

Resource Distribution Under Spatiotemporal Uncertainty of Disease Spread: Stochastic versus Robust Approaches

Beste Basciftci* Xian Yu† Siqian Shen‡

Abstract

We consider the problem of optimizing locations of distribution centers (DCs) and plans for distributing resources such as test kits and vaccines, under spatiotemporal uncertainties of disease spread and demand for the resources. We aim to balance the operational cost (including costs of deploying facilities, shipping, and storage) and quality of service (reflected by demand coverage), while ensuring equity and fairness of resource distribution across multiple populations. We compare a sample-based stochastic programming (SP) approach with a distributionally robust optimization (DRO) approach using a moment-based ambiguity set. Numerical studies are conducted on instances of distributing COVID-19 vaccines in the United States and test kits, to compare SP and DRO models with a deterministic formulation using estimated demand and with the current resource distribution plans implemented in the US. We demonstrate the results over distinct phases of the pandemic to estimate the cost and speed of resource distribution depending on scale and coverage, and show the “demand-driven” properties of the SP and DRO solutions. Our results further indicate that if the worst-case unmet demand is prioritized, then the DRO approach is preferred despite of its higher overall cost. Nevertheless, the SP approach can provide an intermediate plan under budgetary restrictions without significant compromises in demand coverage.

Keywords: COVID-19 pandemic; vaccine distribution; resource allocation; stochastic integer programming; distributionally robust optimization; multi-objective optimization

1 Introduction

With the rapid spread of the coronavirus disease 2019 (COVID-19), testing is central to planning response activities in all countries during and post the pandemic. Effective and efficient testing can lead to early outbreak detection, to quickly isolate and treat infected patients, guide people consciously performing social distancing, and also lock down certain areas/activities if needed.

*Corresponding author; Department of Business Analytics, Tippie College of Business, University of Iowa, Email: beste-basciftci@uiowa.edu;

†Department of Industrial and Operations Engineering, University of Michigan at Ann Arbor, Email: yuxian@umich.edu;

‡Department of Industrial and Operations Engineering, University of Michigan at Ann Arbor, Email: siqian@umich.edu.

Establishing efficient testing systems involves distributing test kits to test centers while considering potential demand uncertainty for testing (see, e.g., Lampariello and Sagratella, 2021; Santini, 2021). A similar resource distribution problem arises in the production and distribution of COVID-19 vaccines that became available in December 2020 – That is, upon needs and orders from different regions, a central government sends limited amounts of vaccines to local agents using the information of regional infection status. This problem was solved in the United States (US) during the first two quarters of Year 2021 for distributing COVID-19 vaccines and is relevant to many countries when vaccines become available globally, or for distributing resources to respond to future outbreaks of other diseases.

In the US, policymakers seek efficient ways to distribute vaccines to states and jurisdictions, and then to local hospitals, clinics, pharmacies, schools, communities and so on. Given certain vaccine allocation policies, how to distribute vaccines from production sites or distribution centers (DCs) to downstream demand is challenging and emergent. Meanwhile, it is hard to quantify the cost and speed trade-off for distributing medical resources such as vaccines or test kits, as there exist hidden costs associated with unsatisfied demand that may come from high-risk population groups. Given the evolving pandemic and rapidly changing demand, it is also challenging to locate DCs and make shipping plans optimally to achieve the best trade-off. Moreover, the scale of the problem can be enormous, involving millions of vaccines (or test kits) and large-scale distributed demand for operating the system for the whole country. A scientific way to conduct vaccine and test-kit distribution needs to rely on mathematical models and algorithms for processing and utilizing information from large datasets about the COVID-19 infection trends and demand for testing/vaccination.

In the supply chain literature, there has been a large body of research studies on how to site retail stores, optimize inventory and production, manage stock levels as demand for products fluctuates by season (Daskin, 2011; Shen et al., 2011; Snyder, 2006; Basciftci et al., 2021). Among them, many focus on facility locations under uncertain demand, using stochastic or robust optimization approaches. Establishing a COVID-19 resource distribution system presents a set of similar challenges under various uncertainties, but in addition, it requires to incorporate more decisions such as resource distribution, inventory control, and demand shortage control into the multi-period facility planning problem. In this work, we present a mathematical framework that encapsulates allocation and distribution stages of disease-control resources (e.g., vaccines and test kits) by considering spatiotemporal uncertainty in the demand and ambiguity of the probability distribution of demand over a multi-period horizon. We optimize the locations of DCs, their capacities, shipment amounts and inventory levels. As the provided framework is presented in a generic form, it aims to address various resource allocation problems during different phases of a pandemic by comparing deterministic, stochastic and distributionally robust decision-making approaches.

The main contributions of the paper are threefold. First, we combine mathematical programming and statistical learning, for optimally locating DCs for disease-control resources and deriving shipping plans, customized for counties and states with diverse demographics and disease spread patterns. Our research can be utilized at the national level to balance the distribution, or at the state or county level to facilitate local operations. Second, through data-informed location opti-

mization, we can effectively identify the most critical and vulnerable groups to prioritize testing or receiving vaccines under constrained amounts of resources and as a result, can protect other population groups. Third, we conduct extensive numerical studies using real COVID-19 infection data in the US and in the State of Michigan, to compare different approaches in terms of the operational cost and speed of resource distribution, by testing a diverse set of instances having different scales and parameter settings.

The rest of the paper is organized as follows. In Section 2, we review the most relevant literature in disease control, facility location, and vaccine supply chain management. In Section 3, we present the stochastic programming model based on a set of demand samples generated from a given probability distribution and in Section 4, we consider that the exact demand distribution is unknown and present the distributionally robust optimization model to optimize decisions against the distributional ambiguity. In Section 5, we conduct numerical studies using COVID-19 infection data to demonstrate results of distributing vaccines and test kits, and compare our solutions with what were implemented by the Centers for Disease Control and Prevention (CDC). In Section 6, we conclude the paper and present future research directions.

2 Literature Review

Resource distribution is of vital importance in many applications, in particular the ones related to disease control and disaster relief, as the resources are usually scarce (Cao and Huang, 2012; Gupta et al., 2016). For COVID-19, personal protective equipment (PPE), test kits, hospital beds and ventilators, and, most recently, vaccines are among the resources that need to be effectively distributed at all levels (Emanuel et al., 2020). Bertsimas et al. (2021a) formulate a deterministic optimization model to improve ventilator allocation by allowing sharing of ventilators between hospitals from different states in the US. Billingham et al. (2020) study a similar problem to optimize ventilator sharing. Lampariello and Sagratella (2021) consider a single-period COVID-19 test-kit allocation problem by maximizing utility functions corresponding to disease detection capabilities in different regions, to determine the amount of test kits to be allocated in each region given a certain budget. Santini (2021) considers the distribution of swabs and reagent to laboratories for maximizing the number of COVID-19 tests processed. The author formulates the problem as a deterministic integer programming model by considering sharing of swabs and reagent among different laboratories over a multi-period planning horizon. Although most of the existing literature consider different stages of resource planning and sharing during the pandemic, they use regression tools to forecast average demand, without explicitly modeling the spatiotemporal infection and demand uncertainties at the resource-planning phase.

Assuming demand being stochastic, Mehrotra et al. (2020); Blanco et al. (2020) propose stochastic programming approaches to address this issue by generating scenarios with different patterns to represent the randomness in the amount of patients due to uncertain disease spread, to re-allocate and share medical resources among different hospitals. Yin et al. (2021) extends this problem to a risk-averse multi-stage stochastic programming setting by incorporating changing transmission

dynamics. We note that these studies consider the resource allocation and redistribution between certain regions with existing facilities, and facility location decisions with their corresponding capacities have not been incorporated to the decision making process. Recently, Parker et al. (2020) use robust optimization by assuming an unknown number of patients for each specific day within a certain range from a nominal value, e.g., the average across a certain period. In addition to the studied settings, the probability distribution of the uncertain demand for resources may be ambiguous during emergency or disaster relief operations due to the inherent and abrupt nature of these events (see, e.g., Liu et al., 2019; Wang and Chen, 2020). This motivates the development of a distributionally robust optimization approach in our paper for robustly allocating resources during an unprecedented event such as the COVID-19 pandemic. **To the best of our knowledge, this paper is the first work that leverages distributionally robust optimization for COVID-19 related medical resource distribution and facility location decision-making problems.**

The supply chain operations of vaccines for infectious disease control involve vaccine production, allocation and distribution stages, with decisions on which vaccines to produce, how many doses to produce, who should be vaccinated, how the vaccines can be distributed, and so on (see Duijzer et al., 2018). Equitable and timely allocation of vaccines to different population groups becomes necessary in order to eliminate infectious diseases worldwide (Tebbens and Thompson, 2009). In that regard, Huang et al. (2017) provide a retrospective study over 2009 H1N1 pandemic by developing deterministic optimization models to ensure fair and equitable allocation of vaccines through first determining coverage levels of each region and then finding an allocation plan within certain tolerance from the targeted levels.

Golan et al. (2020) conduct a comprehensive literature review that focuses on the resilience of vaccine supply chains, and point out that the lack of network-based, modeling-based, quantitative analysis is a major gap that needs to be bridged in order to create methods of real-time analysis and decision tools for the COVID-19 vaccine supply chain. For COVID-19 specifically, Bubar et al. (2020) discuss strategies for vaccine prioritization among different population groups; Babus et al. (2020) estimate occupation-based infection risks and use age-based infection fatality rates in a model to assign priorities over populations with different occupations and ages; Bertsimas et al. (2020) capture vaccine effects and the variability in mortality rates across sub-populations, and then integrate a predictive model into a prescriptive model to optimize vaccine allocation; Chen et al. (2020) focus on how priority and population groups should be identified over time under limited supply. Nevertheless, as the majority of the literature is focusing on COVID-19 vaccine allocation, few studies develop mathematical approaches for solving the operational problems related to vaccine distribution. Among them, Bertsimas et al. (2021b) optimize vaccine site selection and the assignment of population to different sites while ensuring optimal subsequent vaccine allocation. They show that the proposed solution achieves critical fairness objectives and is also highly robust to uncertainties and forecasting errors. Different from most of the existing COVID-19 vaccine allocation studies, in this paper, we study a broader class of vaccine-related operational problems involves locating facilities to produce or store vaccines and shipping them to demand locations in a daily basis with inventory and unmet demand (i.e., lost sales or backlogging) being considered.

