell, U_2, I)$
- (c) (i) $v_2^*(s_0, \dots, s_\ell, U_2, I) = s_0 \wedge I$
(ii) $R^*(s_0, \dots, s_\ell, U_2, I)$ is increasing in its arguments and, for any $\omega \geq 0$,

$$R^*(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega) \leq \omega + R^*(s_0, \dots, s_\ell, U_2, I). \quad (\text{A.5})$$

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). 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 A.1.

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

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

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

Corollary A.2 (Set-aside Policies)

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

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

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

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

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

The interpretation of this result is as discussed in §3. (a) 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 set-aside bounded by the increase in the amount due (ω) and decreasing in the time until due (i). The optimal set-aside 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. This result thus provides the justification for focusing on heuristic policies that have a similar set-aside structure (as discussed in §3.1).

A.3 Proof of Theorem A.1 (Online Supplement)

Proof of Theorem A.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 (A.3)-(A.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 A.1(b) to conclude that the limit as $T \rightarrow \infty$ 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 $T \rightarrow \infty$ 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) 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 (A.3) and (A.4) (thus larger I s lead to lower continuation values) and increasing I increases the size of the feasible sets in (A.3) and (A.4) (thus larger I s lead to smaller minimum values).

We now show that $\phi_t(s_0, \dots, s_{\ell+1}, U_2, I)$ 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 ,

$$\mathbb{E} \left[f_{t+1}(s_1 - v_2, \dots, s_{\ell+1} - v_2, U_2 - v_2, I - v_2 + \tilde{A}) \right],$$

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

Let

$$(s'_0, \dots, s'_\ell, U'_2, I') \quad \text{and} \quad (s''_0, \dots, s''_\ell, U''_2, I'') \tag{A.6}$$

denote two sets of state variables and let

$$v'_1 = v_1^*(s'_0, \dots, s'_\ell, U'_2, I') \quad \text{and} \quad v''_1 = v_1^*(s''_0, \dots, s''_\ell, U''_2, I'') \tag{A.7}$$

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

$$\begin{aligned} s_i^\wedge &= (s'_i + \omega) \wedge s''_i & s_i^\vee &= s'_i \vee (s''_i - \omega) \\ U^\wedge &= (U' + \omega) \wedge U'' & U^\vee &= U' \vee (U'' - \omega) \end{aligned} \tag{A.8}$$

and the corresponding first doses and inventory levels as

$$\begin{aligned} v_1^\wedge &= (s'_\ell + v'_1 + \omega) \wedge (s''_\ell + v''_1) - s_\ell^\wedge & v_1^\vee &= (s'_\ell + v'_1) \vee (s''_\ell + v''_1 - \omega) - s_\ell^\vee \\ I^\wedge &= (I' - v'_1 + \omega) \wedge (I'' - v''_1) + v_1^\wedge & I^\vee &= (I' - v'_1) \vee (I'' - v''_1 - \omega) + v_1^\vee. \end{aligned} \tag{A.9}$$

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

³For 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^{\sharp} -convexity assumption below. Then we have

$$\begin{aligned}
& f_t(s'_0, \dots, s'_\ell, U'_2, I') + f_t(s''_0, \dots, s''_\ell, U'_2, I') \\
&= \phi_t(s'_0, \dots, s'_\ell, s'_\ell + v'_1, U'_2, I' - v'_1) + \phi_t(s''_0, \dots, s''_\ell, s''_\ell + v'_1, U'_2, I' - v'_1) \\
&\geq \phi_t(s^\wedge_0, \dots, s^\wedge_\ell, s^\wedge_\ell + v_1^\wedge, U_2^\wedge, I^\wedge - v_1^\wedge) + \phi_t(s^\vee_0, \dots, s^\vee_\ell, s^\vee_\ell + v_1^\vee, U_2^\vee, I^\vee - v_1^\vee) \\
&\geq f_t(s^\wedge_0, \dots, s^\wedge_\ell, U_2^\wedge, I^\wedge) + f_t(s^\vee_0, \dots, s^\vee_\ell, U_2^\vee, I^\vee).
\end{aligned}$$

The first equality follows from the definition of v'_1 and v''_1 as optimal values in (A.3) for the given states. The next inequality follows from the L^{\sharp} -convexity of ϕ_t . The final inequality follows from the observation that the first dose choices v_1^\wedge and v_1^\vee 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^{\sharp} -convex, thereby completing the proof of part (a).

