

Where to locate COVID-19 mass vaccination facilities?

Dimitris Bertsimas | Vassilis Digalakis Jr | Alexander Jacquillat | Michael Lingzhi Li |

Alessandro Previero

Sloan School of Management and Operations
Research Center, Massachusetts Institute of
Technology, Cambridge, Massachusetts, USA

Correspondence

D. Bertsimas, Sloan School of Management and
Operations Research Center, Massachusetts
Institute of Technology, Cambridge, MA, USA.
Email: dbertsim@mit.edu

[Correction added on 28-July-2021, after first
online publication: In equation 22, β was changed
to $(1-\beta)$.]

History

Accepted by Sanjay Mehrotra, healthcare
management.

Abstract

The outbreak of COVID-19 led to a record-breaking race to develop a vaccine. However, the limited vaccine capacity creates another massive challenge: how to distribute vaccines to mitigate the near-end impact of the pandemic? In the United States in particular, the new Biden administration is launching mass vaccination sites across the country, raising the obvious question of where to locate these clinics to maximize the effectiveness of the vaccination campaign. This paper tackles this question with a novel data-driven approach to optimize COVID-19 vaccine distribution. We first augment a state-of-the-art epidemiological model, called DELPHI, to capture the effects of vaccinations and the variability in mortality rates across age groups. We then integrate this predictive model into a prescriptive model to optimize the location of vaccination sites and subsequent vaccine allocation. The model is formulated as a bilinear, nonconvex optimization model. To solve it, we propose a coordinate descent algorithm that iterates between optimizing vaccine distribution and simulating the dynamics of the pandemic. As compared to benchmarks based on demographic and epidemiological information, the proposed optimization approach increases the effectiveness of the vaccination campaign by an estimated 20%, saving an extra 4000 extra lives in the United States over a 3-month period. The proposed solution achieves critical fairness objectives—by reducing the death toll of the pandemic in several states without hurting others—and is highly robust to uncertainties and forecast errors—by achieving similar benefits under a vast range of perturbations.

KEYWORDS

COVID-19, epidemiological modeling, nonconvex optimization,
vaccine distribution

1 | INTRODUCTION

The outbreak of the COVID-19 pandemic has started a global race to develop vaccines, fueled by extensive investments, governmental support, and scientific breakthroughs. Thanks to these unprecedented efforts, the scientific community delivered the good news that the whole world was eagerly awaiting. By Summer 2020, several vaccines had been developed. By the end of 2020, several vaccines got approved for emergency use and hundreds more were going under development and testing. Whereas vaccine development used to take years and even decades, these results rank, with no doubt,

among the greatest scientific achievements (Graham, 2020; Lurie et al., 2020).

Unfortunately, discovering and developing a vaccine for COVID-19 was just the beginning—it will now take months to produce, distribute, and deliver vaccines at scale. The world has quickly come to the realization that vaccines cannot be made available immediately to everyone, and policy makers need to make tough decisions to pilot vaccine distribution. A global consensus has naturally emerged to prioritize to healthcare workers, other front line workers, and vulnerable populations such as older people and people with comorbidities (see, e.g., National Academies of Sciences,

Engineering, and Medicine, 2020). Within these general principles, each jurisdiction is designing more detailed eligibility guidelines to distribute vaccines effectively and equitably *within a population*, based on demographic, clinical and geographic factors. However, a question remains open: how to plan vaccine distribution *across populations*, that is, how to allocate a limited vaccine supply across communities, across provinces, and even across countries?

In the United States, this question gained prominence in the midst of a presidential transition. In particular, the new Biden administration relies on higher extents of federal coordination in vaccine distribution, as opposed to a more decentralized approach at the state level. In one of its first major decisions, the administration started opening mass vaccination sites, with many more planned over the next few weeks.¹ This environment raises the critical question of where to locate these vaccination sites. Obviously, these decisions need to adhere to a number of political and fairness considerations—most notably, there must be at least one site per state. Yet, there remains flexibility to use mass vaccination sites as a strategic lever to effectively combat the pandemic.

This paper addresses this question with a novel data-driven approach, combining epidemiological modeling and prescriptive analytics, to optimize the location of vaccination sites and the subsequent allocation of vaccines. To this end, we leverage a recent compartmental epidemiological model called DELPHI (Differential Equations Lead to Predictions of Hospitalizations and Infections), which extends Susceptible-Exposed-Infected-Recovered (SEIR) models to capture critical drivers of the COVID-19 pandemic: (i) under-detection due to limited testing, (ii) governmental and societal response, and (iii) declining mortality rates (Li et al., 2020). The DELPHI model has been fitted from historical data at the country level, at the state level in the United States, and at the province level in a few other countries. The DELPHI forecasts have been incorporated into the ensemble forecast from the U.S. Center for Disease Control (2020a) and have been utilized in selecting the Phase III trial locations for the Johnson and Johnson COVID-19 vaccine. Historically, the DELPHI model has featured excellent predictive performance, matching the number of detected cases and deaths with high accuracy across the various waves of the pandemic.

In this paper, we integrate the (predictive) DELPHI model into a (prescriptive) optimization model for vaccine allocation. We first propose an extension of DELPHI, referred to as DELPHI-V, to capture the effects of vaccinations on the dynamics of the pandemic. The DELPHI-V model also disaggregates the dynamics of the pandemic at the subpopulation level to reflect disparities in mortality rates across age groups, which are critical drivers of vaccination strategies. We then formulate an optimization model, referred to as DELPHI-V-OPT, which optimizes the vaccine distribution

strategy (e.g., the deployment of mass vaccination sites at the strategic level, and the subsequent allocation of vaccines at the tactical level) to minimize the death toll of the pandemic. Our focus on mass vaccination centers does not hinder the role that smaller vaccination sites (e.g., pharmacies) have been playing throughout the country to vaccinate the population. Ideally, our modeling approach would consider these various sites jointly. However, given the lack of publicly available information on the vaccines administered in smaller sites and the lack of coordination between the various vaccination sites, we leave this integration for future research.

From a technical standpoint, the DELPHI-V-OPT model relies on time discretization to embed the system of ordinary differential equations governing the DELPHI-V dynamics into an optimization model. The model is formulated as a bilinear (nonconvex) optimization model, due to the SEIR dynamics at the core of DELPHI-V in which the number of new cases is driven by the number of susceptible and infected people. To solve it efficiently in realistic large-scale settings, we propose a coordinate descent algorithm. Starting from a baseline solution, the algorithm iterates, until convergence, between optimizing the vaccine distribution strategy (for given dynamics of the pandemic) and simulating the dynamics of the pandemic (for a given vaccine distribution strategy).

We implement the proposed model and algorithm using real-world data in the United States from the New York Times (2020), the U.S. Census Bureau (2020), and the U.S. Center for Disease Control (2021). We leverage the parameter estimates from the DELPHI model in each U.S. state. One challenge, however, is that DELPHI estimates mortality rates in each state in each time period, while the U.S. Center for Disease Control (2021) reports mortality rates in each age bracket. To develop realistic and consistent estimates for mortality rates in each state, each age group and each time period, we formulate another bilinear optimization model that interpolates these two pieces of information, while ensuring consistency with broader demographic information.

Results suggest that the locations of vaccination sites can have a massive impact on the effectiveness of the vaccination campaign. As compared to several benchmarks based on demographic information (e.g., city and state population) and epidemiological information (e.g., case counts), our optimization approach increases the number of lives saved by the vaccines by 20%, or 4000 lives over a three-month period in the United States. These results underscore the necessity to consider both demographics and epidemiological dynamics when determining the locations of vaccination sites and subsequent vaccine allocation, which is achieved by the combination of our DELPHI-V epidemiological model and our optimization framework. In addition, the optimization approach can ensure equity between states and across vaccination sites, thus alleviating the death toll of the pandemic in some states without hurting others. Finally, these benefits are highly robust to misspecifications and fluctuations in the DELPHI parameters. Practically speaking, even though

¹ www.nbcnews.com/politics/white-house/federal-government-opening-first-mass-covid-19-vaccination-sites-california-n1256611

tactical decisions (e.g., vaccine allocation) need to be revised continuously in response to the latest information available throughout the vaccination campaign, strategic decisions (i.e., the location of vaccination sites) are highly robust to noise and uncertainty.

In summary, this paper makes three contributions. From a modeling standpoint, it formulates a novel optimization model for vaccine allocation, DELPHI-V-OPT, that integrates a state-of-the-art epidemiological model into an optimization model that supports vaccine distribution strategies, in order to mitigate the impact of the pandemic. From a computational standpoint, it develops a scalable coordinate descent algorithm, which converges effectively and in short runtimes. From a practical standpoint, it demonstrates that optimizing the locations of mass vaccination sites can curb the death toll of COVID-19 by a sizeable amount, thus highlighting the critical role of vaccine distribution besides vaccine design and vaccine production in combating the pandemic. Obviously, vaccine distribution involves broad political, economic and social considerations, which lie beyond the scope of this paper; yet, this paper can play a critical role to support ongoing mass vaccination efforts in order to mitigate the impact of the pandemic on public health.

2 | LITERATURE REVIEW

Many pharmaceutical companies and academic institutions have explored different technologies toward a SARS-CoV-2 vaccine (Florindo et al., 2020; Shin et al., 2020). These span (i) inactivated or live-attenuated virus vaccines, which induce an immune response from weakened or killed pathogens (used by the Wuhan Institute of Biological Products, for instance); (ii) viral vector vaccines, which exploit nonreplicating adenoviruses to deliver an antigenic element (used by Johnson and Johnson, for instance); (iii) subunit vaccines, which use a minimal structural component of a pathogen such as a protein (used by Clover Biopharmaceuticals, for instance); (iv) nucleic acid vaccines, which deliver DNA or mRNA of viral proteins (used by Pfizer and Moderna, for instance).

From an operational standpoint, a vast literature studies vaccine supply chains (see Duijzer et al., 2018; Lemmens et al., 2016). A first area involves optimizing vaccine composition (Bandi & Bertsimas, 2020; Cho, 2010; Kornish & Keeney, 2008; Wu et al., 2005). A second area focuses on vaccine production to manage supply-side and demand-side uncertainty and mitigate incentive misalignments between manufacturers and end users (Arifoğlu et al., 2012; Chick et al., 2008; Federgruen & Yang, 2009). Next, vaccine allocation optimizes the management of a vaccine stockpile (Mamani et al., 2013; Sun et al., 2009). Last, vaccine delivery optimizes inventory, distribution and dispensing operations (Aaby et al., 2006; Dai et al., 2016; Jacobson, Sewell, & Proano, 2006). Most of this research focuses on predictable and repeatable epidemics, such as seasonal influenza. For less

predictable epidemics, such as pandemic influenza, advance planning interventions include stockpiling (Jacobson, Sewell, Proano, & Jokela, 2006) and anticipatory vaccination (Ari-naminpathy et al., 2012). Unfortunately, these approaches are not readily applicable to a new disease such as COVID-19.

Our paper deals with centralized vaccine allocation within a population. Early studies established the importance of partitioning the population into risk classes (e.g., age groups) to reflect the impact of an epidemic (Elveback et al., 1976; Longini Jr et al., 1978; Watson, 1972). Emanuel and Wertheimer (2006) propose a life-cycle model that prioritizes the most valuable subpopulations. Within a region, results suggest prioritizing at-risk populations (Chowell et al., 2009; Patel et al., 2005) or active agents who can spread the disease fastest, such as school children (Basta et al., 2009; Dushoff et al., 2007; Lee et al., 2012; Matrajt et al., 2013; Medlock & Galvani, 2009). Across regions, results suggest that vaccines should be allocated to the most infected regions and to those affected the latest by the epidemic (Araz et al., 2012; Keeling & Shattock, 2012).

Methodologically, most studies integrate SEIR or similar epidemiological models into simple optimization routines based on scenario analysis, enumeration, simulation, or simple heuristics (Teytelman & Larson, 2013; Uribe-Sánchez et al., 2011). Tanner et al. (2008) propose a chance-constrained optimization approach to ensure that the post-vaccination reproduction number is lower than one with high probability. Yarmand et al. (2014) formulate a two-stage stochastic programming model to first plan vaccine allocation and then distribute additional doses where the epidemic has not been contained. They model the dynamics of disease propagation by means of a stochastic SEIR model, and define scenarios using Monte Carlo simulation. In contrast, this paper directly embeds SEIR dynamics into an optimization model to support vaccine distribution.

Finally, this paper contributes to the fast-growing field of vaccine distribution in the midst of the COVID-19 pandemic. Recent and ongoing research spans vaccine production (Khamsi, 2020), equity in vaccine distribution (Bae et al., 2020; Muriel & Bauchner, 2021), and public acceptance (Coustasse et al., 2021; Dror et al., 2020). In terms of vaccine distribution, Rastegar et al. (2021) propose a mixed-integer formulation to support influenza vaccine distribution during the COVID-19 pandemic. Matrajt et al. (2020) study which populations to prioritize in a mass vaccination campaign, trading off vaccinating high-risk (older) age-groups versus high-transmission (younger) age-groups in a given location. In contrast, this paper optimizes the distribution of vaccines across locations. This relates to Grauer et al. (2020), who study the spatiotemporal distribution of vaccines, using an SEIR model to test various strategies based on demographic and epidemiological factors.

This paper expands this recent body of work in three major ways. First, we optimize vaccine allocation across regions and risk classes (e.g., age groups), based on data-driven estimates

of infection and mortality rates. Second, we leverage a recent SEIR-inspired epidemiological model that captures dynamics specific to the COVID-19 pandemic, such as under-detection, governmental response, and declining mortality rates. Third, we propose a formal optimization approach and a coordinate descent algorithm to explicitly optimize vaccine distribution strategies, as opposed to relying on enumeration, simulation or simplified heuristics.