This study is closely related to the literature of facility location under uncertainty, which uses different approaches depending on the available information about the uncertain demand (Snyder, 2006). When the distribution of the uncertain parameter is known or can be estimated accurately, the facility location problem can be modeled as a stochastic program by optimizing the expected total cost in its risk-neutral setting (Shen et al., 2011; Bidhandi and Yusuff, 2011; Tolooie et al., 2020). In contrast to the generic stochastic facility location models, in certain problem settings as in this study, operational decisions may need to be integrated into this strategic level problem (Schütz et al. (2009); Georgiadis et al. (2011)). This study extends this line of research by determining facility location and capacity decisions with inventory, shipment and unmet demand amounts within a two-stage stochastic programming framework, while providing extensions of this model to address more complex settings that can involve inventory at distribution centers, lead time of shipping and different types of distribution centers. On the other hand, when the distribution of the uncertain parameter is not fully known (i.e. there exists ambiguity in the distribution), the decision maker can be more conservative by optimizing the problem over the worst-case distribution, which can be modeled using distributionally robust optimization (Lu et al., 2015; Basciftci et al., 2021). Although these studies consider uncertain demand when planning facilities, an integrated decision framework for optimizing the facility location decisions together with multi-period capacity allocation, inventory control and resource distribution plans has not been studied under distributional ambiguity of the multivariate demand, due to the modeling and computational complexity. Thus, this paper also presents contributions to the facility location literature by addressing this complex setting with a distributionally robust optimization approach.

3 A Stochastic Programming Approach

In this paper, we consider the distribution of vaccines or test kits in a given region as a capacitated facility location problem involving multiple periods of shipment planning. A decision maker may need to solve the problem for either national- or state-level operations, to distribute manufactured vaccines and test kits to different states or counties from located DCs. Here, the DCs could also be the manufacturing sites, as manufactures can dispatch productions to states or to counties directly. In addition, the federal or state governments can act to open up new DCs to improve the efficiency of operations and increase demand coverage.

Denote $\mathcal{I}$, $\mathcal{J}$, $\mathcal{T}$ as the sets of potential sites for locating DCs, demand locations and finite periods, respectively. Let c_i^o , c_i^h , c_{ijt}^s , c_{jt}^u , c_{jt}^I be the cost of operating DC i , unit cost of installing capacities in DC i , unit shipping cost from DC i to demand site j in period t , unit penalty cost of unsatisfied demand and unit cost of inventory at demand site j in period t , for all $i \in \mathcal{I}$, $j \in \mathcal{J}$, $t \in \mathcal{T}$, respectively. (Varying cost parameter c_{jt}^u can help ensure fair resource distribution to prioritized demand locations based on their demographics and infection status over time.) Denote B_t as the total capacity of manufacturing resources across all DCs in period t for all $t \in \mathcal{T}$, determined by the total amount of raw materials, space, workers, etc., needed for manufacturing the resources during each period. Let $\mathbf{d}$ be the vector of uncertain demand (i.e., the number of people who

need to be vaccinated or tested) and P be its probability distribution. We use the Monte Carlo sampling approach (Kleywegt et al., 2002) to replace P with an empirical distribution constructed using $|\Omega|$ scenarios with each scenario $\omega \in \Omega$ having an equal probability $p^\omega = 1/|\Omega|$. We consider a finite set Ω of realizations of the random vector $\mathbf{d}$. Specifically, for each scenario $\omega \in \Omega$, we use $d_{jt}(\omega)$ to represent the demand realization at site j in period t for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$, and therefore $\mathbf{d} = [d_{jt}(\omega), \omega \in \Omega, j \in \mathcal{J}, t \in \mathcal{T}]^\top$. Throughout this paper, bold letters are used for representing vectors and matrices. For notational convenience, we also use notation $[m]$ to indicate set $\{1, \dots, m\}$.

We define binary variables $x_i \in \{0, 1\}$, $\forall i \in \mathcal{I}$ such that $x_i = 1$ if DC i is built, and $x_i = 0$ otherwise. For each built DC i and period $t \in \mathcal{T}$, we also decide its capacity $h_{it} \geq 0$ for manufacturing or storing resources. Both variables $\mathbf{x}$ and $\mathbf{h}$ are planning decisions and their values need as determined before realizing uncertain demand $\mathbf{d}$. For each scenario $\omega \in \Omega$, we define variables $s_{ijt}(\omega) \geq 0$ as the amount of resources sent from DC i to demand location j in period t , for all $i \in \mathcal{I}$, $j \in \mathcal{J}$, $t \in \mathcal{T}$. For each demand site $j \in \mathcal{J}$, we allow to keep inventory if the received resources are more than the total demand or to back order otherwise. In particular, we assume that people who cannot receive vaccines or test kits in the current period will wait to be administered in future periods. Accordingly, we define variables $I_{jt}(\omega) \geq 0$ and $u_{jt}(\omega) \geq 0$ as the inventory and backlog recourse variables for each $j \in \mathcal{J}$, $t \in \mathcal{T}$, and $\omega \in \Omega$. Moreover, input parameters I_{j0} and u_{j0} denote the initial inventory and backlog at demand location j for all $j \in \mathcal{J}$, whose values are the same across all scenarios.

We first employ a two-stage stochastic mixed-integer linear programming (SMIP) framework to formulate the problem, where in the first stage, we decide values of variables x_i and h_{it} for all $i \in \mathcal{I}$, $t \in \mathcal{T}$. In the second stage, given each demand value $d_{jt}(\omega)$, we optimize the corresponding shipping and inventory plans using variables $s_{ijt}(\omega)$, $I_{jt}(\omega)$ and $u_{jt}(\omega)$ for each scenario $\omega \in \Omega$.

The SMIP model is given by:

$$\min \sum_{i \in \mathcal{I}} c_i^o x_i + \sum_{i \in \mathcal{I}, t \in \mathcal{T}} c_i^h h_{it} + \sum_{\omega \in \Omega} p^\omega \left(\sum_{i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}} c_{ijt}^s s_{ijt}(\omega) + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} (c_{jt}^I I_{jt}(\omega) + c_{jt}^u u_{jt}(\omega)) \right) \quad (1a)$$

$$\text{s.t. } h_{it} \leq M_i x_i, \forall i \in \mathcal{I}, t \in \mathcal{T}, \quad (1b)$$

$$\sum_{i \in \mathcal{I}} h_{it} \leq B_t, \forall t \in \mathcal{T}, \quad (1c)$$

$$\sum_{j \in \mathcal{J}} s_{ijt}(\omega) \leq h_{it}, \forall i \in \mathcal{I}, t \in \mathcal{T}, \omega \in \Omega, \quad (1d)$$

$$\sum_{i \in \mathcal{I}} s_{ijt}(\omega) + I_{j(t-1)}(\omega) + u_{jt}(\omega) = d_{jt}(\omega) + I_{jt}(\omega) + u_{j(t-1)}(\omega), \forall j \in \mathcal{J}, t \in \mathcal{T}, \omega \in \Omega, \quad (1e)$$

$$x_i \in \{0, 1\}, h_{it}, s_{ijt}(\omega), I_{jt}(\omega), u_{jt}(\omega) \geq 0, \forall i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}, \omega \in \Omega, \quad (1f)$$

where the objective function (1a) minimizes the total cost of opening DCs, installing their capacities, and the total expected cost of shipping, holding inventory and back-orders. Constraints (1b) prohibit assigning any capacity to a DC that is not in use, where M_i is the capacity limit of the DC i for all $i \in \mathcal{I}$. Constraints (1c) impose temporal limitation on the total production capacity

of all DCs. Constraints (1d) link the first-stage variables with the second-stage recourse decisions such that the total shipment from each DC is no more than its installed capacity during any period in any scenario. Constraints (1e) are “flow-balance” constraints to reflect the changes of inventory and back-order levels, depending on the amount of resources received and demand level at each site j , for each period t in each scenario ω . Constraints (1f) require binary valued x -variables and set all the other variables to be non-negative.

3.1 Model extension I: Inventory at DCs

When the costs of production capacity (c_{it}^h) and shipment (c_{ijt}^s) are time-independent, it is without loss of generality to assume that there is no inventory kept in DCs because those additional products can be shipped to and stored at customer sites. However, if these costs vary over time, it may be beneficial to keep some products in DCs and ship them in the future. In this case, we define additional variables $I_{it}^D(\omega)$ as the inventory at DC i in period t with scenario ω for all $i \in \mathcal{I}$, $t \in \mathcal{T}$, $\omega \in \Omega$. Accordingly, Constraints (1d) can be modified as follows:

$$\sum_{j \in \mathcal{J}} s_{ijt}(\omega) + I_{it}^D(\omega) = h_{it} + I_{i(t-1)}^D(\omega), \quad \forall i \in \mathcal{I}, t \in \mathcal{T}, \omega \in \Omega,$$

where $I_{i0}^D(\omega) = I_{i0}^D$, $\forall \omega \in \Omega$ is the (given) initial inventory at DC i . Moreover, the objective function (1a) can be recast as

$$\min \sum_{i \in \mathcal{I}} c_i^o x_i + \sum_{i \in \mathcal{I}, t \in \mathcal{T}} c_{it}^h h_{it} + \sum_{\omega \in \Omega} p^\omega \left(\sum_{i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}} c_{ijt}^s s_{ijt}(\omega) + \sum_{i \in \mathcal{I}, t \in \mathcal{T}} c_{it}^{ID} I_{it}^D(\omega) + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} (c_{jt}^I I_{jt}(\omega) + c_{jt}^u u_{jt}(\omega)) \right),$$

where c_{it}^{ID} is the unit inventory cost at DC i in period t for all $i \in \mathcal{I}$, $t \in \mathcal{T}$.