(b) We first show

$$f(s_0, \dots, s_\ell, U_2, I) \geq f(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega)$$

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

$$\begin{aligned}
& f(s_0, \dots, s_\ell, U_2, I) \\
&= (U_2 - v'_2) + c(s_0 - v'_2) + \mathbb{E} \left[f(s_1 - v'_2, \dots, s_\ell - v'_2, s_\ell + v'_1 - v'_2, U_2 - v'_2, I - v'_1 - v'_2 + \tilde{A}) \right] \\
&= (U_2 - v''_2 + \omega) + c(s_0 - v''_2 + \omega) + \\
&\quad \mathbb{E} \left[f(s_1 - v''_2 + \omega, \dots, s_\ell - v''_2 + \omega, s_\ell + v'_1 - v_2 + \omega, U_2 - v''_2 + \omega, I - v'_1 - v''_2 + \omega + \tilde{A}) \right] \\
&\geq f(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega)
\end{aligned}$$

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 and v''_2 are feasible given $(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega)$ but not necessarily optimal.

The reverse inequality is established similarly. Let v'_1 and v''_2 be optimal given $(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega)$ and let $v'_2 = v''_2 - \omega$. Then

$$\begin{aligned}
& f(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega) \\
&= (U_2 - v''_2 + \omega) + c(s_0 - v''_2 + \omega) + \\
&\quad \mathbb{E} \left[f(s_1 - v''_2 + \omega, \dots, s_\ell - v''_2 + \omega, s_\ell + v'_1 - v_2 + \omega, U_2 - v''_2 + \omega, I - v'_1 - v''_2 + \omega + \tilde{A}) \right] \\
&= (U_2 - v'_2) + c(s_0 - v'_2) + \mathbb{E} \left[f(s_1 - v'_2, \dots, s_\ell - v'_2, s_\ell + v'_1 - v'_2, U_2 - v'_2, I - v'_1 - v'_2 + \tilde{A}) \right] \\
&\geq f(s_0, \dots, s_\ell, U_2, I)
\end{aligned}$$

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 and v'_2 are feasible given $(s_0, \dots, s_\ell, U_2, I)$ but not necessarily optimal. With both inequalities established, we have the desired result.

(c)(i) follows from part (b). If $s_0 \leq I$, then (b) implies

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

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 $s_0 \geq I$, then (b) implies

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

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

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

$$r^*(s_0, \dots, s_\ell, s_{\ell+1}, U_2, I) = \arg \min_{\{R : 0 \leq R \leq I\}} \phi(s_0, \dots, s_\ell, s_{\ell+1}, U_2, R). \quad (\text{A.10})$$

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 A.1(c), we have the following lemma.

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

$$r^*(s_0 + \omega, \dots, s_\ell + \omega, s_{\ell+1} + \omega, U_2 + \omega, I + \omega) \leq \omega + r^*(s_0, \dots, s_\ell, s_{\ell+1}, U_2, I). \quad (\text{A.11})$$

Proof of Lemma 1. L^\sharp -convexity of ϕ was established in Theorem A.1(a) and the feasible set in (A.10) is L^\sharp -convex by Proposition A.1(f). The result then follows from Proposition A.1(e)(i) and (ii). ■

Now consider ordered state variables

$$(s'_0, \dots, s'_\ell, s'_{\ell+1}, U'_2, I') \leq (s''_0, \dots, s''_\ell, s''_{\ell+1}, U''_2, I'')$$

and define

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

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 8.