3 | MODEL FORMULATION

Our model optimizes vaccine distribution strategy. In the U.S. context, this primarily involves the location of mass vaccination sites. However, optimizing these decisions requires to account for subsequent vaccine allocation across the population, in order to further optimize and evaluate the effects of the vaccination campaign. Therefore, we refer to as *vaccine distribution strategy* the set of three decisions: (i) the location of mass vaccination sites, (ii) the allocation of vaccines across vaccination sites, and (iii) the allocation of vaccines within each subpopulation.

We capture the dynamics of the pandemic by means of an epidemiological model, called DELPHI, which forecasts the number of detected cases, hospitalizations and deaths in each U.S. state (Li et al., 2020).² We review it briefly, and augment it to capture the effects of vaccinations—we refer to this model as DELPHI-V. We then embed the DELPHI-V model into a mathematical programming model to optimize vaccine allocation, referred to as DELPHI-V-OPT.

3.1 | DELPHI: Forecasting the dynamics of the COVID-19 pandemic

DELPHI is a compartmental epidemiological model, which extends the widely used SEIR model to account for specificities of the COVID-19 pandemic. The model is governed by a system of ordinary differential equations (ODEs) across 11 states: susceptible (S), exposed (E), infectious (I), undetected cases who will recover (U^R) or die (U^D), hospitalized cases who will recover (H^R) or die (H^D), quarantined cases who will recover (Q^R) or die (Q^D), recovered (R), and dead (D).

DELPHI differs from most other COVID-19 forecasting models (see, e.g., Kissler et al., 2020, Perkins & Espana, 2020, Rodriguez et al., 2020) by capturing three key elements of the pandemic:

- *Under-detection*: Many cases remain undetected due to limited testing, asymptomatic carriers, and detection errors. Ignoring them would underestimate the scale of the pandemic. The DELPHI model captures them through the U^R and U^D states.

- *Governmental and societal response*: Social distancing policies limit the spread of the virus. Ignoring them would overestimate the scale of the pandemic. However, if restrictions are lifted prematurely, a resurgence may occur. We define a governmental and societal response function $\gamma(t)$, which modulates the infection rate and is parameterized as follows:

$$\gamma(t) = 1 + \frac{2}{\pi} \arctan\left(\frac{-(t - t_{\text{int}})}{\omega}\right) + c \exp\left(-\frac{(t - t_{\text{jump}})^2}{2\sigma^2}\right). \quad (1)$$

This parameterization defines four phases (Figure 1). In Phase I, most activities continue normally. In Phase II, the infection rate declines sharply as policies get implemented. The parameters t_{int} and ω can be interpreted as the start time and strength of this response. In Phase III, the decline reaches saturation. The epidemic then experiences a resurgence of magnitude c in Phase IV, due to relaxations in governmental and social restrictions. This is counteracted at time t_{jump} , when restrictions are re-implemented, with σ controlling the duration of this second wave.

- *Declining mortality rates*: The mortality rate of COVID-19 has been declining through the pandemic, due to a better detection of mild cases, enhanced care for COVID-19 patients, and other factors. We model the mortality rate as a monotonically decreasing function of time:

$$m(t) = (m_0 - m_{\text{min}}) \left(1 + \frac{2}{\pi} \arctan(-r_m t)\right) + m_{\text{min}}, \quad (2)$$

where m_0 is the initial mortality rate, m_{min} is the minimum mortality rate and r_m is a decay rate.

Ultimately, DELPHI involves 16 parameters that define the transition rates between the 11 states. We calibrate seven of them from a database on clinical outcomes (Bertsimas et al., 2020). Using nonlinear optimization, we estimate the other nine parameters from historical data on the number of cases and deaths in each region. We refer to Li et al. (2020) for details.

Since its inception in March 2020, DELPHI has been extensively tested and validated against real-world data. Figure 2 reports the historical performance of the model in the United States, during the first wave in the Spring of 2020 and the second wave in the Fall of 2020. As the figure shows, the model has been predicting the magnitude of the pandemic with high accuracy up to 1 month in advance; for instance, as early as April 3, 2020, the model was predicting 1.2–1.4 million cases in the United States by early May, a prediction that became quite accurate a month later (Figure 2a). Obviously, subsequent forecasts, by leveraging more up-to-date information, were able to refine these estimates. As a result, the DELPHI model was incorporated into the ensemble forecast from the U.S. Center for Disease Control (2020a). During the second wave of the pandemic, DELPHI continued to exhibit strong predictive performance, with a mean average percentage error among the lowest of the CDC ensemble forecast (Figure 2b).

²DELPHI is also applied to each country and to other provinces, but this paper focuses on U.S. states.

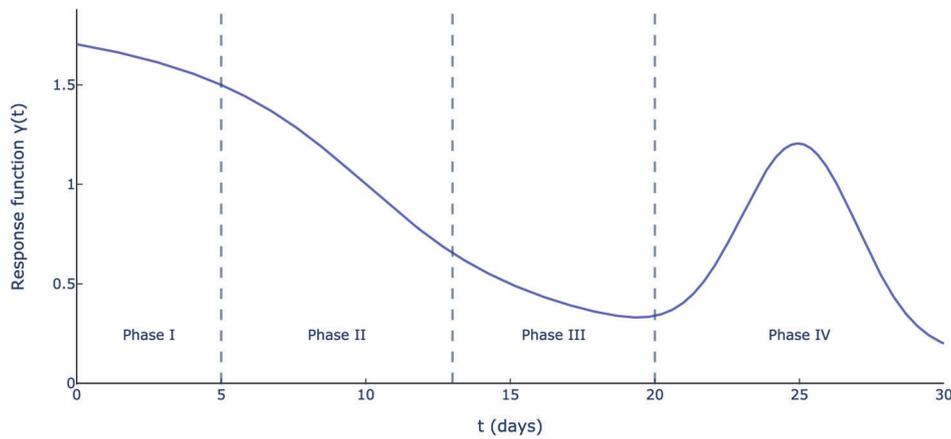


FIGURE 1 Governmental and societal response function $\gamma(t)$ ($\omega = 5$, $t_{\text{int}} = 10$, $c = 1$, $t_{\text{jump}} = 25$, and $\sigma = 2$)

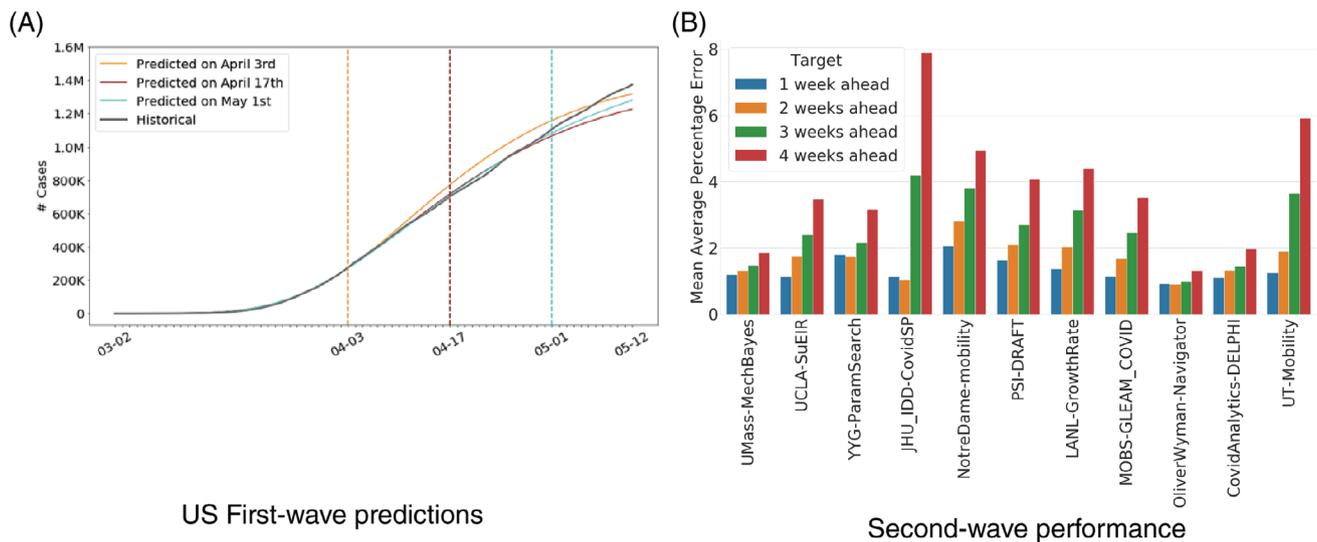


FIGURE 2 Historical performance of the DELPHI predictions in the United States

3.2 | Predictive DELPHI-V: Capturing the effects of vaccination

We now augment the DELPHI model to capture two key aspects of vaccinations:

1. *Disparate impacts of the disease across risk classes.* Age is one of the primary drivers of mortality (Goyal et al., 2020; Petrilli et al., 2020; Wj et al., 2020). The U.S. Center for Disease Control (2020b) reports that the mortality rate among Americans aged 70 and over is two orders of magnitude greater than for those aged 30 and under. We partition the population into *risk classes*, defined as homogeneous groups with comparable health characteristics. We consider age-based risk classes in our experiments, but other categorizations could be used (e.g., based on comorbidities). Accordingly, we replicate the 11 model states for each risk class.
2. *Impact of vaccinations on the dynamics of the pandemic.* A fraction of vaccinated people will be immune to the disease (based on the vaccine's effectiveness). Clinical trials suggest that

early-approved vaccines prevent mortality but not necessarily infections. Therefore, we assume conservatively that all vaccinated people can still transmit the disease. We relax this assumption later on, to show the robustness of our results when a fraction of vaccinated people become fully immune to the disease. We create four new model states: susceptible and vaccinated (S'), exposed and vaccinated (E'), infected and vaccinated (I'), and immune (M).

Figure 3 shows a simplified flow diagram of the DELPHI-V model, with two risk classes (indexed by $k = 1, 2$ and indicated via subscripts). For expositional purposes, we omit dependencies on the region, since the DELPHI-V model is fitted in each region independently. In the remainder of this paper, we also ignore the recovery states, since they do not impact the death-minimization optimization model. Accordingly, we denote the states of undetected, hospitalized and quarantined people who will die from the disease by U , H and Q (as opposed to U^D , H^D and Q^D).

For simplicity, we make three assumptions. First, the effects of vaccines are instantaneous (relaxing this assumption, although straightforward, would merely induce a time lag

model, to minimize (i) the death toll of the pandemic, (ii) the number of exposed people in the termination period, and (iii) the distance between vaccination sites and population centers. The main public health objective is obviously death minimization, so the first objective component is heavily prioritized. However, just considering the number of deaths could result in a waste of vaccines near the end of the planning horizon, as individuals infected in the final periods would not have time to flow to the death state in the epidemiological model. Therefore, the second component of the objective minimizes the number of infections. The last component minimizes geographic disparities. In addition, the model incorporates other equity consideration by means of fairness constraints.

We proceed by time discretization to formulate the optimization model and retain tractability. This reduces to solving the system of ODE equations given in Equations (3)–(13) by a forward difference scheme. We denote by Δt the discretization unit (e.g., 1 day).

Formally, we define the following sets, input parameters, and decision variables.

Sets

\mathcal{L} = set of population centers in the United States
(e.g., counties), $\{1, \dots, L\}$

\mathcal{M} = set of regions in the United States
(e.g., 50 states plus the District of Columbia),
 $\{1, \dots, m\}$

\mathcal{N} = set of candidate vaccination sites, $\{1, \dots, n\}$

\mathcal{N}_j = subset of candidate vaccination sites located in state
 $j \in \mathcal{M}$

\mathcal{N}_l = subset of candidate vaccination sites located
in the same state as county $l \in \mathcal{L}$

\mathcal{K} = set of risk classes in the population (e.g., age groups),
 $\{1, \dots, K\}$

\mathcal{T} = set of discretized time periods (e.g., days), $\{1, \dots, T\}$

Parameters

Pop_l = total population in center $l \in \mathcal{L}$

Δ_{li} = distance from population center $l \in \mathcal{L}$
to candidate vaccination site $i \in \mathcal{N}$

N = number of vaccination sites to be
deployed across the country

B_t = number of vaccines available at time $t \in \mathcal{T}$

β = effectiveness of vaccines

r^I = progression rate of the disease (DELPHI parameter)

r^d = detection rate of the disease (DELPHI parameter)

r^D = death rate of the disease (DELPHI parameter)

r_{jkt}^U = transition rate from I to U , capturing mortality rate
in region $j \in \mathcal{M}$ at time $t \in \mathcal{T}$ for risk class
 $k \in \mathcal{K}$ (DELPHI parameter)

r_{jkt}^H = transition rate from I to H , capturing mortality
rate in region $j \in \mathcal{M}$ at time $t \in \mathcal{T}$ for risk class
 $k \in \mathcal{K}$ (DELPHI parameter)

r_{jkt}^Q = transition rate from I to Q , capturing mortality
rate in region $j \in \mathcal{M}$ at time $t \in \mathcal{T}$ for risk class
 $k \in \mathcal{K}$ (DELPHI parameter)

r_{jkt}^H = transition rate from I to H , capturing mortality rate
 α_j = nominal infection rate in region $j \in \mathcal{M}$
(DELPHI parameter)

γ_{jt} = governmental and societal response in region j
 $\in \mathcal{M}$ at time $t \in \mathcal{T}$ (DELPHI parameter)

Note that the parameters r^U , r^H , and r^Q are defined for each region, risk class and time period, reflecting underlying variations in mortality rates. In contrast, the parameters r^I , r^d , and r^D are treated as uniform characteristics of the disease. In reality, these parameters may vary across risk classes; for instance, the serological estimates from the U.S. Center for Disease Control (2020a) suggest different prevalence of the disease across age groups. We test this hypothesis in our experiments, to verify the robustness of our results to the uniform infection rate assumption.