3.2 Model extension II: Lead time for shipping

The SMIP model (1) assumes that there is no lead time when shipping from DCs to demand locations. In the case of a constant lead time for every DC and demand-location pairs, we can simply shift the optimal production and shipment plans accordingly. However, if the lead time varies by location, it is worthwhile to add lead time into our model. We denote L_{ij} as the lead time in sending resources from DC i to customer location j for all $i \in \mathcal{I}$, $j \in \mathcal{J}$. Then, Constraints (1e) can be adjusted to

$$\sum_{i \in \mathcal{I}: t \geq L_{ij}} s_{ij(t-L_{ij})}(\omega) + I_{jt-1}(\omega) + u_{jt}(\omega) = d_{jt}(\omega) + I_{jt}(\omega) + u_{jt-1}(\omega), \quad \forall j \in \mathcal{J}, t \in \mathcal{T}, \omega \in \Omega. \quad (2)$$

3.3 Model extension III: Different types of DCs

In practice, the total demand could be satisfied by multiple types of resources, and the DCs may have different processing, stocking and distributing costs, depending on their specialties. In such a setting, one can differentiate DCs in terms of their unit production, shipping and inventory costs

and capacities by the types of resources they hold and distribute. For this purpose, we define $\mathcal{L}$ as the set of resource types with different requirements and let x_{il} variable to be 1 if DC i of type l is built, and 0 otherwise. Similarly, we extend capacity h_{itl} , shipment $s_{ijtl}(\omega)$, inventory $I_{jtl}(\omega)$, unmet demand $u_{jtl}(\omega)$ decision variables and their associated costs to this setting by incorporating the type index l . We further introduce a new decision variable $\bar{d}_{jtl}(\omega)$ for representing the amount of the total demand in location j at time t satisfied by the product type l . The resulting extension of SMIP model (1) is given by:

$$\min \sum_{i \in \mathcal{I}, l \in \mathcal{L}} c_{il}^o x_{il} + \sum_{i \in \mathcal{I}, t \in \mathcal{T}, l \in \mathcal{L}} c_{il}^h h_{itl} + \sum_{\omega \in \Omega} p^\omega \left(\sum_{i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}, l \in \mathcal{L}} c_{ijtl}^s s_{ijtl}(\omega) + \sum_{j \in \mathcal{J}, t \in \mathcal{T}, l \in \mathcal{L}} (c_{jtl}^I I_{jtl}(\omega) + c_{jtl}^u u_{jtl}(\omega)) \right) \quad (3a)$$

$$\text{s.t. } h_{itl} \leq M_{il} x_{il}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}, l \in \mathcal{L}, \quad (3b)$$

$$\sum_{i \in \mathcal{I}, l \in \mathcal{L}} h_{itl} \leq B_t, \quad \forall t \in \mathcal{T}, \quad (3c)$$

$$\sum_{j \in \mathcal{J}} s_{ijtl}(\omega) \leq h_{itl}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}, l \in \mathcal{L}, \omega \in \Omega, \quad (3d)$$

$$\sum_{i \in \mathcal{I}} s_{ijtl}(\omega) + I_{j(t-1)l}(\omega) + u_{jtl}(\omega) = \bar{d}_{jtl}(\omega) + I_{jtl}(\omega) + u_{j(t-1)l}(\omega), \quad \forall j \in \mathcal{J}, t \in \mathcal{T}, l \in \mathcal{L}, \omega \in \Omega, \quad (3e)$$

$$\sum_{l \in \mathcal{L}} \bar{d}_{jtl}(\omega) = d_{jt}(\omega), \quad \forall j \in \mathcal{J}, t \in \mathcal{T}, \omega \in \Omega, \quad (3f)$$

$$x_{il} \in \{0, 1\}, h_{itl}, s_{ijtl}(\omega), I_{jtl}(\omega), u_{jtl}(\omega), \bar{d}_{jtl}(\omega) \geq 0, \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, l \in \mathcal{L}, t \in \mathcal{T}, \omega \in \Omega. \quad (3g)$$

Constraints (3e) satisfy the inventory and shipment relationship for each demand location, planning period, product type and demand realization, whereas Constraints (3f) ensure that the total demand is partitioned into different types of resources.

4 A Distributionally Robust Optimization Approach

The stochastic programming approach introduced in Section 3 assumes that the distribution of the underlying uncertainty is perfectly known and one can have access to a large amount of samples from the true distribution. However, these may not be true in practice, specifically during a pandemic, where the distribution may be misspecified under limited information. To model this type of distributional ambiguity, we consider a distributionally robust optimization model, in which optimal solutions are sought for the worst-case probability distribution within a family of candidate distributions, called an ‘‘ambiguity set’’ and denoted by $\mathcal{P}$. We denote the random demand vector by $\boldsymbol{\xi} = [d_{jt}, j \in \mathcal{J}, t \in \mathcal{T}]^\top$ and the unknown probability distribution by $\mathbb{P}$. A distributionally robust analogous model of the SMIP model (1) is:

$$\min \sum_{i \in \mathcal{I}} c_i^o x_i + \sum_{i \in \mathcal{I}, t \in \mathcal{T}} c_i^h h_{it} + \max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}[g(\mathbf{h}, \boldsymbol{\xi})] \quad (4a)$$

$$\text{s.t. } h_{it} \leq M_i x_i, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}, \quad (4b)$$

$$\sum_{i \in \mathcal{I}} h_{it} \leq B_t, \forall t \in \mathcal{T}, \quad (4c)$$

$$x_i \in \{0, 1\}, h_{it} \geq 0, \forall i \in \mathcal{I}, t \in \mathcal{T}, \quad (4d)$$

where

$$g(\mathbf{h}, \boldsymbol{\xi}) = \min \sum_{i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}} c_{ijt}^s s_{ijt} + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} (c_{jt}^I I_{jt} + c_{jt}^u u_{jt}) \quad (5a)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{J}} s_{ijt} \leq h_{it}, \forall i \in \mathcal{I}, t \in \mathcal{T}, \quad (5b)$$

$$\sum_{i \in \mathcal{I}} s_{ijt} + I_{j(t-1)} + u_{jt} = \xi_{jt} + I_{jt} + u_{j(t-1)}, \forall j \in \mathcal{J}, t \in \mathcal{T}, \quad (5c)$$

$$s_{ijt}, I_{j,t}, u_{jt} \geq 0, \forall i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}. \quad (5d)$$

The objective function (4a) minimizes the cost of opening DCs, installing their capacities, and the worst-case expected cost under the set of candidate distributions in the ambiguity set $\mathcal{P}$. The inner problem (5) corresponds to the second-stage problem and minimizes the total cost of shipping, holding inventory and backlogging given first-stage decision $\mathbf{h}$.

To express the form of uncertainty in distribution, two main classes of ambiguity sets can be defined for distributionally robust optimization models. They are (i) statistical distance-based ambiguity sets that consider distributions within a certain distance to a reference distribution (Jiang and Guan, 2016; Mohajerin Esfahani and Kuhn, 2018), and (ii) moment-based ambiguity sets that consider distributions based on moment information (Delage and Ye, 2010; Zhang et al., 2018a). Since a reference distribution is essential in constructing the former class of ambiguity sets, such a distribution might be misleading in the case of a pandemic with abrupt and unprecedented changes in events and requires further data points to ensure its accuracy with a high level of confidence. Therefore, we focus on the moment-based ambiguity sets for describing possible distributions corresponding to the underlying demand uncertainty. To construct our ambiguity set, we bound a set of moment functions of $\boldsymbol{\xi}$ by certain parameters. Specifically, we consider m different moment functions $\mathbf{f} := (f_1(\boldsymbol{\xi}), \dots, f_m(\boldsymbol{\xi}))^\top$. We assume that the random vector $\boldsymbol{\xi}$ has a finite support set S containing possible realizations $S = \{\boldsymbol{\xi}^1, \dots, \boldsymbol{\xi}^K\}$, for a given distribution $\mathbb{P}$. [In this paper, we consider a finite support set for computational tractability. If using continuous support sets, the reformulation of the problem becomes as a semi-infinite integer program with an infinite number of constraints and mixed-integer binary variables, and thus cannot be optimized directly using the state-of-the-art solvers. For example, Zhang et al. \(2018b\) develop approximation algorithms for solving 0-1 semi-definite program for the problem of distributionally robust surgery planning and we leave the continuous support case for our problem for future research but will test different sizes of the discrete support case to justify its proximity and result sensitivity.](#)

Consequently, we reformulate $\mathbb{E}[g(\mathbf{h}, \boldsymbol{\xi})]$ as $\sum_{k \in [K]} p_k g(\mathbf{h}, \boldsymbol{\xi}^k)$ and the inner problem (5) can be represented for each realization from the set $\{\boldsymbol{\xi}^1, \dots, \boldsymbol{\xi}^K\}$. Then for each $s \in [m]$, the corresponding moment function $f_s(\boldsymbol{\xi}) = \prod_{j \in \mathcal{J}, t \in \mathcal{T}} \xi_{jt}^{k_{sjt}}$, where k_{sjt} is a non-negative integer indicating the power of ξ_{jt} for the s -th moment function. The lower and upper bounds are defined by $\mathbf{l} := (l_1, \dots, l_m)^\top$

and $\mathbf{u} := (u_1, \dots, u_m)^\top$, respectively. Correspondingly, we specify the ambiguity set $\mathcal{P}$ as follows.

$$\mathcal{P} := \left\{ \mathbf{p} \in \mathbb{R}_+^K \mid \mathbf{l} \leq \sum_{k \in [K]} p_k \mathbf{f}(\boldsymbol{\xi}^k) \leq \mathbf{u} \right\}. \quad (6)$$

Note that to guarantee that the ambiguity set (6) defines a set of probability distributions, one of the moment functions of $\mathbf{f}$ and its corresponding lower and upper bound values l_s, u_s , can be set as $\sum_{k \in [K]} p_k = 1$. The following theorem demonstrates a reformulation of Model (4) given the moment-based ambiguity set in (6).

Theorem 1. If the ambiguity set defined in (6) is non-empty, then Model (4) can be reformulated as a single-level problem in the form:

$$\min \quad \sum_{i \in \mathcal{I}} c_i^o x_i + \sum_{i \in \mathcal{I}, t \in \mathcal{T}} c_i^h h_{it} - \boldsymbol{\alpha}^\top \mathbf{l} + \boldsymbol{\beta}^\top \mathbf{u} \quad (7a)$$

$$\text{s.t.} \quad (4b)-(4d)$$

$$(-\boldsymbol{\alpha} + \boldsymbol{\beta})^\top \mathbf{f}(\boldsymbol{\xi}^k) \geq g(\mathbf{h}, \boldsymbol{\xi}^k), \quad \forall k \in [K], \quad (7b)$$

$$\boldsymbol{\alpha}, \boldsymbol{\beta} \geq 0. \quad (7c)$$

Proof. Explicitly stating the constraints in the ambiguity set $\mathcal{P}$, we first express the inner maximization problem as

$$\max \quad \sum_{k \in [K]} p_k g(\mathbf{h}, \boldsymbol{\xi}^k) \quad (8a)$$

$$\text{s.t.} \quad \sum_{k \in [K]} p_k \mathbf{f}(\boldsymbol{\xi}^k) \geq \mathbf{l}, \quad (8b)$$

$$\sum_{k \in [K]} p_k \mathbf{f}(\boldsymbol{\xi}^k) \leq \mathbf{u}, \quad (8c)$$

$$p_k \geq 0, \quad \forall k \in [K]. \quad (8d)$$

Letting $\boldsymbol{\alpha}, \boldsymbol{\beta}$ be the dual variables associated with the lower- and upper-bound constraints in (8), respectively, we obtain its dual formulation as

$$\min_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \quad -\boldsymbol{\alpha}^\top \mathbf{l} + \boldsymbol{\beta}^\top \mathbf{u} \quad (9a)$$

$$\text{s.t.} \quad (-\boldsymbol{\alpha} + \boldsymbol{\beta})^\top \mathbf{f}(\boldsymbol{\xi}^k) \geq g(\mathbf{h}, \boldsymbol{\xi}^k), \quad \forall k \in [K], \quad (9b)$$

$$\boldsymbol{\alpha}, \boldsymbol{\beta} \geq 0. \quad (9c)$$

Following the strong duality between (8) and (9), and replacing $\max_{\mathbf{p} \in \mathcal{P}} \sum_{k \in [K]} p_k g(\mathbf{h}, \boldsymbol{\xi}^k)$ with (9) in the outer optimization problem (4), we obtain the desired result. $\square$