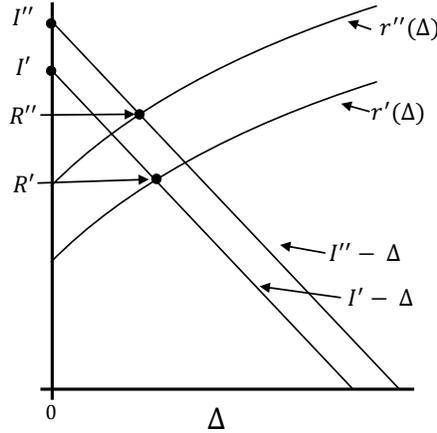


Figure 8: Illustration for proof of Theorem A.1(c)

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, $R' = R^*(s'_0, \dots, s'_\ell, U'_2, I')$, 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 8) 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 (A.5) of Theorem (c)(ii) define

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

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 8. The lemma also says that $r''(\Delta) \leq r'(\Delta) + \omega$. The optimal values $R' = R^*(s_0, \dots, s_\ell, U_2, I)$ and $R'' = R^*(s_0 + \omega, \dots, s_\ell + \omega, U_2 + \omega, I + \omega)$ 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

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

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 A.2. (a) follows immediately from Theorem A.1(c)(i) given that $s_0 = d_0$.

(b) follows from (A.5). For (i) and (ii), first note that the changes contemplated do not affect the number of first doses given, so the properties of differences in S^* reflect the difference in R^* . The result of (A.5) says that increasing all arguments of R^* by ω (as on the left side of (A.5)) 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 $(s_i, \dots, s_\ell, U_2)$ to $(s_i + \omega, \dots, s_\ell + \omega, U_2 + \omega)$, 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. ■

⁴Note 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.

References

- Amit, S., et al. 2021. Early rate reductions of SARS-CoV-2 infection and COVID-19 in BNT162b2 vaccine recipients. *The Lancet* URL [https://doi.org/10.1016/S0140-6736\(21\)00448-7](https://doi.org/10.1016/S0140-6736(21)00448-7).
- 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>.
- CDC. 2021a. Information about the Pfizer-VioNTech COVID-19 vaccine URL <https://www.cdc.gov/coronavirus/2019-ncov/vaccines/different-vaccines/Pfizer-BioNTech.html>.
- CDC. 2021b. Information about the Moderna COVID-19 vaccine URL <https://www.cdc.gov/coronavirus/2019-ncov/vaccines/different-vaccines/Moderna.html>.
- CDC. 2021c. Interim clinical considerations for use of mRNA COVID-19 vaccines currently authorized in the united states URL <https://www.cdc.gov/vaccines/covid-19/info-by-product/clinical-considerations.html>.
- CDC. 2021d. State moderna allocation totals URL <https://data.cdc.gov/Vaccinations/State-Moderna-Allocation-Totals/spzz-he63>.
- CDC. 2021e. State pfizer allocation totals URL <https://data.cdc.gov/Vaccinations/State-Pfizer-Allocation-Totals/y77q-5ru3>.
- 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/>.
- 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>.
- Goodin-Smith, O. 2021. Coronavirus: 15toKnow KoP Mall drive-through vaccine site scrambles for second doses from the state. *Philadelphia Inquirer* URL <https://www.inquirer.com/health/coronavirus/15toknow-vaccine-king-of-prussia-mall-coronavirus-second-dose-pennsylvania-20210310.html>.
- 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/>.
- 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/NEJMc1de2101987>.

- 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.
- 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>.
- 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-20210201-1sx5bh6rjfh7ra3rhpjh247kxi-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.
- 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.
- Rowland, C. 2021. Biden wants to squeeze an extra shot of vaccine out of every Pfizer vial. it won't be easy. *The Washington Post* URL <https://www.washingtonpost.com/business/2021/01/22/pfizer-vaccine-doses-syringes/>.
- Smith, J. E., K.F. McCardle. 2002. Structural properties of stochastic dynamic programs. *Operations Research* **50**(5) 796–809.
- 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.