We also assume a single vaccine effectiveness value β . In theory, vaccine effectiveness might also vary across risk classes. More importantly, there are now several vaccines available, each with different clinical characteristics. Ideally, we could introduce an additional set to capture vaccine heterogeneity, and let the vaccine effectiveness vary across vaccine types. This approach however, may be somewhat impractical in practice, as it may be difficult to strategically allocate different vaccines to different populations based on vaccine effectiveness. For equity, we therefore assume conservatively that the mix of vaccines remains identical across vaccination sites. Under this restriction, the mix of vaccines can be reduced to a representative vaccine with average effectiveness.

Primary decision variables

$x_i = \begin{cases} 1 & \text{if vaccination site } i \in \mathcal{N} \text{ is selected} \\ 0 & \text{otherwise} \end{cases}$

C_{it} : number of vaccines distributed to site
 $i \in \mathcal{N}$ at time $t \in \mathcal{T}$

$W_{li} = \begin{cases} 1 & \text{if population center } l \in \mathcal{L} \text{ is assigned} \\ & \text{to vaccination site } i \in \mathcal{N} \\ 0 & \text{otherwise} \end{cases}$

\bar{S}_{jkt} : number of eligible people to region $j \in \mathcal{M}$
in risk class $k \in \mathcal{K}$ at time $t \in \mathcal{T}$

V_{jkt} : number of vaccines allocated to region $j \in \mathcal{M}$
in risk class $k \in \mathcal{K}$ at time $t \in \mathcal{T}$

To track the impact of vaccine allocation on the resulting dynamics of the pandemic, we create indirect variables,

corresponding to all the states in the DELPHI-V model shown in Figure 3.

The vaccine distribution problem, referred to as (P) , is then formulated in Equations (14)–(39).

$$\min \sum_{j=1}^M \sum_{k=1}^K (D_{jkt} + H_{jkt} + Q_{jkt}) + \lambda_E \sum_{j=1}^M \sum_{k=1}^K E_{jkt} + \lambda_D \sum_{l=1}^L \sum_{i=1}^n Pop_l \Delta_{li} W_{li} \quad (14)$$

$$\text{s.t. } \sum_{i=1}^n x_i = N \quad (15)$$

$$\sum_{i \in \mathcal{N}_j} x_i \geq 1, \quad \forall j \in \mathcal{M} \quad (16)$$

$$W_{li} \leq x_i, \quad \forall l \in \mathcal{L}, i \in \mathcal{N} \quad (17)$$

$$\sum_{i \in \mathcal{N}_l} W_{li} = 1, \quad \forall l \in \mathcal{L} \quad (18)$$

$$\sum_{i=1}^n C_{it} \leq B_t, \quad \forall t \in \mathcal{T} \quad (19)$$

$$C_{it} \leq B_t x_i, \quad \forall i \in \mathcal{N}, t \in \mathcal{T} \quad (20)$$

$$\sum_{k=1}^K V_{jkt} \leq \sum_{i \in \mathcal{N}_j} C_{it}, \quad \forall j \in \mathcal{M}, t \in \mathcal{T} \quad (21)$$

$$\bar{S}_{j,k,t+1} \leq \bar{S}_{jkt} - (1 - \beta)V_{jkt} - (S_{jkt} - S_{j,k,t+1}), \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (22)$$

$$V_{jkt} \leq \bar{S}_{jkt}, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (23)$$

$$|C_{i,t+1} - C_{it}| \leq \theta_S C_{it}, \quad \forall i \in \mathcal{N}, t \in \mathcal{T} \quad (24)$$

$$\frac{\sum_{l \in \mathcal{L}_j} Pop_l}{\sum_{j'=1}^m \sum_{l \in \mathcal{L}_{j'}} Pop_l} N - \Theta_L \leq \sum_{i \in \mathcal{N}_j} x_i \leq \frac{\sum_{l \in \mathcal{L}_j} Pop_l}{\sum_{j'=1}^m \sum_{l \in \mathcal{L}_{j'}} Pop_l} N + \Theta_L, \quad \forall j \in \mathcal{M} \quad (25)$$

$$\frac{B_t}{N(1 + \theta_V)} x_i \leq C_{it} \leq \frac{B_t}{N} (1 + \theta_V) x_i, \quad \forall i \in \mathcal{N}, t \in \mathcal{T} \quad (26)$$

$$\sum_{i \in \mathcal{N}_j} C_{it} \leq \left(\frac{\sum_{l \in \mathcal{L}_j} Pop_l}{\sum_{j'=1}^m \sum_{l \in \mathcal{L}_{j'}} Pop_l} + \theta_P \right) B_t, \quad \forall j \in \mathcal{M}, t \in \mathcal{T} \quad (27)$$

$$S_{j,k,t+1} \geq S_{jkt} - \beta V_{jkt} - \alpha_j \gamma_{jt} (S_{jkt} - \beta V_{jkt}) \left(\sum_{k=1}^K I_{jkt} + I'_{jt} \right) \Delta t, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (28)$$

$$S'_{j,t+1} \geq S'_{jt} + \beta \sum_{k=1}^K V_{jkt} - \alpha_j \gamma_{jt} \left(S'_{jt} + \beta \sum_{k=1}^K V_{jkt} \right) \left(\sum_{k=1}^K I_{jkt} + I'_{jt} \right) \Delta t, \quad \forall j \in \mathcal{M}, t \in \mathcal{T} \quad (29)$$

$$E_{j,k,t+1} \geq E_{jkt} + \left(\alpha_j \gamma_{jt} (S_{jkt} - \beta V_{jkt}) \left(\sum_{k=1}^K I_{jkt} + I'_{jt} \right) - r^J E_{jkt} \right) \Delta t, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (30)$$

$$E'_{j,t+1} \geq E'_{jt} + \left(\alpha_j \gamma_{jt} \left(S'_{jt} + \beta \sum_{k=1}^K V_{jkt} \right) \left(\sum_{k=1}^K I_{jkt} + I'_{jt} \right) - r^J E'_{jt} \right) \Delta t, \quad \forall j \in \mathcal{M}, t \in \mathcal{T} \quad (31)$$

$$I_{j,k,t+1} \geq I_{jkt} + (r^J E_{jkt} - r^d I_{jkt}) \Delta t, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (32)$$

$$I'_{j,t+1} \geq I'_{jt} + (r^J E'_{jt} - r^d I'_{jt}) \Delta t, \quad \forall j \in \mathcal{M}, t \in \mathcal{T} \quad (33)$$

$$U_{j,k,t+1} \geq U_{jkt} + (r^U I_{jkt} - r^D U_{jkt}) \Delta t, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (34)$$

$$H_{j,k,t+1} \geq H_{jkt} + (r^H I_{jkt} - r^D H_{jkt}) \Delta t, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (35)$$

$$Q_{j,k,t+1} \geq Q_{jkt} + (r^Q I_{jkt} - r^D Q_{jkt}) \Delta t, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (36)$$

$$D_{j,k,t+1} = D_{jkt} + r^D (U_{jkt} + H_{jkt} + Q_{jkt}) \Delta t, \quad \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T} \quad (37)$$

$$M_{j,t+1} = M_{jkt} + r^d I'_{jt}, \quad \forall j \in \mathcal{M}, t \in \mathcal{T} \quad (38)$$

$$\mathbf{x}, \mathbf{W} \text{ binary}, \quad \mathbf{C}, \mathbf{V}, \mathbf{S}, \mathbf{S}', \bar{\mathbf{S}}, \mathbf{E}, \mathbf{E}', \mathbf{I}, \mathbf{I}', \mathbf{U}, \mathbf{H}, \mathbf{Q}, \mathbf{D}, \mathbf{M} \geq \mathbf{0}. \quad (39)$$

Equation (14) formalizes the three objectives of the model. The first term corresponds to our primary objective of minimizing the number of deaths over the planning horizon, across all regions and risk classes. This number includes people in the absorbing state D , as well as the transient states H and Q (we ignore undetected deaths). The next terms minimize, as lower-priority objectives, the number of exposed people at the end of the horizon and the distance to the vaccination sites. The hyperparameters λ_E and λ_D are set to small values to prioritize the death-minimization objective.

Next, the constraints capture practical considerations surrounding vaccine distribution:

- *Number of vaccination sites:* We impose a total budget of N vaccination sites (Equation (15)). For obvious reasons, there needs to be at least one site in every state (Equation (16)). We consider $N = 100$ in our experiments, which leaves flexibility to strategically deploy 49 sites.
- *Assignment:* Equation (17) ensures that people get assigned to vaccination sites that have been selected. Equation (18) assigns each population center to exactly one site, in the same state. These assignment constraints are used to compute the distance term in the objective function.
- *Inter-regional vaccine capacity:* Due to restrictions in vaccine manufacturing and distribution networks, a limited number of vaccines can be allocated in each time period. Equation (19) ensures that the total number of vaccines allocated lies within the available budget in each period.
- *Consistency:* Equation (20) ensures that vaccines only get distributed to selected sites. Similarly, Equation (21) ensures that the number of people vaccinated in each state (across risk classes) does not exceed the number of vaccines allocated that state. This constraint involves two assumptions. A first, conservative assumption is that people can only get vaccinated in the state that they live in, which is required in practice for traceability purposes. Another, optimistic assumption is that the vaccine allocation constraint applies to each state, as opposed to each vaccination site. In other words, the model assumes vaccines can be reallocated between vaccination sites within a state, thus maintaining a degree of freedom in intra-state vaccine distribution.
- *Eligibility:* We prevent people from being vaccinated twice: a patient who has been vaccinated but remains susceptible cannot be vaccinated again. Equation (22) defines the number eligible people as the previous number of eligible people minus the number of people for whom the vaccine was effective and the number of people who got exposed to the disease. Equation (22) then ensures that the number of vaccinated people lies below the number of eligible people.
- *Smoothness:* Large fluctuations in the number of vaccines allocated to each region from day to day would likely cause problems from a supply chain management perspective—both to deliver and to administer the vaccines. Equation (24) ensures that such fluctuations remain minimal. The hyperparameter θ_S controls the trade-off between efficiency and smoothness.
- *Fairness:* To be politically and socially viable, vaccine distribution must not neglect any region, even if it is not a virus “hot spot.” This also enhances the robustness of the solution, given that inter-regional transmission can occur in practice. Equation (25) promotes inter-state fairness at the strategic level, by ensuring that the fraction of vaccination sites in each state does not deviate too much from its population share. Equation (26) promotes inter-site fairness, by ensuring that vaccine distribution across sites does not deviate too much from uniform distribution. Finally,

Equation (27) promotes inter-state fairness at the tactical level, by ensuring that no state receives a fraction of vaccines that exceeds its population share by a wide margin. The hyperparameters Θ_L , θ_V , and θ_P control the trade-off between efficiency and fairness. As the results will show, even tight fairness constraints leave critical flexibility when locating vaccination sites and allocating vaccines.

- *DELPHI–V dynamics:* Equations (28)–(38) capture the dynamics of the DELPHI–V model in a discretized time space (Equations (3)–(13)).
- *Domain of definition:* Equation (39) defines the domain of each variable.

3.4 | Model structure

Problem (\mathcal{P}) is a nonlinear program, due to the bilinear terms in Equations (28)–(31), which reflect the fact that the number of new infections result from the interactions between susceptible and infected populations—a key characteristic of all SEIR-based compartmental models. These bilinear terms result in nonconvex constraints, thus in a highly challenging optimization model.

The latest Gurobi 9.0 release includes a solver for nonconvex quadratic problems (Gurobi Optimization, 2020). Yet, general-purpose technologies are limited to small-scale instances. In our setting, Problem (\mathcal{P}) includes $2m(K + 1)T$ nonconvex constraints each involving $2(K + 1)$ bilinear terms, for a total of $4mT(K + 1)^2$ bilinear terms. A realistically-sized problem with $M = 51$ (50 U.S. states plus Washington, DC), $K = 6$ (6 age groups) and $T = 90$ (a 3 month planning horizon with daily discretization) would result in nearly 900 000 bilinear terms. Problem (\mathcal{P}) remains intractable with existing commercial solvers, motivating the development of a tailored algorithm.

4 | SOLUTION ALGORITHM

We propose an iterative coordinate descent algorithm to solve Problem (\mathcal{P}) in short computational times—consistent with practical requirements. We describe the algorithm in this section. We also present three baselines replicating reasonable strategies that could be implemented in the absence of our data-driven optimization model. These baselines are used for two purposes: (i) to provide an initial feasible solution in the coordinate descent algorithm, and (ii) as benchmarks to evaluate the benefits of the data-driven optimization approach proposed in this paper.