Denote the empirical first and second moments of the uncertain demand parameter at location

j and period t as μ_{jt} and S_{jt} , respectively. We examine a special form of ambiguity set (6), where for the first and second moments of each demand parameter, we consider their lower and upper bounds as follows:

$$\mathcal{P} = \left\{ \mathbf{p} \in \mathbb{R}_+^K \mid \sum_{k \in [K]} p_k = 1, \right. \quad (10a)$$

$$\mu_{jt} - \epsilon_{jt}^\mu \leq \sum_{k \in [K]} p_k \xi_{jt}^k \leq \mu_{jt} + \epsilon_{jt}^\mu, \quad \forall j \in \mathcal{J}, t \in \mathcal{T}, \quad (10b)$$

$$S_{jt} \underline{\epsilon}_{jt}^S \leq \sum_{k \in [K]} p_k (\xi_{jt}^k)^2 \leq S_{jt} \bar{\epsilon}_{jt}^S, \quad \forall j \in \mathcal{J}, t \in \mathcal{T} \left. \right\}. \quad (10c)$$

Here, Constraint (10a) is a normalization constraint to ensure that $\{p_k\}_{k \in [K]}$ form a probability distribution. Constraints (10b) bound the mean of parameter ξ_{jt} in an ϵ_{jt}^μ -interval of the empirical mean μ_{jt} , and constraints (10c) bound the second moment of parameter ξ_{jt} via scaling the empirical second moment S_{jt} with parameters $0 \leq \underline{\epsilon}_{jt}^S \leq 1 \leq \bar{\epsilon}_{jt}^S$ for all $j \in \mathcal{J}$, $t \in \mathcal{T}$. We note that under the perfect knowledge assumption of first and second moment information, the robustness parameters can be set as $\epsilon_{jt}^\mu = 0$, and $\underline{\epsilon}_{jt}^S = \bar{\epsilon}_{jt}^S = 1$, for all $j \in \mathcal{J}$, $t \in \mathcal{T}$. Under the uncertainty of a pandemic, adjusting these parameters helps decision makers make inferences for representing spatiotemporal demand by considering different phases of the pandemic and taking into account potential deviations from predictive results.

To obtain a mixed-integer linear programming reformulation of the single-level formulation (7) given ambiguity set (10), we first describe some intermediate steps and results. In this regard, we analyze certain properties of the function $g(\mathbf{h}, \boldsymbol{\xi}^k)$.

Proposition 1. Function $g(\mathbf{h}, \boldsymbol{\xi}^k)$ is a convex piecewise linear function in $\mathbf{h}$ for every $k \in [K]$.

Proof. First, note that problem (5) is always feasible as it has complete recourse. Letting θ_{it} , γ_{jt} be the dual variables, we obtain the dual of problem (5) as follows.

$$\max \quad \sum_{i \in \mathcal{I}} \sum_{t \in \mathcal{T}} h_{it} \theta_{it} + \sum_{j \in \mathcal{J}} \sum_{t=2}^T \xi_{jt}^k \gamma_{jt} + \sum_{j \in \mathcal{J}} (\xi_{j1}^k + u_{j0} - I_{j0}) \gamma_{j1} \quad (11a)$$

$$\text{s.t.} \quad \theta_{it} + \gamma_{jt} \leq c_{ijt}^s, \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}, \quad (11b)$$

$$-\gamma_{jt} + \gamma_{j(t+1)} \leq c_{jt}^I, \quad \forall j \in \mathcal{J}, t \in [T-1], \quad (11c)$$

$$-\gamma_{jT} \leq c_{jT}^I, \quad \forall j \in \mathcal{J}, \quad (11d)$$

$$\gamma_{jt} - \gamma_{j(t+1)} \leq c_{jt}^u, \quad \forall j \in \mathcal{J}, t \in [T-1], \quad (11e)$$

$$\gamma_{jT} \leq c_{jT}^u, \quad \forall j \in \mathcal{J}, \quad (11f)$$

$$\theta_{it} \leq 0 \quad \forall i \in \mathcal{I}, t \in \mathcal{T}. \quad (11g)$$

As the dual problem (11) is feasible, at least one of the optimal solutions of this problem is at one of its extreme points. Considering the fact that objective function of this problem is linear in $\mathbf{h}$, the

resulting optimal objective value can be represented as the maximum of the extreme-point-based function values. Hence, $g(\mathbf{h}, \boldsymbol{\xi}^k)$ becomes a convex piecewise linear function in $\mathbf{h}$. $\square$

Lemma 1. Denote $g(\mathbf{h}, \boldsymbol{\xi}^k)$ by a compact form $\min_{\mathbf{y}} \{\mathbf{c}^\top \mathbf{y} : \mathbf{A}\mathbf{y} \geq \mathbf{h}, \mathbf{D}\mathbf{y} \geq \mathbf{f}\}$. Then, set $\Upsilon^1 := \{(\mathbf{h}, m) : g(\mathbf{h}, \boldsymbol{\xi}^k) \leq m\}$ for each $\boldsymbol{\xi}^k \in S$ has a polyhedral representation $\Upsilon^2 := \{(\mathbf{h}, m) : \exists \mathbf{y} \text{ s.t. } \mathbf{c}^\top \mathbf{y} \leq m, \mathbf{A}\mathbf{y} \geq \mathbf{h}, \mathbf{D}\mathbf{y} \geq \mathbf{f}\}$.

Proof. Demonstrated in Proposition 1, $g(\mathbf{h}, \boldsymbol{\xi}^k)$ is a convex function. Consider any $(\mathbf{h}, m) \in \Upsilon^1$. The optimal solution of $g(\mathbf{h}, \boldsymbol{\xi}^k)$, say $\mathbf{y}^1$, satisfies $\mathbf{c}^\top \mathbf{y}^1 \leq m, \mathbf{A}\mathbf{y}^1 \geq \mathbf{h}, \mathbf{D}\mathbf{y}^1 \geq \mathbf{f}$, proving that $(\mathbf{h}, m) \in \Upsilon^2$. For the other direction of the proof, consider any $(\mathbf{h}, m) \in \Upsilon^2$. Then, there exists a vector $\mathbf{y}^2$ such that it satisfies $\mathbf{c}^\top \mathbf{y}^2 \leq m, \mathbf{A}\mathbf{y}^2 \geq \mathbf{h}, \mathbf{D}\mathbf{y}^2 \geq \mathbf{f}$. As $\mathbf{y}^2$ is a feasible solution of the problem $\min_{\mathbf{y}} \{\mathbf{c}^\top \mathbf{y} : \mathbf{A}\mathbf{y} \geq \mathbf{h}, \mathbf{D}\mathbf{y} \geq \mathbf{f}\}$, $g(\mathbf{h}, \boldsymbol{\xi}^k) \leq \mathbf{c}^\top \mathbf{y}^2 \leq m$, and therefore, $(\mathbf{h}, m) \in \Upsilon^1$. $\square$

Combining Lemma 1 with the single-level formulation (7), we propose a mixed-integer linear programming reformulation of Model (4) as follows.

Theorem 2. Using the ambiguity set defined in (10), Model (4) is equivalent to the following mixed-integer linear program:

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{I}} c_i^o x_i + \sum_{i \in \mathcal{I}, t \in \mathcal{T}} c_i^h h_{it} - \alpha_1 - \sum_{j \in \mathcal{J}, t \in \mathcal{T}} \alpha_{2jt} (\mu_{jt} - \epsilon_{jt}^\mu) - \sum_{j \in \mathcal{J}, t \in \mathcal{T}} \alpha_{3jt} (\mu_{jt}^2 + \sigma_{jt}^2) \bar{\epsilon}_{jt}^S \\ & + \beta_1 + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} \beta_{2jt} (\mu_{jt} + \epsilon_{jt}^\mu) + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} \beta_{3jt} (\mu_{jt}^2 + \sigma_{jt}^2) \bar{\epsilon}_{jt}^S \end{aligned} \quad (12a)$$

s.t. (4b)–(4d)

$$- \alpha_1 + \beta_1 + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} \xi_{jt}^k (-\alpha_{2jt} + \beta_{2jt}) + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} (\xi_{jt}^k)^2 (-\alpha_{3jt} + \beta_{3jt}) \geq \Phi^k, \quad \forall k \in [K], \quad (12b)$$

$$\Phi^k = \sum_{i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}} c_{ijt}^s s_{ijt}^k + \sum_{j \in \mathcal{J}, t \in \mathcal{T}} \left(c_{jt}^I I_{jt}^k + c_{jt}^u u_{jt}^k \right), \quad \forall k \in [K], \quad (12c)$$

$$\sum_{j \in \mathcal{J}} s_{ijt}^k \leq h_{it}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}, \quad \forall k \in [K], \quad (12d)$$

$$\sum_{i \in \mathcal{I}} s_{ijt}^k + I_{j(t-1)}^k + u_{jt}^k = \xi_{jt}^k + I_{jt}^k + u_{j(t-1)}^k, \quad \forall j \in \mathcal{J}, t \in \mathcal{T}, \quad \forall k \in [K], \quad (12e)$$

$$s_{ijt}^k, I_{j,t}^k, u_{jt}^k \geq 0, \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, t \in \mathcal{T}, \quad \forall k \in [K], \quad (12f)$$

$$\boldsymbol{\alpha}, \boldsymbol{\beta} \geq \mathbf{0}. \quad (12g)$$

Proof. To obtain the mixed-integer linear programming reformulation, we first revise the single-level formulation (7) under the ambiguity set defined in (10). Then, we derive a polyhedral representation of the set $\{(\mathbf{h}, m) : g(\mathbf{h}, \boldsymbol{\xi}^k) \leq m\}$ where m is the left-hand side of constraint (7b). Using Lemma 1 and the definition of the function $g(\mathbf{h}, \boldsymbol{\xi}^k)$, we derive the resulting formulation (12). $\square$

5 Case Studies of COVID-19 Vaccine and Test Kit Distribution

In this section, we present two comprehensive case studies of the presented optimization frameworks to conduct resource distribution under uncertain spatiotemporal demand for COVID-19 testing and vaccination. In Section 5.1, we consider COVID-19 vaccine distribution in the US, and in Section 5.2, we consider COVID-19 test kit distribution in the State of Michigan. For both types of instances, we compute solutions of three different approaches, namely, the deterministic (DT) approach where the SMIP model (1) only contains one representative scenario in set Ω , which can be the mean values of demand, the stochastic programming (SP) approach described in Section 3, and the distributionally robust optimization (DRO) approach described in Section 4. For the SP and DRO approaches, we construct the optimization models based on 100 in-sample scenarios that are independently and identically generated from a given nominal distribution of the demand, and parameters in the nominal distribution follow estimated demand mean and variance based on real-world COVID-19 infection data in the corresponding regions of interest that we will describe later. Different solutions will be evaluated in out-of-sample scenarios to see their performance in terms of operational cost and possible lost sale/backlog (if demand values in the out-of-sample tests become extremely high). Note that the out-of-sample scenarios can be generated from the same nominal distribution for generating the in-sample data, or be a different one. The latter represents the case when future infection trends and thus future demand values become significantly different from what has been observed. [We will test both cases in Section 5.1.2 to examine how the differences in data-generating distributions affect the solution performance.](#) We will also describe details about how to generate out-of-sample scenarios and the results for the vaccine-distribution instances and test-kit distribution instances separately.

We use Gurobi 9.0.3 coded in Python 3.6.8 for solving all mixed-integer programming models. Our numerical tests are conducted on a Windows 2012 Server with 128 GB RAM and an Intel 2.2 GHz processor.

5.1 Vaccine Distribution in the US

5.1.1 Experimental design and setup

We test different vaccine-allocation phases following the government and CDC’s guidelines, where in the earlier phases only healthcare workers and prioritized population groups (e.g., seniors more than 65 years old) are targeted for vaccination and in the later phases, larger segments of the population are recommended for vaccination to stop virus transmission. Specifically, we follow a recent epidemiological study by Wang et al. (2020) and consider three phases to distribute the vaccines, aiming to vaccinate 6.18%, 41.97%, 51.85% of the overall adult population in the US, respectively. We also take into account certain levels of vaccine hesitancy in the population during each phase, and follow the results by Malik et al. (2020) who surveyed the US adult population to understand the acceptance of COVID-19 vaccines, among 10 Department of Health and Human Services (DHHS) regions listed below.

- Region 1 – Boston (Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont);
- Region 2 – New York (New Jersey, New York, and Puerto Rico);
- Region 3 – Philadelphia (Delaware, District of Columbia, Maryland, Pennsylvania, Virginia, and West Virginia);
- Region 4 – Atlanta (Alabama, Florida, Georgia, Kentucky, Mississippi, North Carolina, South Carolina, and Tennessee);
- Region 5 – Chicago (Illinois, Indiana, Michigan, Minnesota, Ohio, and Wisconsin);
- Region 6 – Dallas (Arkansas, Louisiana, New Mexico, Oklahoma, and Texas);
- Region 7 – Kansas City (Iowa, Kansas, Missouri, and Nebraska);
- Region 8 – Denver (Colorado, Montana, North Dakota, South Dakota, Utah, and Wyoming);
- Region 9 – San Francisco (Arizona, California, Hawaii, Nevada);
- Region 10 – Seattle (Alaska, Idaho, Oregon, and Washington).

We select the representative cities in the DHHS regions as our demand sites, which aggregate the demand in each region. Based on the acceptance rate estimates and sample sizes of the surveys conducted by Malik et al. (2020), we compute a 90% confidence interval of the acceptance rate for each region, depicted in Table 1. We then multiply these values by the total adult population in each region to obtain lower and upper bounds on the number of people to be vaccinated. As each person needs to get two doses of vaccines, we multiply these bounds by 2 to determine the number of doses needed (i.e., potential demand). To represent the demand uncertainty, we sample scenarios following uniform distributions between the demand lower and upper bounds (i.e., $\underline{\mu}_{jt}$, $\bar{\mu}_{jt}$). Table 2 summarizes the mean values of the estimated demand during the three phases for each DHHS region.

Table 1: Confidence intervals of COVID-19 vaccines’ acceptance rates in 10 DHHS regions

Regions	Region 1	Region 2	Region 3	Region 4	Region 5	Region 6	Region 7	Region 8	Region 9	Region 10
Acceptance	68.06%	43.14%	72.04%	59.76%	39.13%	74.42%	72.22%	80.00%	67.61%	70.00%
90% CI LB	59.02%	31.73%	64.39%	50.85%	22.39%	68.95%	54.86%	65.29%	58.47%	60.99%
90% CI UB	77.09%	54.55%	79.70%	68.66%	55.87%	79.89%	89.59%	94.71%	76.74%	79.01%

Table 2: Estimated demand mean values during the three phases in 10 DHHS regions

Regions	Region 1	Region 2	Region 3	Region 4	Region 5	Region 6	Region 7	Region 8	Region 9	Region 10	Total
Phase 1	1.0M	1.3M	2.2M	3.9M	2.0M	3.0M	970.0K	920.4K	3.3M	968.5K	19.4M
Phase 2	6.8M	9.0M	14.7M	26.3M	13.4M	20.1M	6.6M	6.3M	22.6M	6.6M	132.4M
Phase 3	8.4M	11.2M	18.2M	32.5M	16.6M	24.8M	8.1M	7.7M	27.9M	8.1M	163.5M
Total	16M	22M	35M	63M	32M	48M	16M	15M	54M	16M	316M

As the sizes of populations to be vaccinated vary during different phases, we assume the lengths of the three phases to be 1, 2, and 3 months, respectively, with one period in our models being 2 weeks, and the demand of each phase is evenly distributed for each period. To determine candidate locations for siting DCs, we first consider the ones being used for vaccine production in the US currently. There are 5 DCs used by Pfizer-BioNTech and Moderna in the US. As these facilities are already available, we set their corresponding x -values to 1 in all the models in the default case, and examine two more cases with this restriction being relaxed or strengthened later. We select representative cities of the 10 DHHS regions as potential locations to open additional DCs, and Table 3 depicts all 15 DC locations and 10 demand locations. For each DC $i \in \mathcal{I}$, we calculate the capacity upper bound M_i by assuming the daily maximum production to be 500,000, 750,000 and 1,000,000 doses for Phases 1, 2 and 3, respectively, as manufacturers will raise their daily production capacity as the demand increases. For each $t \in \mathcal{T}$, the temporal capacity of manufacturing vaccines B_t is set as the sum of the maximum capacities over all DCs. We also examine the case when the supply chain of vaccines experiences certain disruptions and set the corresponding temporal capacities B_t as 10% of the normal capacities. The operating cost c_i^o for each DC $i \in \mathcal{I}$ is estimated as 10000 times the local unit warehouse rental price per square foot, by assuming a standard warehouse is about 10000 square feet. The unit capacity cost c_i^h is set as \$25 for each $i \in \mathcal{I}$, which is the average market price of one dose of COVID-19 vaccine. The unit inventory cost c_{jt}^I is estimated as the sum of low-temperature inventory cost and energy cost of common refrigerators for storing vaccines, which is \$0.00008 for each $j \in \mathcal{J}$, $t \in \mathcal{T}$. The unit shipping cost c_{ijt}^s for each region $j \in \mathcal{J}$ and supplier $i \in \mathcal{I}$ consists of two parts: The first part is the shipping cost for trucks calculated as \$3 per mile times the distance traveled from i to j in miles divided by 230,400, following the fact that each truck on average can carry 230,400 doses of Moderna vaccines (Moderna, Inc., 2021), and the second part is the refrigerated trucks' overall cost per liter of vaccine transported (PATH, World Health Organization (WHO), 2013). Via sensitivity analysis of different penalty values for unit demand shortage, the penalty cost c_{jt}^u of each unit of unsatisfied demand is set to \$100. The sources used for estimating these parameters are further summarized in Table 13 in the Appendix.

5.1.2 Sensitivity Analysis

We employ DT, SP, and DRO approaches to optimize facility location and resource distribution for Phases 1, 2, and 3, based on the same in-sample data. At the default setting, we randomly sample 100 scenarios (i.e., $K = 100$) for the in-sample computation from a uniform distribution with the upper and lower bounds calculated based on Table 1. For DT, we set demand mean values μ_{jt} , $j \in \mathcal{J}$, $t \in \mathcal{T}$ as the empirical mean values of the 100 in-sample scenarios. For DRO, we set parameter ϵ_{jt}^μ in (10b) as $\epsilon_{jt}^\mu = 0.5\mu_{jt}$, and bounding parameters in (10c) as $\underline{\epsilon}_{jt}^S = 0.1$, $\bar{\epsilon}_{jt}^S = 2$ for all $j \in \mathcal{J}$, $t \in \mathcal{T}$. The SP approach uses all the 100 scenarios to form set Ω in the SMIP model (1) and the DRO approach uses the in-sample scenarios as its support set for computing Phase 1's solutions. We independently generate 1000 out-of-sample scenarios for evaluating solution performance of the three approaches. In the last three columns in Table 3, we display the DCs that

Table 3: Locations of 10 customers sites, 15 candidate DCs and optimal solutions for Phase 1 given by Deterministic (DT), Stochastic (SP) and Distributionally Robust Optimization (DRO) approaches

City	Customer sites	Status of candidate DCs	DT	SP	DRO
Kalamazoo, MI		already opened			
Pleasant Prairie, WI		already opened		✓	
Bloomington, IN		already opened			✓
Norwood, MA		already opened	✓	✓	✓
Saint Louis, MO		already opened		✓	
Boston, MA	Region 1	candidate			
New York City, NY	Region 2	candidate	✓	✓	
Philadelphia, PA	Region 3	candidate	✓		✓
Atlanta, GA	Region 4	candidate	✓	✓	✓
Chicago, IL	Region 5	candidate	✓		✓
Dallas, TX	Region 6	candidate	✓	✓	✓
Kansas City, KS	Region 7	candidate	✓		✓
Denver, CO	Region 8	candidate	✓		✓
San Francisco, CA	Region 9	candidate	✓	✓	
Seattle, WA	Region 10	candidate	✓		✓

are open (in addition to the existing 5 DCs) in the optimal solutions of the three approaches for Phase 1’s problem.

To evaluate the impact of data-generating distributions, we present the in-sample and out-of-sample performance of the three approaches using normal and uniform distributions in Table 4, where in the normal distributions, we set the mean values to $(\underline{\mu}_{jt} + \bar{\mu}_{jt})/2$ and set the standard deviation to $(\bar{\mu}_{jt} - \underline{\mu}_{jt})/6$, respectively. The in-sample and out-of-sample scenarios are generated according to the same distributions in Table 4. We also record the percentage changes of the out-of-sample cost compared to the in-sample cost in the bracket. In practice, it is possible that the true data follows a distribution that is different than what we assume and therefore, we generate some out-of-sample scenarios by shifting an assumed distribution to the right. (That is, the true distributions have mean values of $1.01\mu_{jt}$ with μ_{jt} being the mean values of the assumed distributions.) The corresponding results are summarized in Table 5.