4.1 | Algorithm design

Our algorithm relies on two key observations: (i) aside from Equations (28)–(31), the objective function and all other constraints in (\mathcal{P}) are linear, and (ii) given a fixed vaccine

Algorithm 1. Coordinate descent algorithm for DELPHI-V-OPT (Problem (P))**input :** Prescriptive DELPHI-V-OPT data; termination tolerance ε Initialization: $i \leftarrow 0, \mathbf{x}^{(i)} \leftarrow \mathbf{0}, \mathbf{C}^{(i)} \leftarrow \mathbf{0}, \mathbf{W}^{(i)} \leftarrow \mathbf{0}, \mathbf{V}^{(i)} \leftarrow \mathbf{0}, \mathbf{I}^i \leftarrow \mathbf{0}, Z^i \leftarrow \infty$ $i \leftarrow i + 1$ $(\mathbf{x}^{(i)}, \mathbf{C}^{(i)}, \mathbf{W}^{(i)}, \mathbf{V}^{(i)}, \mathbf{I}^{(i)}, Z^{(i)}) \leftarrow \text{GenerateFeasibleSolution}$ **while** $|Z^{(i)} - Z^{(i-1)}| / Z^{(i-1)} > \varepsilon'$ **do** $i \leftarrow i + 1$ Run `Simulate` $(\mathbf{x}^{(i-1)}, \mathbf{C}^{(i-1)}, \mathbf{W}^{(i-1)}, \mathbf{V}^{(i-1)})$. Update $\hat{\mathbf{I}}^{(i)} \leftarrow \hat{\mathbf{I}}$, where $\hat{\mathbf{I}}$ is the output of `Simulate` $(\mathbf{x}^{(i-1)}, \mathbf{C}^{(i-1)}, \mathbf{W}^{(i-1)}, \mathbf{V}^{(i-1)})$.Run `Optimize` $(\hat{\mathbf{I}}^{(i)})$. Update $(\mathbf{x}^{(i)}, \mathbf{C}^{(i)}, \mathbf{W}^{(i)}, \mathbf{V}^{(i)}) \leftarrow (\mathbf{x}, \mathbf{C}, \mathbf{W}, \mathbf{V})$, where $(\mathbf{x}, \mathbf{C}, \mathbf{W}, \mathbf{V})$ is the output of`Optimize` $(\hat{\mathbf{I}}^{(i)})$. Update the objective function:

$$Z^{(i)} \leftarrow \sum_{j=1}^M \sum_{k=1}^K (D_{jkt} + H_{jkt} + Q_{jkt}) + \lambda_E \sum_{j=1}^M \sum_{k=1}^K E_{jkt} + \lambda_D \sum_{l=1}^L \sum_{i=1}^n \text{Pop}_l \Delta_{li} W_{li}$$

end**output:** Vaccine distribution strategy $(\mathbf{x}^{(i)}, \mathbf{C}^{(i)}, \mathbf{W}^{(i)}, \mathbf{V}^{(i)})$

distribution strategy, the discretized DELPHI-V model can be solved efficiently. Therefore, we proceed by coordinate descent, alternating between two modules: one that optimizes the vaccination distribution strategy given the infection dynamics, and one that simulates the bilinear dynamics of the pandemic for a given vaccination distribution strategy. The optimization part reduces to a linear program, which can be solved very efficiently. Using the resultant vaccine distribution, the simulation part re-estimates the infected population under bilinear dynamics, using a forward discretization scheme. Specifically, the two modules are defined as follows:

1. *Simulate*: Based on a vaccine allocation solution \mathbf{V} , we compute the DELPHI-V dynamics from $t = 0$ to $t = T$ (Section 3.2) by solving the ODE system (Equations (3)–(13)) using a forward difference scheme in a discretized time space. This terminates in $\mathcal{O}(MKT)$ operations. We denote the total infected population (across all risk classes and vaccinated people) in region j at time $t \in \mathcal{T}$ by $\hat{I}_{jt} = \sum_{k=1}^K I_{jkt} + I'_{jt}$. We refer to this procedure as `Simulate` $(\mathbf{x}, \mathbf{C}, \mathbf{W}, \mathbf{V})$.
2. *Optimize*: Given the infectious population estimates $\hat{\mathbf{I}}$, we can approximate Equation (28)–(31) by the following linear constraints. The problem can then be efficiently solved as a linear programming model. We refer to this module as `Optimize` $(\hat{\mathbf{I}})$.

$$S_{j,k,t+1} \geq S_{jkt} - \beta V_{jkt} - \alpha_j \gamma_{jt} (S_{jkt} - \beta V_{jkt}) \hat{I}_{jt} \Delta t, \\ \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T}$$

$$S'_{j,t+1} \geq S'_{jt} + \beta \sum_{k=1}^K V_{jkt} - \alpha_j \gamma_{jt} \left(S'_{jt} + \beta \sum_{k=1}^K V_{jkt} \right) \hat{I}_{jt} \Delta t, \\ \forall j \in \mathcal{M}, t \in \mathcal{T}$$

$$E_{j,k,t+1} \geq E_{jkt} + (\alpha_j \gamma_{jt} (S_{jkt} - \beta V_{jkt}) \hat{I}_{jt} - r^I E_{jkt}) \Delta t, \\ \forall j \in \mathcal{M}, k \in \mathcal{K}, t \in \mathcal{T}$$

$$E'_{j,t+1} \geq E'_{jt} + \left(\alpha_j \gamma_{jt} \left(S'_{jt} + \beta \sum_{k=1}^K V_{jkt} \right) \hat{I}_{jt} - r^I E'_{jt} \right) \Delta t, \\ \forall j \in \mathcal{M}, t \in \mathcal{T}$$

We iterate between the `Simulate` and `Optimize` modules, until convergence. Specifically, the algorithm terminates when the variation in the objective function value remains minimal from one iteration to the next. The pseudocode summarizing this approach is presented in Algorithm 1. We turn next to the generation of an initial feasible solution.

4.2 | Baselines

We propose three simple and interpretable baselines for generating a feasible solution to (P). By design, these baselines are heuristics that solely rely on the inputs of the optimization models, as opposed to requiring the full model and algorithm developed in this paper.

4.2.1 | Top-cities baseline

This approach prioritizes cities based on population. Specifically, it deploys vaccination sites in the most populous cities, while accounting for the constraint that each state must have at least one center (Equation (16)). Subsequently, it allocates an equal fraction of the daily vaccine budget to each vaccination site: we fix the variables \mathbf{x} and \mathbf{C} , run the model to optimize vaccine allocation within each state, and estimate the resulting number of deaths. This baseline corresponds to a city-level approach based on demographic information alone.

4.2.2 | Population-based baseline

Under this approach, the number of vaccination sites deployed in each state is based on the state's population share. This is formulated as follows, where y_j is a decision variable denoting the number of vaccination sites in state j .

$$\begin{aligned} \min \sum_{j=1}^m \left| \frac{\sum_{l \in \mathcal{L}_j} \text{Pop}_l}{\sum_{j'=1}^m \sum_{l \in \mathcal{L}_{j'}} \text{Pop}_l} N - y_j \right|, \\ \text{s.t. } \sum_{j=1}^m y_j = 100, \quad y_j \geq 1, \forall j \in \mathcal{M}, \quad \mathbf{y} \text{ integer.} \end{aligned}$$

We then solve DELPHI-V-OPT, while fixing the aggregate number of vaccination sites per state, i.e., $\sum_{i \in \mathcal{N}_j} x_i = y_j$, and assuming equal allocation of vaccines across vaccination sites, i.e., $C_{it} = \frac{1}{N} B_t, \forall t \in \mathcal{T}$. The model allocates vaccines within each state and estimates the number of deaths. This baseline corresponds to a state-level approach based on demographic information alone.

4.2.3 | Case-based baseline

Under this approach, the number of vaccination sites deployed in each state is based on the number of COVID-19 cases at the beginning of the planning horizon. This is formulated as follows, where y_j is a decision variable denoting the number of vaccination sites in state j and $Cases_j$ denotes the case count in state j .

$$\begin{aligned} \min \sum_{j=1}^m \left| \frac{Cases_j}{\sum_{j'=1}^m Cases_{j'}} N - y_j \right|, \\ \text{s.t. } \sum_{j=1}^m y_j = 100, \quad y_j \geq 1, \forall j \in \mathcal{M}, \quad \mathbf{y} \text{ integer.} \end{aligned}$$

We then proceed as with the population-based baseline, by fixing the number of vaccination sites per state, assuming equal vaccine allocation across sites, and re-solving the model. This baseline corresponds to a state-level approach based on epidemiological information alone.

By design, these baselines satisfy all constraints of Problem (P), and thus provide valid initializations into our coordinate descent algorithm. They also provide sensible and equitable benchmarks based on readily-available demographic information (e.g., census data) and epidemiological information (e.g., case counts), hence easily implementable. Comparisons between our optimized solution and these benchmarks thus estimate the benefits of vaccine distribution optimization.

5 | EXPERIMENTAL SETUP

We implement the proposed model and algorithm in the United States. We select $N = 100$ vaccination sites out of the 500 most populous cities in the United States as candidate locations (set \mathcal{N}). We define the set \mathcal{M} as 51 “states” (the 50 states plus the District of Columbia) and the set \mathcal{L} as the 3006 counties. We define six risk classes based on six relatively coarse age groups: 0–9 years, 10–49 years,

50–59 years, 60–69 years, 70–79 years, and 80 years and above. These simplified risk classes facilitate the practical implementation of the solution while capturing broad trends in mortality rates. We define the time horizon \mathcal{T} as the three-month period from February to April 2021, consistently with the ongoing planning horizon of the U.S. federal government.

5.1 | Data sources

We calibrate the model using multiple data sources (Figure 4). First, we estimate the parameters of the DELPHI model (without vaccinations) independently for each state, using historical data on cases and deaths from the New York Times (2020). We obtain a granular population breakdown by age for each state from the U.S. Census Bureau (2020). We then run DELPHI (still, without vaccinations) to derive the initial susceptible, exposed and infected populations (on January 30, 2021), which we distribute among the risk classes proportionally to their size.

The next input is an estimate of the mortality rate per region, risk class and time period. We make use of the data from the U.S. Center for Disease Control (2021), which report the total number of confirmed COVID-19 cases, hospitalizations and deaths by age group in the United States until the end of January 2021. In contrast, the DELPHI model fits a time-dependent mortality rate within each region. To the best of our knowledge, there exists no other data source for nationwide cases and deaths by age group. This leaves a discrepancy between the time-independent estimates at the risk class level from the CDC, and the time-varying estimates at the region-level from DELPHI. To reconcile these data, we employ an optimization procedure that interpolates the mortality rate per region, risk class and time period. We present this approach in the next section.

5.2 | Mortality rate estimation

Our procedure to estimate mortality rates starts from two sets of inputs:

- *DELPHI predictions*: Let \hat{X}_{jt} denote the estimated number of new detected cases in region $j \in \mathcal{M}$ and time period $t \in \mathcal{T}$. Let $\hat{D}_{j,t+\tau_j}^d$ denote the number of deaths, where τ_j is the median death time after detection in region j . These quantities are aggregated across risk classes.
- *CDC data*: Let X_k^{CDC} and D_k^{CDC} denote the number of cases and deaths for risk class $k \in \mathcal{K}$. These quantities are aggregated across regions and time periods.

We define the reference mortality rate \bar{m}_{jk} of risk class $k \in \mathcal{K}$ based in region j as follows:

$$\bar{m}_{jk} = \frac{D_k^{\text{CDC}}}{X_k^{\text{CDC}}} \left(\frac{\sum_{t=1}^T \hat{D}_{j,t+\tau_j}^d}{\sum_{t=1}^T \hat{X}_{jt}} \right) \left(\frac{\sum_{l=1}^K X_l^{\text{CDC}}}{\sum_{l=1}^K D_l^{\text{CDC}}} \right) \quad (40)$$

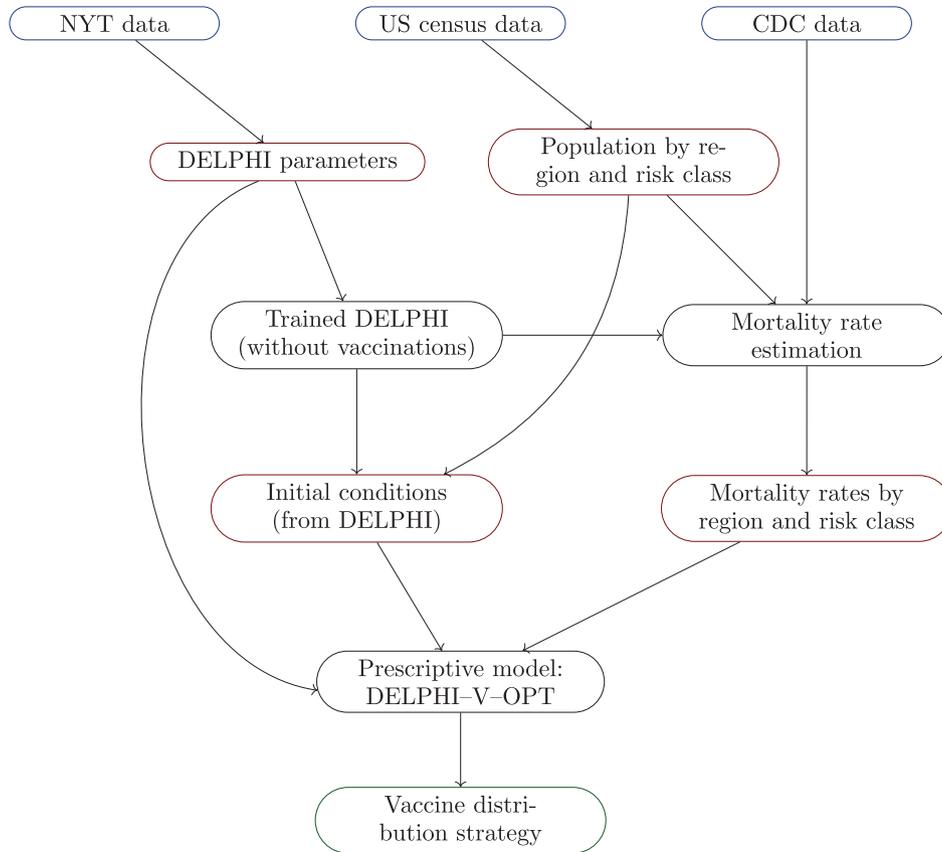


FIGURE 4 Simulation environment: Raw data (blue), processed data (red), models (black) and outputs (green)

By design, this expression preserves the ratio of mortality rates between different risk classes from the CDC data, while correcting the mean reference mortality rate in each region across the planning horizon to be in line with the DELPHI projections.