Table 4: In-sample and out-of-sample performance comparison of the three approaches using different demand distribution assumptions and unshifted mean values.

Distribution	Approach	Unmet Demand		Overall Cost	
		In-sample	Out-of-sample	In-sample	Out-of-sample
Normal	DT	0	20K	97M	100M (+2.15%)
	SP	3K	3K	98M	98M (+0.02%)
	DRO	0	34	99M	99M (+0.003%)
Uniform	DT	0	294K	489M	518M (+6.02%)
	SP	72K	73K	504M	504M (+0.02%)
	DRO	6	752	520M	520M (+0.01%)


Table 5: In-sample and out-of-sample performance comparison of the three approaches using different demand distribution assumptions and shifted mean values.

Distribution	Approach	Unmet Demand		Overall Cost	
		In-sample	Out-of-sample	In-sample	Out-of-sample
Normal	DT	0	29K	97M	100M (+3.01%)
	SP	3K	5K	98M	99M (+0.21%)
	DRO	1	28	99M	99M (+0.003%)
Uniform	DT	0	3M	489M	767M (+56.89%)
	SP	72K	2M	504M	656M (+30.21%)
	DRO	6	784K	520M	599M (+15.08%)


From Table 4, the gaps between the out-of-sample and in-sample objective costs given by the DRO approach are the smallest across all three approaches and are always within 0.01%, while the DT approach attains the highest gaps at 6.02% when assuming uniformly distributed demand. These gaps further increase in Table 5 due to the inconsistency of assumed and true distributions. Notably, even with shifted means, the DRO approach can still achieve an objective cost gap of 0.003% when assuming Normal distributed demand, showing the result robustness when having distributional ambiguity. Comparing the two demand distribution assumptions, following the Normal distribution generates much smaller in-sample and out-of-sample gaps. This is mainly because most of the scenarios generated by the Normal distribution are centered around the mean, and even with shifted means, the in-sample and out-of-sample scenarios are close to each other. However, when using the Uniform distribution to generate demand samples, scenarios are scattered between lower and upper bounds uniformly, and shifting the bounds can change scenarios significantly. Because of that, we only assume Normal distribution as the true demand distribution in the testing kit distribution problem in Section 5.2.

Next, we vary the number of in-sample scenarios K from 10 to 100, and present the overall cost of the DRO approach on the same 1000 out-of-sample scenarios and the computational time comparison in Figure 1. We also compare two different ambiguity sets in the DRO approach, where in “First+Second Moment”, we use both first and second moments’ information (10b)–(10c) to construct the ambiguity set, and in “First Moment”, we only use the first moment’s information (10b). From Figure 1, using only the first moment can obtain nearly the same out-of-sample performance as using both first and second moment information, and it scales better with the number of in-sample scenarios K . As we include more in-sample scenarios in the discrete support of the DRO approach, there are more choices for the inner-max problem to select the worst-case distribution and thus we become more conservative. Although the out-of-sample cost is fluctuating due to the differences between out-of-sample and in-sample scenarios, we can still see an increasing trend from Figure 1a. In practice, to trade off between the computational complexity and conservatism, one can choose a reasonable K between 10 and 100. In our test for solving Phases 2 and 3’s problems, we set K to 10 for DRO approach to produce solutions within a reasonable time limit.

We also examine the impact of the bounding parameters in the DRO approach, where we present




(a) Out-of-sample overall cost comparison



(b) Computational time comparison

Figure 1: Out-of-sample cost and CPU time of the DRO approach with varying K over different ambiguity sets.

the out-of-sample performance with varying ϵ_μ in Figure 2a, and the out-of-sample performance with varying ϵ_S in Figure 2b, respectively. From Figure 2, the out-of-sample cost is not very



(a) Results with varying ϵ_μ

(b) Results with varying ϵ_S

Figure 2: Out-of-sample cost and unmet demand changes with varying ϵ_μ and ϵ_S .

sensitive to both bounding parameters, as the percentage changes are all within 0.03%. When we increase ϵ_S , there is an increasing trend in both unmet demand and overall cost because we enlarge the ambiguity set and the results become more conservative.

5.1.3 Results

For solving problems in Phases 2 and 3, in order to obtain solutions for the DRO approach within a reasonable time limit (e.g., 24 hours), we reduce the number of in-sample scenarios to 10 to form the support set. Also, in our in-sample computation, we assume that the supply of vaccines matches the demand and therefore try to minimize the unmet demand only against its uncertainty rather than also taking into account resource scarcity. We present the unsatisfied demand and overall cost

of different approaches under full and scarce resources (the latter being 10% of the former) in Table 6, where we mark the percentage increase compared with the optimal one in parentheses. The last column in Table 6 records the in-sample computational time in seconds for each method.

Table 6: Out-of-sample performance of different approaches in terms of unmet demand and overall cost for Phases 1, 2, 3 of vaccine distribution in the US.

Phases	Approach	Unmet Demand		Overall Cost		Time (sec.)
		Full Resource	Scarce Resource	Full Resource	Scarce Resource	
Phase 1	DT	372K ($+1.26 \times 10^4\%$)	372K ($+5.31 \times 10^3\%$)	\$524M ($+4.11\%$)	\$524M ($+4.04\%$)	0.1
	SP	64K ($+2.11 \times 10^3\%$)	68K ($+8.85 \times 10^2\%$)	\$503M	\$503M	6.5
	DRO	3K	7K	\$516M ($+2.50\%$)	\$513M ($+1.95\%$)	2148.8
Phase 2	DT	2M ($+2.37 \times 10^3\%$)	173M	\$4B ($+4.42\%$)	\$19B ($+8.13 \times 10^{-5}\%$)	0.1
	SP	166K ($+77\%$)	173M	\$3B	\$19B	1.3
	DRO	94K	173M	\$3B ($+0.40\%$)	\$19B ($+4.67 \times 10^{-5}\%$)	903.2
Phase 3	DT	3M ($+7.38 \times 10^2\%$)	131M	\$4B ($+4.40\%$)	\$16B ($+4.23 \times 10^{-4}\%$)	0.2
	SP	313K	131M	\$4B	\$16B	2.5
	DRO	450K ($+43\%$)	131M	\$4B ($+0.39\%$)	\$16B ($+7.43 \times 10^{-4}\%$)	907.2

From Table 6, in Phase 1, DRO obtains the least amount of unsatisfied demand and the second highest cost overall, while DT performs the worst in terms of both demand satisfaction and overall cost. In terms of the overall cost, SP always outperforms the other two approaches. This is because in the ambiguity set (10), we allow deviations from the predictive results μ_{jt} , S_{jt} and the worst-case scenario usually achieves a demand mean higher than the empirical mean μ_{jt} . Compared to SP model, DRO always produces a higher manufacturing capacity against the worst-case scenario (as can be seen from Tables 14–15) and thus leads to a higher overall cost but lower unsatisfied demand. When facing scarce resource (i.e., 10% of the regular capacity), the results do not change significantly in Phase 1 because all three approaches do not use up all the resources. However, when the resource capacity becomes tight (e.g., in Phases 2 and 3), the differences between the three approaches are almost negligible as there is not much flexibility in making different production and shipment plans due to limited temporal resources. We note that the raw cost breakdowns of each approach are detailed in Tables 14 and 15 in the Appendix. Next, we focus on the analysis of optimal solutions of different methods under the setting of full resources.

The average shipments in the optimal solutions of SP and DRO for Phase 1’s vaccine distribution are presented in Figures 3 and 4, respectively. From Figure 3, seven DCs are selected in the optimal solution, and among them four are in regions with high demand volumes, and three are from the existing Pfizer-BioNTech and Moderna open DCs. On the other hand, the other two existing DCs do not produce or ship any vaccines either because they are too close to some other open DCs in the optimal solution or they reside in the areas having low demand volumes. Moreover, the top five largest shipments are from DCs to their nearest demand sites (i.e., New York City to Philadelphia, Pleasant Prairie to Chicago, Saint Louis to Kansas City, Dallas to Denver, and San Francisco to Seattle), and each requires up to five trucks if fully occupied and the delivery time will be all within two days. All the other longer-distance shipments are in smaller volumes, requiring only one truck. For example, the longest-distance shipment is from Norwood to Seattle, which may take 4 to 5 days, but it only carries 25 doses of vaccines. Comparing Figure 4 with Figure 3, DRO selects more




Figure 3: Average shipments of Phase 1’s vaccines in 1000 out-of-sample scenarios given by the optimal solution of the SP approach. The background color of 10 DHHS regions indicates the total demand across all three phases where darker color represents higher demand volume. The black and gray circles stand for open DCs and demand sites, respectively, while DCs located in New York City, Atlanta, Dallas and San Francisco are also demand sites themselves and the sizes of these black circles represent the amount of vaccines they produce for satisfying their own demand. The black lines represent shipments from open DCs to demand locations with wider lines meaning higher volumes.

DCs to open, among them only two are existing open ones. The number of shipment routes also increases significantly. As San Francisco is not an open DC anymore, the largest shipment occurs between Seattle and San Francisco and requires 8 trucks and 2 days to transit. The second and third largest shipments are from Philadelphia to New York City and Norwood to Boston, respectively, both requiring 3 trucks and 1 day.

Furthermore, we present Phase 1’s unsatisfied demand mean, standard deviation, and 75 to 95 percentile values in the 1000 out-of-sample scenarios according to the three approaches in Figure 5a, and average unsatisfied demand percentages in each region in Figure 5b, respectively. From Figure 5, DRO attains the least amount of unsatisfied demand while DT performs the worst. Comparing the results in different regions, Region 5 – Chicago does not meet over 4% of the demand in DT approach, and Region 8 – Denver does not cover about 4% of the demand in SP approach, both of which have relatively low demand volumes according to Table 2.

In the above results, we assume that the existing 5 open DCs for Pfizer-BioNTech and Moderna have to be used but we can open additional ones if necessary. Next we examine two other cases where we either relax this constraint and allow to open facilities in any DC location options, or only open the 5 DCs. We refer to these two cases as the “best-case” and “most-restrictive” settings,





Figure 4: Average shipments of Phase 1’s vaccines in 1000 out-of-sample scenarios given by the optimal solutions of the DRO approach. The background color of 10 DHHS regions indicates the total demand across all three phases where darker color represents higher demand volume. The black and gray circles stand for open DCs and customer sites, respectively, while DCs located in Philadelphia, Chicago, Atlanta, Dallas, Kansas City, Denver and Seattle are also demand sites themselves and the size of these black circles represent the amount of vaccines they produced for satisfying their own demand. The black lines represent shipments from open DCs to demand locations with wider lines meaning higher volumes.

respectively, and compare their Phase 2’s results with the default setting in Table 7. In Column “Overall Cost”, we display the total cost for each approach and present their percentage changes from their corresponding costs in the default setting. The last column records the total number of DCs open by each approach.


Table 7: Out-of-sample performance of different approaches for Phase 2 of vaccine distribution in the US under different DC settings.

DC Settings	Model	Operating	Capacity	Shipping	Inventory	Penalty	Overall Cost	# DCs
Default	DT	\$194K	\$3B	\$145K	\$259	\$231M	\$4B	14
	SP	\$175K	\$3B	\$126K	\$752	\$17M	\$3B	13
	DRO	\$107K	\$3B	\$545K	\$519	\$9M	\$3B	8
Best-case	DT	\$154K	\$3B	\$142K	\$260	\$231M	\$4B ($-1.21 \times 10^{-3}\%$)	10
	SP	\$154K	\$3B	\$118K	\$744	\$17M	\$3B (+0.01%)	10
	DRO	\$204K	\$3B	\$473K	\$3K	\$7M	\$3B (-0.08%)	14
Most-restrictive	DT	\$60K	\$3B	\$1M	\$143	\$272M	\$4B (+1.19%)	5
	SP	\$60K	\$3B	\$1M	\$522	\$18M	\$3B (+0.06%)	5
	DRO	\$60K	\$3B	\$1M	\$455	\$10M	\$3B (+0.03%)	5

From Table 7, after relaxing the constraints on the existing 5 open DCs for Pfizer-BioNTech



(a) Unsatisfied demand statistics



(b) Regional unsatisfied demand

Figure 5: Unsatisfied demand given by vaccine distribution in Phase 1 of different approaches.

and Moderna, both DT and SP reduce the operational and shipping cost in the best-case setting due to better options of DCs to open. Moreover, the number of open DCs increases from 8 to 14 in the DRO approach because it has more freedom to choose DCs. In the best-case setting, DRO opens more DCs than the other two benchmarks as the optimal solution is chosen against the worst-case scenario where the demand mean is higher than the empirical one. Comparing the most-restrictive setting where only the existing 5 DCs are open with the default setting, all approaches have increased shipping, penalty and overall costs.

Next, we vary the resource scarcity level between 10% and 30% of the full capacity and present the out-of-sample performance of different approaches for Phase 2 in Table 8. From the table, the results with 30% scarcity level are similar to the ones with full capacity. When the scarcity level decreases, the capacity cost drops while penalty and overall costs increase drastically. Moreover, results of the three approaches are more similar given low capacity.

We also compare the vaccine allocation results given by the SP and DRO approaches with the current vaccination distribution status reported by CDC (Centers for Disease Control and Prevention, 2021c). From CDC’s website (Centers for Disease Control and Prevention, 2021b,a), Pfizer-BioNTech has started to provide first-dose vaccines since December 14, 2020, while Moderna started to ship first-dose vaccines on December 21, 2020. Consequently, we let December 14, 2020 be the starting date of our planning horizon and denote December 14, 2020 – January 11, 2021 as Phase 1, January 11, 2021 – March 8, 2021 as Phase 2, and March 8, 2021 – May 31, 2021 as Phase

Table 8: Out-of-sample performance of different approaches for Phase 2 of vaccine distribution in the US under different resource scarcity levels.

Scarcity	Model	Operating	Capacity	Shipping	Inventory	Penalty	Overall Cost
30%	DT	\$194K	\$3B	\$145K	\$260	\$231M	\$4B
	SP	\$175K	\$3B	\$126K	\$752	\$17M	\$3B
	DRO	\$131K	\$3B	\$419K	\$536	\$10M	\$3B
20%	DT	\$194K	\$3B	\$150K	\$15	\$2B	\$5B
	SP	\$175K	\$3B	\$117K	\$14	\$2B	\$5B
	DRO	\$178K	\$3B	\$159K	\$12	\$2B	\$5B
10%	DT	\$105K	\$2B	\$48K	\$0	\$17B	\$19B
	SP	\$96K	\$2B	\$42K	\$0	\$17B	\$19B
	DRO	\$105K	\$1B	\$41K	\$0	\$17B	\$19B

3. We report the total number of doses that our SP and DRO approaches have distributed and CDC has delivered/administered by each checkpoint in Table 9 as of September 27, 2021. We note that both of our SP and DRO models finish distributing vaccines by the end of Phase 3 (i.e., May 31, 2021) and as a result, the total number of doses distributed until September 27, 2021 remains the same with the one until May 31, 2021, while CDC still distributes vaccines during this period.

Table 9: Total number of doses distributed in the optimal solutions of SP/DRO and CDC’s delivered and administered data for each phase as of September 27, 2021

Approaches	Until 1/11/2021	Until 3/8/2021	Until 5/31/2021	Until 9/27/2021
SP Distributed	20M	155M	321M	321M
DRO Distributed	21M	157M	322M	322M
CDC Delivered	24M	116M	366M	472M
CDC Administered	9M	92M	296M	391M

From Table 9, compared with SP and DRO approaches, CDC allocates much more vaccine doses in Phase 1 and but fewer in Phase 2. By the end of Phase 3, the SP and DRO approaches have distributed about 321 millions of doses, which are much closer to the real demand (i.e., 296M) as reported by CDC’s administered data. As of September 27, 2021, CDC delivered 472 millions of doses while the total number of doses received by all people in the US is only 391M, leaving about 81 millions of doses unused, becoming either disposal or inventory. Both of the SP and DRO approaches are demand-driven, as they can improve CDC vaccine allocation plans by allocating more vaccines to appropriate regions when demand is high, or by reducing allocation amounts to avoid unnecessary shipping and inventory cost when demand is low.

5.2 COVID-19 Test Kit Allocation in Michigan, USA

5.2.1 Experimental Design and Setup

To demonstrate that the presented generic approaches are well-suited for distributing various types of resources in different geographic scales, in this section we consider the COVID-19 test kit resource

allocation problem in the State of Michigan, and divide the state into eight emergency preparedness regions according to Figure 6 and consider aggregated demand in these regions. (Regions 2N, 2S, 3 are renamed as Regions 2, 3 and 4 for notation simplicity in the case study.) We pick 29 test-kit suppliers in Michigan as candidate locations for setting up DCs.



Figure 6: Michigan Emergency Preparedness Regions (source: Michigan state government).

For predicting demand and generating uncertain scenarios, we mainly follow Pei and Shaman (2020), in which the authors estimate the daily new county-level confirmed cases in the US based on a meta-population Susceptible-Exposed-Infectious-Removed (SEIR) model. The demand projections are reported in the 2.5, 25, 50, 75, 97.5 percentiles for each county everyday, and are regularly updated to capture different trends in virus transmission and intervention policies during the pandemic. To identify the demand of each emergency preparedness region, we first calculate the percentiles of the total projected confirmed cases, and divide them by our targeted positivity rate (i.e., 5%), meaning that we aim to test 100 people if there are 5 confirmed cases. Having obtained the percentiles of the projected demand, we extract the mean and variance according to Wan et al. (2014) by assuming the 2.5 and 97.5 percentile values being lower and upper bounds of the support of the underlying random demand for estimation.

We consider different planning horizons to represent various phases of virus transmission in Michigan – the number of infected cases reached the first peak in April, 2020; it was then controlled in June, 2020 with the regulations from governor’s executive orders and closures of businesses and public services; around December, winter holiday gatherings increased virus spread, and thus resulted in another peak; most recently, the re-opening of schools in September, 2021 leads to continued transmissions. Consequently, we select April 5, 2020, June 4, 2020, December 20, 2020 and September 19, 2021 as the starting dates of the planning horizon, where we consider a two-

week planning horizon with one period being one week. As a result, we examine the presented approaches over (i) the first peak (i.e., April 5, 2020 – April 19, 2020), (ii) off-peak periods where we can still observe the effects of lockdown procedures (i.e., June 4, 2020 – June 18, 2020), (iii) the second peak (i.e., December 20, 2020 – January 3, 2021) and (iv) continued transmission period (i.e., September 19, 2021 – October 3, 2021).

In terms of parameter choices, we set the upper bound M_i of the capacity h_{it} to the maximum production amount of supplier $i \in \mathcal{I}$ at period $t \in \mathcal{T}$ depending on the size of each supplier, and the temporal capacity of manufacturing test kits B_t to the sum of maximum capacity over all DCs for all $t \in \mathcal{T}$. We also test the case where there are only scarce resources available (usually in earlier phases of the pandemic), and set the capacity B_t to 10% of the regular capacities for all $t \in \mathcal{T}$. The operating cost c_i^o for each supplier $i \in \mathcal{I}$ is similar to the one in the vaccine distribution case. We set the unit capacity cost c_i^h to \$20 for all $i \in \mathcal{I}$, and the unit inventory cost c_{jt}^I to the average pallet storage cost divided by the number of test kits a pallet can store for all $j \in \mathcal{J}$, $t \in \mathcal{T}$. The unit shipping costs c_{ijt}^s for all region $j \in \mathcal{J}$ and supplier $i \in \mathcal{I}$ are calculated as \$3 per mile times the distance in miles divided by the number of test kits a truck can carry. In the state-level test kit allocation problem, we examine three cases of the unit penalty c_{jt}^u for unmet demand such that we can prioritize population groups or regions during different time periods. In particular, in Case (i), we set c_{jt}^u to a constant, i.e., $c_{jt}^u = 100$ for all j and t ; in Case (ii), we let c_{jt}^u be linearly dependent on the projected demand median for region j at period t and set $c_{jt}^u = d_{jt}^{\text{median}} + 10$, indicating that we add a small constant to the median as some regions' demand medians are 0 in early phases and also penalize demand loss; and in Case (iii), we set $c_{jt}^u = 0.001d_j^{\text{elder}}$ with d_j^{elder} being the number of people above 65 years old in region j . In the latter two cases, we prioritize regions with more infections and more elderly or high-risk population groups, respectively.