We then estimate the mortality rate for each region $j \in \mathcal{M}$, risk class $k \in \mathcal{K}$, and time period $t \in \mathcal{T}$, denoted by m_{jkt} . We also introduce additional decision variables X_{jkt} and $D_{j,k,t+\tau_j}$, reflecting the number of detected cases and the number of future deaths in region j assigned to risk class $k \in \mathcal{K}$ and time period $t \in \mathcal{T}$. Given that the fitting procedure is done separately in each region $j \in \mathcal{M}$, we decouple the problem at the region level—thereby considerably reducing the size of each problem instance. Specifically, we formulate the optimization problem given in Equations (41)–(48).

for all $j \in \mathcal{M}$:

$$\min \sum_{t=1}^T \sum_{k=1}^K \left[\left(\frac{m_{jkt} - \bar{m}_{jk}}{\bar{m}_{jk}} \right)^2 + \lambda_M \left| \frac{X_{jkt} - p_{jk} \hat{X}_{jt}}{p_{jk} \hat{X}_{jt}} \right| \right] \quad (41)$$

$$\text{s.t.} \sum_{k=1}^K D_{j,k,t+\tau_j} = \hat{D}_{j,t+\tau_j}, \forall t \in \mathcal{T} \quad (42)$$

$$\sum_{k=1}^K X_{jkt} = \hat{X}_{jt}, \forall t \in \mathcal{T} \quad (43)$$

$$m_{jkt} \cdot X_{jkt} = D_{j,k,t+\tau_j}, \forall k \in \mathcal{K}, \forall t \in \mathcal{T}, \quad (44)$$

$$m_{jkt} \geq m_{j,k,t+1}, \forall k \in \mathcal{K}, \forall t \in \mathcal{T} \quad (45)$$

$$m_{jkt} \geq m_{jlt}, \forall k, l \in \mathcal{K} : \bar{m}_k \geq \bar{m}_l, \quad (46)$$

$$m_{jkt} \leq 1, \forall k \in \mathcal{K}, \forall t \in \mathcal{T} \quad (47)$$

$$\mathbf{m}, \mathbf{C}, \mathbf{D} \geq \mathbf{0}. \quad (48)$$

The first term in Equation (41) minimizes the squared relative error between the mortality rate estimates and their reference values (Equation (40)). The second term is a regularization penalty that minimizes deviations between the proportion of detected cases and the proportion p_{jk} of the population, in each risk class. The parameter λ_M trades off these two objectives (we use $\lambda_M = 0.1$ in our experiments). Equations (42) and (43) ensure consistency with the DELPHI predictions. Equation (44) defines the mortality rate as the ratio between the number of deaths and cases. Equations (45) and (46) ensure that mortality rates are decreasing over time and monotonic with risk classes. Finally, Equations (47) and (48) define the domain of the variables.

Note that the problem is nonlinear due to the bilinear term in Equation (44). Yet, thanks to the decoupling at the region level, we can solve the problem using the quadratic solver in Gurobi 9.0 (Gurobi Optimization, 2020), which addresses nonconvexities using branching and cutting planes algorithms. In practice, a solution within a 1% optimality gap is generally obtained within minutes.

Table 1 Calibrated monthly mortality rates, averaged by risk class and over all states

Month	0–9	10–49	50–59	60–69	70–79	80+
February	0.008%	0.119%	0.882%	2.271%	6.101%	15.027%
March	0.007%	0.116%	0.859%	2.212%	5.942%	14.636%
April	0.007%	0.114%	0.844%	2.171%	5.832%	14.365%

The output of this algorithm is an estimate of the mortality rate at the level of each state and each risk class, in each time period. We report aggregated statistics in Table 1.

5.3 | Implementation details

Our DELPHI–V–OPT model relies on a forward difference scheme to simulate the dynamics of the pandemic, for any vaccine allocation. If the discretization is too coarse, the algorithm will introduce truncation errors. If, however, the discretization is too granular, computational times will be prohibitively long. After extensive experimentations, we found $\Delta t = 1$ day to yield the best compromise. It is also practical choice as it yields a day-by-day plan.

The coordinate descent scheme terminates when the change in the objective function value lies below a pre-determined threshold. We set the termination tolerance to $\epsilon = 0.1\%$.

Regarding DELPHI–V–OPT, we vary the hyperparameters λ_E , λ_V , θ_S , Θ_L , θ_V , and θ_P to balance efficiency with practical and fairness considerations, and perform sensitivity analyses on these parameters. We consider a baseline vaccine effectiveness of 90% (in line with the values from the first vaccine approvals) and a baseline budget of 1 million vaccines per day (which can be viewed as 1 million single-dose vaccines, 2 million double-dose vaccines or a combination thereof). Given the strong underlying uncertainty, we perform sensitivity analyses to demonstrate the robustness of the benefits of our optimization outputs to vaccine effectiveness and vaccine budget.

All optimization models are implemented in Gurobi 9.0, with 2.3 GHz processor and four cores. We use a barrier method to solve each linear program, with a barrier convergence tolerance of 10^{-6} .

6 | EXPERIMENTAL RESULTS

We now present results obtained with the modeling and algorithmic framework developed in this paper. Our main focus is on the location of the 100 vaccination sites—decisions that have to be made immediately—as opposed to vaccine allocation—decisions that can be revisited continuously over time as more information becomes available. We first evaluate the benefits of the optimized solution against baseline approaches, and then establish the robustness of these benefits.

6.1 | Benefits of optimizing the vaccine distribution strategy

Figure 5 shows the heatmap of vaccination sites across all 51 states (the exact list of proposed locations is reported in the Appendix). Recall that, by design, the optimized solution deploys at least one vaccination site per state (Equation (16)). In addition, one of the fairness constraints (Equation (25)) anchors the number of vaccination sites per state to its population share. Other than that, the vaccination sites are deployed strategically to curb the spread of the pandemic. As a result, the optimized number of vaccination sites varies significantly per state, as a function of underlying demographic and epidemiological factors. For instance, the four largest states by population (California, Texas, Florida and New York) receive the most vaccination sites (6–10 each). In contrast, the smaller states receive only one vaccination site. Ultimately, the heatmap suggests that the optimized solution targets large population centers and some of the hot spots of the pandemic, while ensuring equity nationwide.

Then, could we have achieved a similar solution with some of our benchmarks (Section 4.2), which determine the locations of the vaccination sites based on demographic and epidemiological data but do not make use of our epidemiological and optimization models? To investigate this question, we evaluate the dynamics of the pandemic under each of the three baselines. When it comes to the optimization, we derive three solutions: (i) an “optimized locations” solution, which optimizes the site locations decisions freely (without fairness constraints) but then allocates vaccines uniformly across the 100 selected vaccination sites; (ii) a “fully optimized” solution, which optimizes the site locations decisions and vaccine allocation decisions freely (without fairness constraints); and (iii) a “proposed” solution, which optimizes the site locations decisions and vaccine allocation decisions with fairness constraints. For each solution, Table 2 reports the number of lives saved by the vaccination campaign, as compared to a solution without vaccinations.

Let us start with the main observation: the optimized solutions provides significant benefits, as compared to all benchmarks. Comparing the “optimized locations” solution to the top-cities benchmark, we find that optimizing locations alone can save around 4500 lives over the three-month period under consideration, enhancing the impact of the vaccination campaign by 24%. Moving to the “fully optimized” solution, we note that jointly optimizing locations and vaccine allocation achieves further benefits, by saving an extra 2000 lives and increasing the benefits over the benchmark to 35%. These

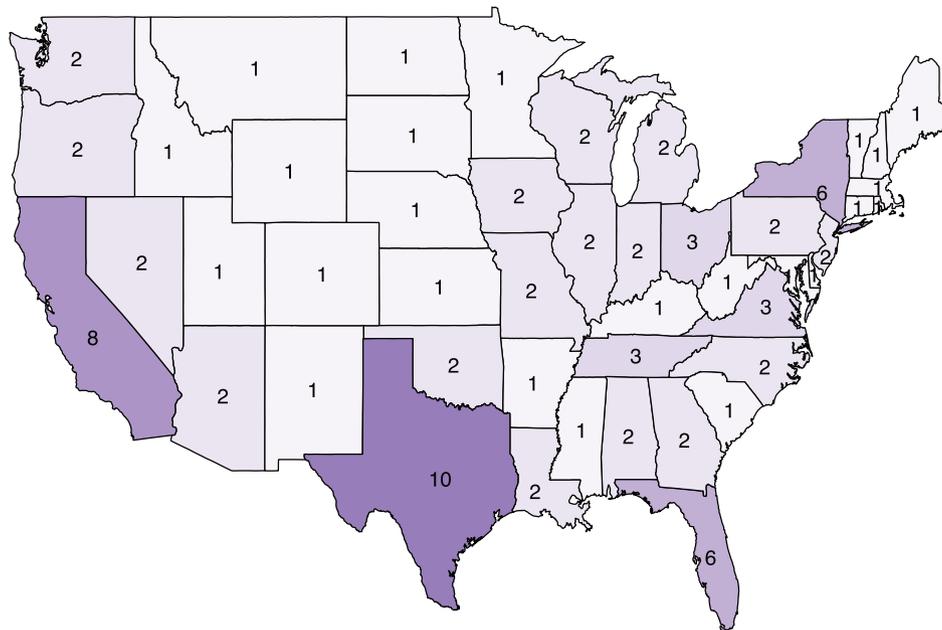


FIGURE 5 Number of centers per state in the proposed solution

Table 2 Comparison of optimized solutions and benchmarks

Method	Solution	Site locations	Vaccine distribution	Saved lives
Benchmarks	Top-cities	Most populous cities	Uniform across centers	19 045 (base)
	Population	Pro-rata population	Uniform across centers	19 709 (+3.5%)
	Cases	Pro-rata active cases	Uniform across centers	21 037 (+10.5%)
Optimization	Locations	Optimized: minimizes deaths	Uniform across centers	23 622 (+24.0%)
	Full	Optimized: minimizes deaths	Optimized: minimizes deaths	25 615 (+34.5%)
	Proposed	Optimized: minimizes deaths	Optimized: minimizes deaths	23 000 (+20.8%)
		Fair: interstate equity	Fair: intercenter equity	

results underscore the potential of the proposed optimization approach, which leverages available vaccines strategically to target the regions that need them most. Another observation is that determining the locations of vaccination sites does not achieve all potential benefits of the vaccination campaign, underscoring the role of downstream vaccine allocation as an extra lever to combat the pandemic.

A downside of these two solutions (“locations optimized” and “fully optimized”), however, is that they can result in very inequitable outcomes between states and between vaccination sites. The last (“proposed”) solution circumvents this challenge by imposing all fairness constraints (Equations (25)–(27)), and choosing tight values for the corresponding hyperparameters Θ_L , θ_V , and θ_P . As such, the proposed solution ensures fairness across various dimensions (site locations, vaccine allocation across states, and vaccine allocation across sites). Remarkably, even when constraining the optimization as such, the resulting proposed solution still results in 20% benefits, as compared to the top-cities benchmark—saving 4000 deaths over the three-month horizon.

In comparison, the benchmarks cannot reach the same impact of vaccinations. The top-cities and population-based benchmarks achieve similar outcomes, with 19 000–20 000

lives saved as compared to a no-vaccination baseline. The case-based baseline performs better, by saving 21 000 lives (10% more than the top-cities baseline). Yet, these numbers remain significantly lower than those achieved with our solution. The top-cities and population-based baselines, by relying on demographic information exclusively, fail account for disparities in the severity of the across the country. The case-based baseline captures the status of the pandemic, but not its dynamics, thus treating similarly a state where the pandemic is already waning and a state where it is rising fast—although vaccinations have a stronger impact in the latter state than the former.

These results underscore the main takeaways from this paper: optimization provides significant benefits, as compared to simple benchmarks based on readily-available information at the outset of the vaccination campaign. Instead, the optimized approaches design vaccine distribution strategies based on both demographics and epidemiological dynamics. The edge of optimization can be significant, by saving an extra 20% of lives with the same vaccine capacity. Stated differently, under the proposed optimization approach, each vaccine is effectively “worth” 1.2 vaccines.

A natural question to ask is whether optimization induces strong geographic disparities, by merely shifting the benefits

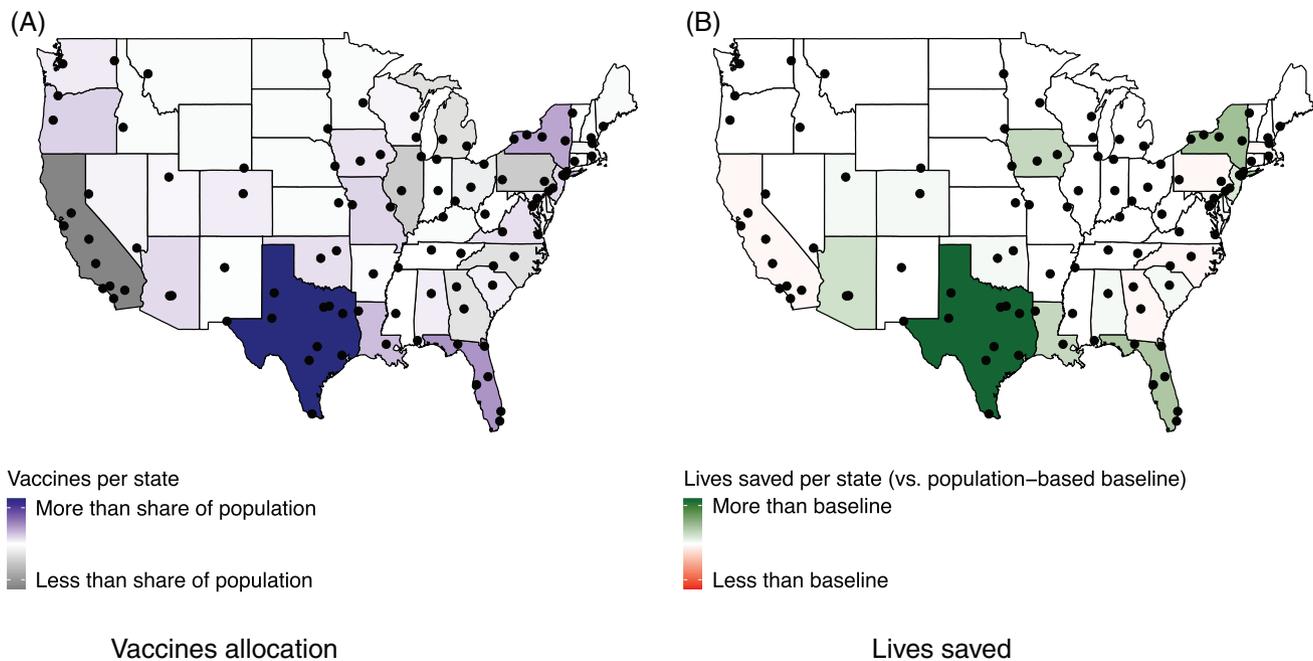


FIGURE 6 Vaccine allocation (vs. population share) and lives saved (vs. population-based baseline) per state

of the vaccines from one state to another. To explore this question, Figure 6 plots the number of vaccines distributed to each state, as compared to its population share (Figure 6a) and the number of lives saved, as compared to the top-cities benchmark (Figure 6b).

Recall that the optimization imposes fairness constraints in vaccine allocation (Equations (25)–(27)); yet, the remaining flexibility can be used strategically to target the populations where vaccines can have the strongest impact. In fact, as Figure 6a shows, vaccines do not get distributed proportionally to each state’s population. For instance, states like Texas, Florida and New York receive a higher share of vaccines, whereas states like California, Pennsylvania and Illinois receive a lower share. This is expected, given the significantly higher number of lives saved under the optimized (“proposed”) solution as compared to the population-based benchmark (Table 2).

However, Figure 6b shows that these disparities in vaccine allocation do not result in sharp disparities in public health outcomes. Specifically, the optimized solution saves hundreds of extra lives (as compared to the population-based benchmark) in seven states with very different epidemiological profiles. Texas benefits the most from optimization (with an additional 1450 lives saved), followed by New York (440), Florida (380) and Iowa (290). At the same time, the optimized solution does not increase the death toll in any state by more than 100. Pennsylvania is the most negatively impacted state, with an estimated 85 additional deaths—well within the margin of error of DELPHI–V. Overall, the proposed optimization approach can thus distribute vaccine capacities to combat the pandemic in some critical states without hurting others.

The obvious question then is: how does the proposed solution compare to the vaccination plan that was applied

in practice? Unfortunately, it is hard to perform a complete apples-to-apples comparison. Indeed, the deployment of vaccination site was conducted in a more ad hoc fashion in practice than modeled in this paper, based on local capacity, heterogeneity across “big” and “small” sites, and so on. Moreover, since this paper was written, the Federal government has updated its plan, so that the Federal facilities take a supporting role while delegating most authority to state-driven vaccination efforts. Therefore, the full plan was not implemented and instead FEMA has opened just 24 facilities in 13 states, as compared to the original 100 planned ones. We provide details in Table A2 in the Appendix.

Nonetheless, we provide two sources of evidence supporting the results from our analysis. First, we observe that FEMA allocated a higher share to states such as Florida than would be otherwise allocated on a population pro rata basis. Conversely, other states like California were underweighted. This is consistent with the main recommendations outlined in the paper. Second, even though our recommendations were not applied directly to the location of vaccination centers, the subsequent allocation of vaccines to each state was influenced by the model’s recommendations. This underscores a second tactical lever in our model-based recommendation—vaccine allocation—beyond its main strategic lever—location of vaccination sites.

6.2 | Sensitivity and robustness

A core challenge in vaccine distribution lies in the significant uncertainty regarding the dynamics of the pandemic and the effect of vaccinations. To address this challenge, we assess the sensitivity and the robustness of the optimized solution when the structure of the DELPHI–V model and some of its key

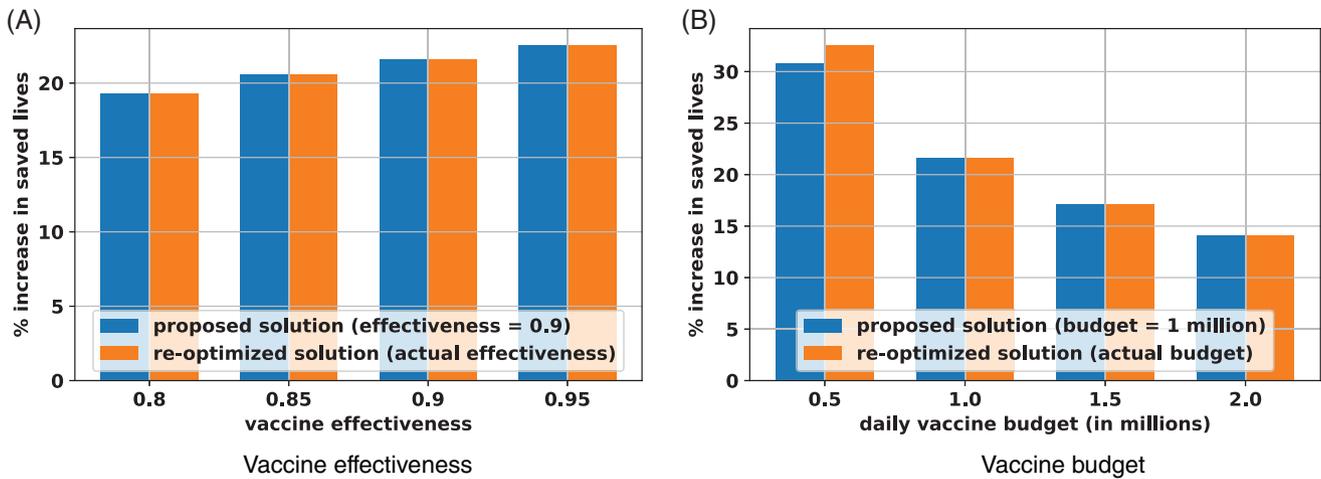


FIGURE 7 Sensitivity and robustness of results with varying vaccine effectiveness and vaccine budget

parameters are perturbed. For each perturbation, we compare three solutions:

- the top-cities baseline, evaluated with the new perturbed model
- the re-optimized solution, optimized and evaluated with the new perturbed model
- the proposed solution, optimized with the original model and evaluated with the new perturbed model. Specifically, we first run the optimization with the initial inputs. We then introduce perturbations, and re-optimize vaccine allocation decisions (variables C and V), while fixing the vaccination sites (variable x). Indeed, in practice, vaccination sites are determined once and for all, whereas vaccine allocation can be re-adjusted as information becomes available.

We then compare the number of deaths under these three solutions, estimated with the perturbed DELPHI-V model. Comparisons between the top-cities benchmark and the re-optimized solution estimate the sensitivity of the benefits of optimization to the perturbations. Comparisons between the top-cities benchmark and the proposed solution estimate the robustness of the solution.

We first vary the two main drivers of the vaccination campaign, that are the vaccine effectiveness (parameter β) and the vaccine budget (parameters B_t). Figure 7 reports the percent-wise increase in saved lives, as compared to the top-cities benchmark, for both the proposed and the re-optimized solutions. The main takeaways fall under three categories. First, the benefits of optimization increase with vaccine effectiveness. This is expected, as a higher vaccine effectiveness increases the impact of all strategies (in the extreme example where $\beta = 0$, all strategies have the same null performance). Second, the benefits of optimization decrease with vaccine budget. Indeed, in the extreme scenario with an infinite vaccine budget, any distribution strategy can immediately end the pandemic, leaving essentially no space for optimization. As the budget becomes more scarce, the decisions of *who* receives a vaccine and *when*

become increasingly complex, so the optimized strategy has a positive and significant impact on the spread of the disease. It is interesting to note that this monotonic relationship would get inverted in the other regime with a small vaccine budget (again, all solutions perform identically when the budget gets to zero). Yet, with the current vaccine capacities, the benefits of optimization are very strong, saving an extra 15–35% of lives.

The third takeaway from Figure 7 is the high degree of robustness of the proposed solution. In all but one experiment, the proposed solution remains optimal under the perturbed parameters (indicated by exactly the same benefits obtained with the proposed solution and the re-optimized solution). In the last case (with a daily budget of 500 000 vaccines), the proposed solution is dominated by the re-optimized solution, but remains within 2% of the new optimum. Obviously, the new values of vaccine effectiveness and vaccine budgets impact downstream vaccine allocations and the dynamics of the pandemic. However, the location of vaccination sites is highly robust to variations in vaccine effectiveness and vaccine budgets.

Next, we test the robustness of the proposed solution to the dynamics of the pandemic. One assumption of DELPHI-V is that the infection rate (captured by the parameters α_j and γ_{jt}) is identical across age groups. However, serological evidence suggests potential disparities in infections across sub-populations. To test this, we vary the infection rates with the risk class k , by adjusting the relative variations based on the serological estimates from the U.S. Center for Disease Control (2020a). Another assumption is that vaccines prevent mortality but not infection and transmission. In practice, vaccines may still reduce the risk of infection and the propensity to transmit the disease. To test this, we perturb the DELPHI-V model by assuming that a fraction β of vaccinated people transition directly to the immune state M (as opposed to the susceptible and vaccinated state S').

Figure 8 reports the results from these new perturbations, using the same nomenclature as Figure 7. The main takeaway is identical, in that the proposed solution remains

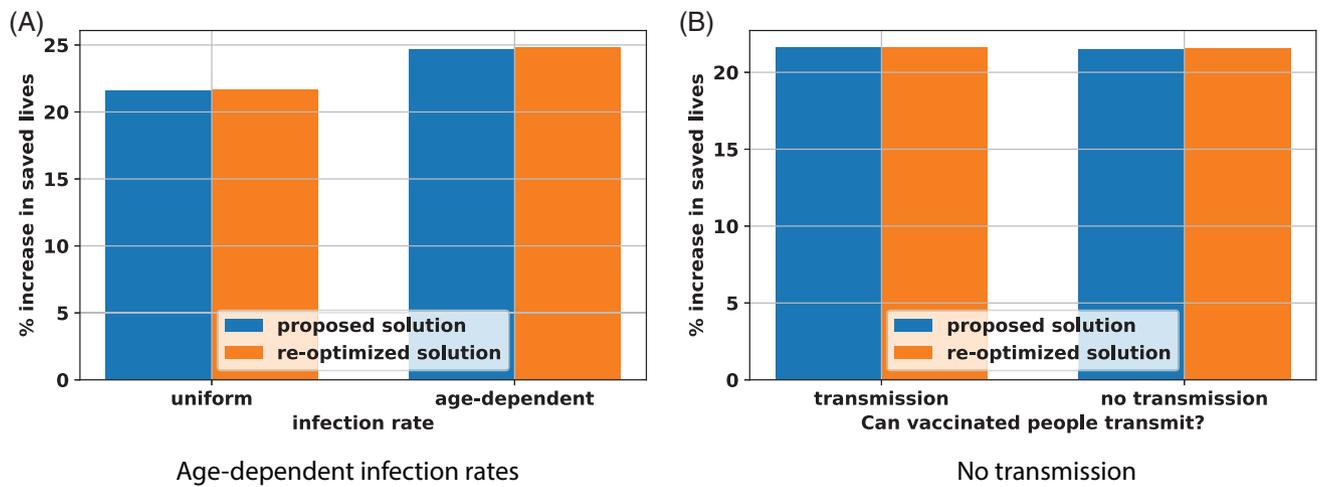


FIGURE 8 Robustness of results with age-dependent infection rates and no transmission from vaccinated people

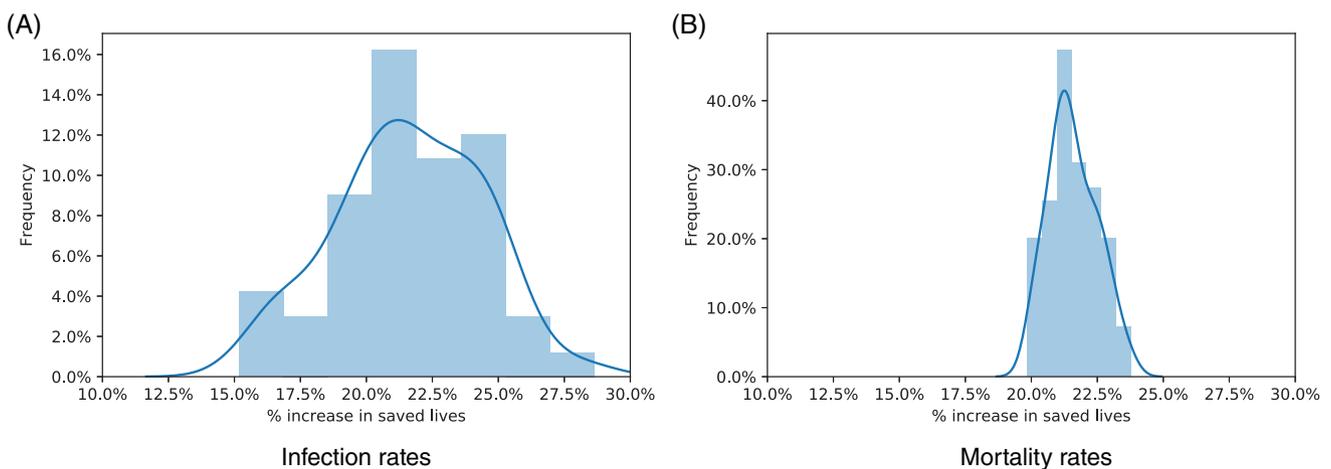


FIGURE 9 Robustness of the benefits of optimization with infection and mortality rates

near-optimal under the two perturbations. In both tests, the proposed solution results in less than 50 extra deaths than the re-optimized solution—a very small amount in view of the 20 000 lives saved by the optimization. In other words, the proposed solution is not only robust to vaccination characteristics (effectiveness and budget) but also to the dynamics of the COVID-19 pandemic.

Finally, we vary two key parameters in the DELPHI-V dynamics: the infection and mortality rates. Although these parameters are fitted from historical data, there remains considerable uncertainty regarding the dynamics of the pandemic and people's behaviors through a period of mass vaccination. Therefore, we introduce a random perturbation in the infection or mortality rate, in each state. Specifically, we define 50 perturbation scenarios; in each one, we sample each parameter in each state independently, following a uniform distribution centered around the nominal value and spanning $\pm 20\%$. Therefore, the full range of infection and mortality rates spans 40% of the nominal value—thus capturing instances where the estimated infection and mortality rates in the DELPHI-V model are subject to very large errors.

Figure 9 reports the distribution of the benefits of optimization under these 50 scenarios, with perturbed infection rates

(Figure 9a) and mortality rates (Figure 9b). We compare here the proposed solution (obtained with the nominal values of the infection and mortality rates) and the top-cities benchmark, both evaluated with the perturbed infection and mortality rates. As the results show, the benefits of optimization remains highly significant, and robust to the perturbations. Recall that the benefits of optimization were estimated at 20% under the nominal infection and mortality rates (Table 2). After perturbations, they span 12.5–30% with perturbations in the infection rates and 17.5%–25% with the perturbed mortality rates. In other words, even if the DELPHI-V model makes large errors when estimating the key dynamics of the pandemic, the proposed optimized vaccination sites still increase the impact of the vaccination campaign by over 10%. In fact, the variations can go either way: in over half the simulations, the relative benefits of the optimized solution are even higher under the perturbed parameters than the nominal ones.

Note, finally, that the infection rate has a more significant impact on the relative benefits of optimized vaccine allocation than the mortality rate. This suggests that the impact of the pandemic depends mainly on how vaccines can curb infections at the upstream level—as opposed to mortality at the downstream level. This, again, illustrates the effects

of the nonlinear SEIR dynamics on the spread of the disease, and how we can leverage an epidemiological model such as DELPHI-V to allocate resources strategically in order to combat the pandemic most effectively.

From a practical standpoint, the results from this section are highly significant. Indeed, the DELPHI-V model, like any epidemiological model, only provides a rough approximation of the dynamics of the pandemic and the effects of vaccinations. Yet, the robustness of the solution provides guarantees that the locations of mass vaccination sites, even optimized against this approximate model, remain highly robust when evaluated against alternative dynamics. This is obviously not to say that we can commit to a full-scale vaccine distribution strategy that spans the entire vaccination campaign. However, the “here-and-now” decisions (i.e., the location of vaccination sites) are likely to remain near-optimal (or even optimal) in the next phases of the pandemic, ultimately enabling the vaccination campaign to save an extra 15–35% of lives.

7 | CONCLUSION

This paper has presented a new prescriptive approach to optimize vaccine distribution strategies in response to the COVID-19 pandemic. The approach starts with a state-of-the-art epidemiological model called DELPHI, which augments SEIR models by capturing dynamics specific to COVID-19 (under-detection, governmental and societal response, and declining mortality rates). This paper has first proposed an extension, named DELPHI-V, which captures the effects of vaccinations and reflects the disaggregated impact of COVID-19 on mortality across risk classes (e.g., age groups). Then, this paper has embedded the predictive DELPHI-V model into an optimization model, termed DELPHI-V-OPT, to support vaccine distribution. DELPHI-V-OPT is formulated as a bilinear optimization model, and solved using a tailored algorithm based on coordinate descent.

We applied the model and algorithm to one of the priorities of the new Biden administration in the United States: determining the locations of mass vaccination sites across the country. We formalize the problem by selecting locations strategically to minimize the death toll of the pandemic, subject to practical and fairness constraints. Experimental results using real-world data suggest that the proposed optimization approach can yield significant benefits, as compared to benchmark solutions that locate vaccination sites based on readily-available demographic and epidemiological information. Specifically, the model can increase the effectiveness of the vaccination campaign by 20%, saving an extra 4000 lives over a three-month period. Remarkably, the proposed solution achieves critical fairness objectives—by significantly reducing the death toll of the pandemic in several states without hurting others—and is highly robust to uncertainties and

forecast errors—by achieving similar benefits under a vast range of perturbations.

Obviously, the optimization approach developed in this paper is not without limitations. For instance, our experiments have only partitioned the population according to age groups, thus ignoring other objectives such as prioritizing allocations to healthcare workers, other essential workers, or patients with comorbidities. Moreover, our epidemiological model does not capture heterogeneity in population mixing across subpopulations (e.g., interactions may be more intense in urban areas and between young people than in rural areas and between older populations). Similarly, our model only captures the first-order effects of vaccinations, ignoring for instance heterogeneity across multiple vaccine types and double-dose vaccines. Finally, our methodology relies on a time discretization approximation and a coordinate descent approach, which do not yield theoretical guarantees on solution quality.

Whereas these limitations undoubtedly motivate further research, this paper lays one of the first data-driven bricks on the optimal distribution of COVID-19 vaccines at a macroscopic level. At a time where vaccine development and vaccine production are going full speed, the results from this paper highlight the critical role of vaccine distribution strategies to combat the pandemic. Obviously, it is essential to make every effort possible to develop new vaccines, enhance vaccine effectiveness, and produce as many vaccines as possible. But this paper highlights another lever that can be pulled to curb the effect of the pandemic: strategically managing vaccine stockpiles to prevent the spread of the pandemic and mitigate its impact. As such, this paper can provide critical decision-making support to governmental agencies as they are currently planning mass vaccinations around the globe.

ACKNOWLEDGMENTS

The authors gratefully acknowledge T. Greg McKelvey Jr. for feedback and discussions that have greatly improved the quality of the paper.

REFERENCES

- Aaby, K., Herrmann, J. W., Jordan, C. S., Treadwell, M., & Wood, K. (2006). Montgomery county's public health service uses operations research to plan emergency mass dispensing and vaccination clinics. *Interfaces*, 36(6), 569–579.
- Araz, O. M., Galvani, A., & Meyers, L. A. (2012). Geographic prioritization of distributing pandemic influenza vaccines. *Health Care Management Science*, 15(3), 175–187.
- Arifoğlu, K., Deo, S., & Irvani, S. M. (2012). Consumption externality and yield uncertainty in the influenza vaccine supply chain: Interventions in demand and supply sides. *Management Science*, 58(6), 1072–1091.
- Arinaminpathy, N., Ratmann, O., Koelle, K., Epstein, S. L., Price, G. E., Viboud, C., Miller, M. A., & Grenfell, B. T. (2012). Impact of cross-protective vaccines on epidemiological and evolutionary dynamics of influenza. *Proceedings of the National Academy of Sciences*, 109(8), 3173–3177.

- Bae, J., Gandhi, D., Kothari, J., Shankar, S., Bae, J., Patwa, P., ... et al. (2020). Challenges of equitable vaccine distribution in the covid-19 pandemic. *arXiv preprint arXiv:2012.12263*.
- Bandi, H., & Bertsimas, D. (2020). *Optimizing influenza vaccine composition: From predictions to prescriptions*. In Machine Learning for Healthcare Conference (pp. 121–142). PMLR.
- Basta, N. E., Chao, D. L., Halloran, M. E., Matrajt, L., & Longini, I. M., Jr. (2009). Strategies for pandemic and seasonal influenza vaccination of schoolchildren in the United States. *American Journal of Epidemiology*, *170*(6), 679–686.
- Bertsimas, D., Boussioux, L., Wright, R. C., Delarue, A., Digalakis, Jr V., Jacquillat, A., ... , et al. (2020) From predictions to prescriptions: A data-driven response to COVID-19. *arXiv preprint arXiv:2006.16509*.
- Chick, S. E., Mamani, H., & Simchi-Levi, D. (2008). Supply chain coordination and influenza vaccination. *Operations Research*, *56*(6), 1493–1506.
- Cho, S. H. (2010). The optimal composition of influenza vaccines subject to random production yields. *Manufacturing & Service Operations Management*, *12*(2), 256–277.
- Chowell, G., Viboud, C., Wang, X., Bertozzi, S. M., & Miller, M. A. (2009). Adaptive vaccination strategies to mitigate pandemic influenza: Mexico as a case study. *PLoS One*, *4*(12), e8164.
- Coustasse, A., Kimble, C., & Maxik, K. (2021). Covid-19 and vaccine hesitancy: A challenge the United States must overcome. *The Journal of Ambulatory Care Management*, *44*(1), 71–75.
- Dai, T., Cho, S. H., & Zhang, F. (2016). Contracting for on-time delivery in the us influenza vaccine supply chain. *Manufacturing & Service Operations Management*, *18*(3), 332–346.
- Dror, A. A., Eisenbach, N., Taiber, S., Morozov, N. G., Mizrachi, M., Zigron, A., Srouji, S., & Sela, E. (2020). Vaccine hesitancy: The next challenge in the fight against covid-19. *European Journal of Epidemiology*, *35*(8), 775–779.
- Duijzer, L. E., van Jaarsveld, W., & Dekker, R. (2018). Literature review: The vaccine supply chain. *European Journal of Operational Research*, *268*(1), 174–192.
- Dushoff, J., Plotkin, J. B., Viboud, C., Simonsen, L., Miller, M., Loeb, M., & David, J. (2007). Vaccinating to protect a vulnerable subpopulation. *PLoS Medicine*, *4*(5), e174.
- Elveback, L. R., Fox, J. P., Ackerman, E., Langworthy, A., Boyd, M., & Gatewood, L. (1976). An influenza simulation model for immunization studies. *American Journal of Epidemiology*, *103*(2), 152–165.
- Emanuel, E. J., & Wertheimer, A. (2006). Who should get influenza vaccine when not all can? *Science*, *312*(5775), 854–855.
- Federal Emergency Management Agency. (2021a). *FEMA COVID-19 vaccination update*. Retrieved from <https://www.fema.gov/press-release/20210402/fema-covid-19-vaccination-update>
- Federal Emergency Management Agency. (2021b). *FEMA supporting vaccination centers nationwide*. Retrieved from <https://www.fema.gov/press-release/20210226/fema-supporting-vaccination-centers-nationwide>.
- Federgruen, A., & Yang, N. (2009). Competition under generalized attraction models: Applications to quality competition under yield uncertainty. *Management Science*, *55*(12), 2028–2043.
- Florindo, H. F., Kleiner, R., Vaskovich-Koubi, D., Acúrcio, R. C., Carreira, B., Yeini, E., Tiram, G., Liubomirski, Y., & Satchi-Fainaro, R. (2020). Immune-mediated approaches against COVID-19. *Nature Nanotechnology*, *15*, 630–645.
- Goyal, P., Choi, J. J., Pinheiro, L. C., Schenck, E. J., Chen, R., Jabri, A., Satlin, M. J., Campion, T. R., Jr., Nahid, M., Ringel, J. B., Hoffman, K. L., Alshak, M. N., Li, H. A., Wehmeyer, G. T., Rajan, M., Reshetnyak, E., Hupert, N., Horn, E. M., Martinez, F. J., ... Safford, M. M. (2020). Clinical characteristics of Covid-19 in New York City. *New England Journal of Medicine*, *382*, 2372–2374.
- Graham, B. S. (2020). Rapid COVID-19 vaccine development. *Science*, *368*(6494), 945–946.
- Grauer, J., Löwen, H., & Liebchen, B. (2020). Strategic spatiotemporal vaccine distribution increases the survival rate in an infectious disease like covid-19. *Scientific Reports*, *10*(1), 1–10.
- Gurobi Optimization. (2020). *Gurobi optimizer reference manual*. Retrieved <http://www.gurobi.com>.
- Jacobson, S. H., Sewell, E. C., & Proano, R. A. (2006). An analysis of the pediatric vaccine supply shortage problem. *Health Care Management Science*, *9*(4), 371–389.
- Jacobson, S. H., Sewell, E. C., Proano, R. A., & Jokela, J. A. (2006). Stockpile levels for pediatric vaccines: How much is enough? *Vaccine*, *24*(17), 3530–3537.
- Keeling, M. J., & Shattock, A. (2012). Optimal but unequitable prophylactic distribution of vaccine. *Epidemics*, *4*(2), 78–85.
- Khamsi, R. (2020). If a coronavirus vaccine arrives, can the world make enough. *Nature*, *580*(7805), 578–580.
- Kissler, S. M., Tedijanto, C., Goldstein, E., Grad, Y. H., & Lipsitch, M. (2020). Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. *Science*, *368*, 860–868.
- Kornish, L. J., & Keeney, R. L. (2008). Repeated commit-or-defer decisions with a deadline: The influenza vaccine composition. *Operations Research*, *56*(3), 527–541.
- Lee, S., Golinski, M., & Chowell, G. (2012). Modeling optimal age-specific vaccination strategies against pandemic influenza. *Bulletin of Mathematical Biology*, *74*(4), 958–980.
- Lemmens, S., Decouttere, C., Vandaele, N., & Bernuzzi, M. (2016). A review of integrated supply chain network design models: Key issues for vaccine supply chains. *Chemical Engineering Research and Design*, *109*, 366–384.
- Li, M. L., Tazi Bouardi, H., Skali Lami, O., Trikalinos, T. A., Trichakis, N. K., & Bertsimas, D. (2020). Forecasting COVID-19 and analyzing the effect of government interventions. *medRxiv*. Retrieved from <https://www.medrxiv.org/content/early/2020/06/24/2020.06.23.20138693>
- Longini, I. M., Jr., Ackerman, E., & Elveback, L. R. (1978). An optimization model for influenza a epidemics. *Mathematical Biosciences*, *38*(1–2), 141–157.
- Lurie, N., Saville, M., Hatchett, R., & Halton, J. (2020). Developing covid-19 vaccines at pandemic speed. *New England Journal of Medicine*, *382*(21), 1969–1973.
- Mak, H. Y., Dai, T., & Tang, C. S. (2021). *Managing two-dose covid-19 vaccine rollouts with limited supply*. Available at SSRN 3790836.
- Mamani, H., Chick, S. E., & Simchi-Levi, D. (2013). A game-theoretic model of international influenza vaccination coordination. *Management Science*, *59*(7), 1650–1670.
- Matrajt, L., Eaton, J., Leung, T., & Brown, E. R. (2020). Vaccine optimization for covid-19: Who to vaccinate first? *Science Advances*, *7*, <https://advances.sciencemag.org/content/7/16/eabf1374>.
- Matrajt, L., Halloran, M. E., & Longini, I. M., Jr. (2013). Optimal vaccine allocation for the early mitigation of pandemic influenza. *PLoS Computational Biology*, *9*(3), e1002964.
- Medlock, J., & Galvani, A. P. (2009). Optimizing influenza vaccine distribution. *Science*, *325*(5948), 1705–1708.
- Muriel, J. J., & Bauchner, H. (2021). Vaccine distribution—Equity left behind? *JAMA*, *325*(9), 829–830.

- National Academies of Sciences, Engineering, and Medicine. (2020). *Framework for equitable allocation of COVID-19 vaccine*. New York Times (2020). Coronavirus in the U.S.: Latest Map and Case Count. *New York Times*. Retrieved from <https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html>
- Patel, R., Longini, I. M., Jr., & Halloran, M. E. (2005). Finding optimal vaccination strategies for pandemic influenza using genetic algorithms. *Journal of Theoretical Biology*, 234(2), 201–212.
- Perkins, A., & Espana, G. (2020). *Notredame-fred covid-19 forecasts*.
- Petrilli, C. M., Jones, S. A., Yang, J., Rajagopalan, H., O'Donnell, L. F., Chernyak, Y., Tobin, K., Cerfolio, R. J., Francois, F., & Horwitz, L. I. (2020). Factors associated with hospitalization and critical illness among 4,103 patients with COVID-19 disease in New York City. *medRxiv*.
- Rastegar, M., Tavana, M., Meraj, A., & Mina, H. (2021). An inventory-location optimization model for equitable influenza vaccine distribution in developing countries during the covid-19 pandemic. *Vaccine*, 39(3), 495–504.
- Rodriguez, A., Tabassum, A., Cui, J., Xie, J., Ho, J., Agarwal, P., Adhikari, B., & Prakash, B. A. (2020). Deepcovid: An operational deep learning-driven framework for explainable real-time covid-19 forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 15393–15400. <https://ojs.aaai.org/index.php/AAAI/article/view/17808>.
- Shin, M. D., Shukla, S., Chung, Y. H., Beiss, V., Chan, S. K., Ortega-Rivera, O. A., Wirth, D. M., Chen, A., Sack, M., Pokorski, J. K., et al. (2020). COVID-19 vaccine development and a potential nanomaterial path forward. *Nature Nanotechnology*, 15, 646–655.
- Sun, P., Yang, L., & De Véricourt, F. (2009). Selfish drug allocation for containing an international influenza pandemic at the onset. *Operations Research*, 57(6), 1320–1332.
- Tanner, M. W., Sattenspiel, L., & Ntamo, L. (2008). Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Mathematical Biosciences*, 215(2), 144–151.
- Teytelman, A., & Larson, R. C. (2013). Multiregional dynamic vaccine allocation during an influenza epidemic. *Service Science*, 5(3), 197–215.
- U.S. Census Bureau. (2020). *County population by characteristics: 2010–2019*. Retrieved from <https://www.census.gov/data/tables/time-series/demo/popest/2010s-counties-detail.html>
- U.S. Center for Disease Control. (2020a). *COVID-19 forecasts*. Retrieved from <https://www.cdc.gov/coronavirus/2019-ncov/covid-data/forecasting-us.html>
- U.S. Center for Disease Control. (2020b). *COVID-19 hospitalization and death by age*. <https://www.cdc.gov/coronavirus/2019-ncov/covid-data/investigations-discovery/hospitalization-death-by-age.html>
- U.S. Center for Disease Control. (2021). *Provisional death counts for coronavirus disease 2019 (COVID-19)*. Retrieved from https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly
- Uribe-Sánchez, A., Savachkin, A., Santana, A., Prieto-Santa, D., & Das, T. K. (2011). A predictive decision-aid methodology for dynamic mitigation of influenza pandemics. *OR Spectrum*, 33(3), 751–786.
- Watson, R. K. (1972). On an epidemic in a stratified population. *Journal of Applied Probability*, 9(3), 659–666.
- Wj, G., Ni Zy, H. Y., Liang, W., Cq, O., He, J., Liu, L., Shan, H., Cl, L., Hui, D. S., et al. (2020). Clinical characteristics of coronavirus disease 2019 in China. *New England Journal of Medicine*, 382(18), 1708–1720.
- Wu, J. T., Wein, L. M., & Perelson, A. S. (2005). Optimization of influenza vaccine selection. *Operations Research*, 53(3), 456–476.
- Yarmand, H., Ivy, J. S., Denton, B., & Lloyd, A. L. (2014). Optimal two-phase vaccine allocation to geographically different regions under uncertainty. *European Journal of Operational Research*, 233(1), 208–219.

How to cite this article: Bertsimas, D., Digalakis Jr, V., Jacquillat, A., Li, M. L., & Previero, A. (2022). Where to locate COVID-19 mass vaccination facilities?. *Naval Research Logistics (NRL)*, 69(2), 179–200. <https://doi.org/10.1002/nav.22007>

APPENDIX A

Table A1 presents the proposed center allocation; for each state, we show the number of centers allocated, the number of vaccines per day, and the cities where the proposed centers are located.

In Table A2, we outline the realized center allocations which has so far been implemented by the Federal Emergency Management Agency (FEMA) (Federal Emergency Management Agency, 2021a, 2021b). In Table A3, we compare the realized center allocation (R) with the proposed center allocation (Pr). For the latter, we calculate percentages by taking into account only the 48 centers that were allocated to the 13 states under consideration. We also show the population-based center allocation (Pop), for which we calculate percentages by taking into account the 54 centers that were allocated to the 13 states under consideration. The last column in Table A3 indicates whether (R) and (Pr) agree, that is, whether the fraction of centers allocated by both to any given state is greater/equal/less than the population share of the state.

Similar to Tables A3 and A4 compare the number of vaccines per day (v/d, in thousands) allocated to each state in the realized solution (R) with that in the proposed solution (Pr) and the population-based baseline solution (Pop). For (Pop), we take the (Pop) center allocation shown in Table A3 and assume that the vaccines are distributed uniformly across centers. More specifically, the number of vaccines that each center gets is equal to the average number of vaccines allocated to each center in (R), namely, in the realized solution.

Both tables show that, despite obvious differences given the rollout plan and a number of practical considerations, the model-based recommendations are generally in line with the center and vaccine allocation decisions that were made in practice. Specifically, the realized and proposed solutions agree in the majority of cases, both in terms of centers and in terms of vaccines. These results underscore the role of our model as a strategic tool to support the deployment of vaccination sites as well as a tactical tool to support vaccine distribution.

Table A1 Proposed center allocation

State	# centers	# vaccines/day	Selected cities
Alabama	2	14 458	Birmingham, Mobile
Alaska	1	6592	Anchorage
Arizona	2	29 663	Mesa, Scottsdale
Arkansas	1	14 094	Little Rock
California	8	52 734	Bakersfield, Chula Vista, Citrus Heights, Fremont, Fresno Glendale, Indio, Irvine
Colorado	2	28 289	Centennial, Colorado Springs
Connecticut	1	8699	Waterbury
Delaware	1	6592	Wilmington
District of Columbia	1	6592	Washington
Florida	6	88 989	Boca Raton, Jacksonville, Miami, Orlando, Tallahassee, Tampa
Georgia	2	13 357	Sandy Springs, Warner Robins
Hawaii	1	6592	Honolulu
Idaho	1	6592	Meridian
Illinois	2	13 184	Cicero, Springfield
Indiana	2	20 807	Indianapolis, South Bend
Iowa	2	16 356	Cedar Rapids, Des Moines
Kansas	1	6592	Topeka
Kentucky	1	6592	Louisville
Louisiana	2	29 386	Baton Rouge, Shreveport
Maine	1	6592	Portland
Maryland	1	12 167	Baltimore
Massachusetts	1	11 756	Lowell
Michigan	2	13 184	Farmington Hills, Wyoming
Minnesota	1	6592	Brooklyn Park
Mississippi	1	8991	Jackson
Missouri	2	19 139	Lee's Summit, O'Fallon
Montana	1	6592	Missoula
Nebraska	1	6592	Omaha
Nevada	2	13 184	Henderson, Sparks
New Hampshire	1	6592	Nashua
New Jersey	2	29 663	Camden, Newark
New Mexico	1	11 670	Rio Rancho
New York	5	70 791	Buffalo, Rochester, Schenectady, Syracuse, Yonkers
North Carolina	2	13 184	Cary, Charlotte
North Dakota	1	6592	Fargo
Ohio	3	19 775	Cincinnati, Cleveland, Columbus
Oklahoma	2	20 835	Broken Arrow, Edmond
Oregon	1	13 395	Portland
Pennsylvania	2	13 184	Pittsburgh, Reading
Rhode Island	1	6592	Pawtucket
South Carolina	1	10 979	Columbia
South Dakota	1	6592	Sioux Falls
Tennessee	3	19 775	Knoxville, Memphis, Nashville
Texas	11	163 146	Austin, Corpus Christi, Dallas, El Paso, Fort Worth, Houston, Lubbock, McAllen, Midland, San Antonio, Tyler
Utah	1	10 818	West Valley City
Vermont	1	6804	Burlington
Virginia	3	22 841	Alexandria, Newport News, Roanoke
Washington	2	26 447	Kent, Spokane
West Virginia	1	8370	Charleston
Wisconsin	2	13 184	Appleton, Waukesha
Wyoming	1	6592	Cheyenne

Table A2 Realized center allocation

State	# centers	# vaccines/day	Selected cities
California	2	12 000	Oakland, Los Angeles
Florida	4	12 000	Miami, Tampa, Orlando, Jacksonville
Illinois	1	6000	Chicago
Indiana	1	3000	Gary
Massachusetts	1	5000	Boston
Missouri	1	3000	St. Louis
New Jersey	1	5000	Newark
New York	6	12 000	Brooklyn, Queens, Yonkers, Rochester, Buffalo, Albany
North Carolina	1	3000	Greensboro
Pennsylvania	1	6000	Philadelphia,
Texas	3	12 000	Houston, Dallas, Arlington
Virginia	1	5000	Norfolk
Washington	1	1000	Yakima

Table A3 Comparison between realized and proposed center allocations

State	Realized		Proposed		Population		(R) ⇔ (Pr)
	#	%	#	%	#	%	
California	2	0.08	8	0.17	12	0.22	1
Florida	4	0.17	6	0.13	6	0.11	1
Illinois	1	0.04	2	0.04	4	0.07	1
Indiana	1	0.04	2	0.04	2	0.04	1
Massachusetts	1	0.04	1	0.02	2	0.04	0
Missouri	1	0.04	2	0.04	1	0.02	1
New Jersey	1	0.04	2	0.04	2	0.04	1
New York	6	0.25	5	0.1	6	0.11	0
North Carolina	1	0.04	2	0.04	3	0.06	1
Pennsylvania	1	0.04	2	0.04	4	0.07	1
Texas	3	0.13	11	0.23	8	0.15	0
Virginia	1	0.04	3	0.06	2	0.04	0
Washington	1	0.04	2	0.04	2	0.04	1
Total	24	0.99	48	0.99	54	1.01	9

Table A4 Comparison between realized and proposed vaccine allocations

State	Realized		Proposed		Population		(R) ⇔ (Pr)
	#	%	#	%	#	%	
California	12	0.14	53	0.1	42.5	0.22	1
Florida	12	0.14	89	0.16	21.25	0.11	1
Illinois	6	0.07	13	0.02	14.167	0.07	0
Indiana	3	0.04	21	0.04	7.083	0.04	1
Massachusetts	5	0.06	12	0.02	7.083	0.04	0
Missouri	3	0.04	19	0.03	3.542	0.02	1
New Jersey	5	0.06	30	0.05	7.083	0.04	1
New York	12	0.14	71	0.13	21.25	0.11	1
North Carolina	3	0.04	13	0.02	10.625	0.06	1
Pennsylvania	6	0.07	13	0.02	14.167	0.07	0
Texas	12	0.14	163	0.3	28.333	0.15	0
Virginia	5	0.06	23	0.04	7.083	0.04	0
Washington	1	0.01	26	0.05	7.083	0.04	0
Total	85	1.01	546	0.98	191.249	1.01	7