To compare the severity of disease transmission in different regions and phases, we display the weekly projected demand median d_{jt}^{median} for 8 regions in Michigan during the four phases as well as 65 or older populations in each region in Table 10. As can be seen from Table 10, Regions 7 and

Table 10: Projected demand median during different phases and population above 65 for 8 regions

Regions	Region 1	Region 2	Region 3	Region 4	Region 5	Region 6	Region 7	Region 8
Demand in Apr. 2020	153	3.6K	2.2K	107	2	407	31	0
Demand in June 2020	43	236	516	61	112	443	0	0
Demand in Dec. 2020	55.5K	100.0K	93.3K	72.7K	52.1K	76.5K	18.8K	12.7K
Demand in Sept. 2021	48.5K	76.0K	72.2K	51.0K	56.4K	83.3K	14.3K	15.8K
Population above 65	181.8K	360.2K	349.2K	216.9K	143.5K	237.2K	84.2K	66.5K

8 have significantly fewer number of new cases each day and senior populations than other regions, while Regions 2 and 3 have the highest projected demand and senior population and thus will be most prioritized when we follow the penalty cost cases (ii) and (iii).

5.2.2 Results

We first obtain optimal solutions of the three models (DT, SP, DRO) based on 100 in-sample scenarios and 1000 out-of-sample scenarios following the same normal distribution with the estimated mean and variance. For DRO, we set parameters in the ambiguity set as $\epsilon_{jt}^\mu = 0.5\mu_{jt}$, $\epsilon_{jt}^S = 0.1$, $\bar{\epsilon}_{jt}^S = 2$ for all $j \in \mathcal{J}$, $t \in \mathcal{T}$ in June 2020, December 2020 and September 2021 and set $\epsilon_{jt}^\mu = \mu_{jt}$, $\epsilon_{jt}^S = 0.01$, $\bar{\epsilon}_{jt}^S = 10$ for all $j \in \mathcal{J}$, $t \in \mathcal{T}$ in April 2020. Letting the penalty parameter $c_{jt}^u = 0.001d_j^{\text{elder}}$ for each region $j \in \mathcal{J}$ (i.e., Case (iii) penalty setting), we summarize out-of-sample performance given by solutions of the three approaches under different resource settings in Table 11, where we mark the percentage increase compared with the optimal one in parentheses. The last column in Table 11 records the in-sample computational time in seconds. The raw cost breakdowns of each approach are presented in Tables 16 and 17 in the Appendix.

Table 11: Out-of-sample performance of different approaches in the COVID-19 test kit distribution example

Phases	Approach	Unsatisfied Demand		Overall Cost		Time (sec.)
		Full Resources	Scarce Resources	Full Resources	Scarce Resources	
Apr. 2020	DT	76K (+5.86 × 10 ⁵ %)	76K (+3.89 × 10 ³ %)	\$29M (+68.27%)	\$29M (+68.27%)	0.1
	SP	3K (+2.00 × 10 ⁴ %)	3K (+36%)	\$17M	\$17M	9.5
	DRO	13	2K	\$20M (+17.02%)	\$17M (+1.04%)	12.1
Jun. 2020	DT	5K (+4.24 × 10 ⁴ %)	5K (+4.24 × 10 ⁴ %)	\$2M (+28.79%)	\$2M (+28.79%)	0.1
	SP	691 (+5.78 × 10 ³ %)	691 (+5.78 × 10 ³ %)	\$1M	\$1M	9.4
	DRO	12	12	\$2M (+18.69%)	\$2M (+18.69%)	20.6
Dec. 2020	DT	40K (+2.41 × 10 ⁴ %)	150K	\$22M (+9.75%)	\$34M	0.1
	SP	6K (+3.89 × 10 ³ %)	150K	\$20M	\$34M (+2.80 × 10 ⁻³ %)	5.2
	DRO	164	150K	\$22M (+10.49%)	\$34M (+5.03 × 10 ⁻³ %)	12.0
Sept. 2021	DT	39K	55K	\$21M (+9.44%)	\$22M (+1.32 × 10 ⁻³ %)	0.1
	SP	10K	55K	\$19M	\$22M	9.8
	DRO	0	55K	\$21M (+12.17%)	\$22M (+9.17 × 10 ⁻³ %)	13.0

From Table 11, in the full resource level, DRO always attains the least amount of unsatisfied demand, while DT has the most. Under scarce resources (i.e., 10% of the full resource level), the results almost remain the same in April and June 2020, because most of the approaches do not use up 10% of the total resources. However, when the resource capacity becomes tight such as in December 2020 and September 2021, the results of the three approaches are almost identical as there is not much flexibility in producing and shipping test kits. In terms of overall cost, SP outperforms the other two models in most of the settings.

We display the detailed unsatisfied demand percentages in each region under different cases of penalty cost patterns in Figure 7. Comparing Figure 7 with Table 10, for Cases (ii) and (iii) penalty settings, the demand in Regions 2, 3, 4 and 6 is all satisfied, as these regions have more projected demand in December and more senior populations. On the other hand, Regions 7 and 8 always have the highest unsatisfied demand as they are farther away from the DCs and have less projected demand.

Next, we present the cost breakdown over different phases in Figure 8. From Figure 8, under full resources, DRO can satisfy almost all the demand, while DT still leaves thousands of people untested. However, when we limit the resources to be 10% of the regular amounts (see gray dashed




Figure 7: Unsatisfied demand percentages in each region in December with full resources given Cases (i), (ii), (iii) unit penalty.

line in Figure 8), all capacity costs that exceed the gray line will be truncated to the one that uses only 10% of the original B_t .

6 Conclusion

In this paper, we presented a generic framework to model resource distribution for epidemic response under spatiotemporal demand uncertainties for these resources. Depending on the level of statistical information available for characterizing the uncertainties, which are impacted by various factors including infection trends and demographic behavior, we proposed a stochastic programming and a distributionally robust optimization approach to optimize the locations and capacities of DCs, together with shipping, inventory and demand loss variables. As the proposed optimization frameworks can be applied to different scales of resource distribution, we presented two case studies using instances of COVID-19 vaccine distribution in the US and the analogous test kit distribution in the State of Michigan. Furthermore, we considered different phases of the pandemic, ample or scarce resources, and three cases of unmet demand penalty settings to compare the results of




Figure 8: Cost breakdown of test kit distribution over different phases under ample and scarce resources (in dollars).

deterministic, stochastic and robust optimization approaches.

Our approaches aim efficient and fair distribution of resources by allowing prioritization of regions with more vulnerable population groups or with higher infection susceptibility. The case studies demonstrated the importance of incorporating demand uncertainty in these planning problems as the stochastic programming and distributionally robust optimization approaches outperform the deterministic one in terms of cost and demand coverage. The distributionally robust approach provided a better most-restrictive performance in terms of unvaccinated or untested people who qualify, with an overall cost higher than the one of the stochastic programming approach. Our comparison of the proposed approaches with the CDC’s current distribution status of COVID-19 vaccines shows that our approaches are demand-driven, satisfy more demand in the earlier phases and also prevent unnecessary production, shipping and inventory costs in the later phases of the pandemic when demand drops. Furthermore, the case study of test kit distribution demonstrated the prioritization of resource allocation depending on the infection trend and vulnerable population percentages over the studied regions.

As a limitation of the work, we did not incorporate the time difference of demand for dose 1 and dose 2 of the Pfizer-BioNTech or Moderna vaccines into the optimization models, but only consider the aggregated total demand. For future research, it will be interesting to consider different uncertainties for dose-1 and dose-2 demand, and also incorporate the uncertain time between taking

two doses into the SP and DRO models. One can also build divergence-based ambiguity sets rather than using the moment information [or consider continuous support case](#) in the DRO approach.

Acknowledgement

The authors gratefully acknowledge the partial support from U.S. National Science Foundation (NSF) grants #CMMI-1727618, 2041745 and Department of Energy (DoE) grant #DE-SC0018018.

References

- Babus, A., Das, S., and Lee, S. (2020). The optimal allocation of Covid-19 vaccines. *medRxiv*.
- Basciftci, B., Ahmed, S., and Shen, S. (2021). Distributionally robust facility location problem under decision-dependent stochastic demand. *European Journal of Operational Research*, 292(2):548–561.
- Bertsimas, D., Boussioux, L., Wright, R. C., Delarue, A., Digalakis, V., Jacquillat, A., Kitane, D. L., Lukin, G., Li, M. L., Mingardi, L., Nohadani, O., Orfanoudaki, A., Papalexopoulos, T., Paskov, I., Pauphilet, J., Lami, O. S., Stellato, B., Bouardi, H. T., Carballo, K. V., Wiberg, H., and Zeng, C. (2021a). From predictions to prescriptions: A data-driven response to covid-19. *Health Care Management Science*.
- Bertsimas, D., Digalakis Jr, V., Jacquillat, A., Li, M. L., and Previero, A. (2021b). Where to locate Covid-19 mass vaccination facilities? *arXiv preprint arXiv:2102.07309*.
- Bertsimas, D., Ivanhoe, J. K., Jacquillat, A., Li, M. L., Previero, A., Lami, O. S., and Bouardi, H. T. (2020). Optimizing vaccine allocation to combat the Covid-19 pandemic. *medRxiv*.
- Bidhandi, H. M. and Yusuff, R. M. (2011). Integrated supply chain planning under uncertainty using an improved stochastic approach. *Applied Mathematical Modelling*, 35(6):2618–2630.
- Billingham, S., Widrick, R., Edwards, N. J., and Klaus, S. A. (2020). Covid-19 (sars-cov-2) ventilator resource management using a network optimization model and predictive system demand. *medRxiv*.
- Blanco, V., Gázquez, R., and Leal, M. (2020). Reallocating and Sharing Health Equipments in Sanitary Emergency Situations: The COVID-19 Case in Spain. *arXiv preprint arXiv:2012.02062*.
- Bubar, K. M., Kissler, S. M., Lipsitch, M., Cobey, S., Grad, Y., and Larremore, D. B. (2020). Model-informed Covid-19 vaccine prioritization strategies by age and serostatus. *medRxiv*.
- Cao, H. and Huang, S. (2012). Principles of scarce medical resource allocation in natural disaster relief: A simulation approach. *Medical Decision Making*, 32(3):470–476.

- Centers for Disease Control and Prevention (2021a). COVID-19 Vaccine Distribution Allocations by Jurisdiction - Moderna. <https://data.cdc.gov/Vaccinations/COVID-19-Vaccine-Distribution-Allocations-by-Juris/b7pe-5nws>.
- Centers for Disease Control and Prevention (2021b). COVID-19 Vaccine Distribution Allocations by Jurisdiction - Pfizer. <https://data.cdc.gov/Vaccinations/COVID-19-Vaccine-Distribution-Allocations-by-Juris/saz5-9hgg>.
- Centers for Disease Control and Prevention (2021c). COVID Data Tracker: COVID-19 Vaccinations in the United States. <https://covid.cdc.gov/covid-data-tracker/#vaccinations>.
- Chen, X., Li, M., Simchi-Levi, D., and Zhao, T. (2020). Allocation of covid-19 vaccines under limited supply. *Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3678986>*.
- Daskin, M. S. (2011). *Network and discrete location: Models, algorithms, and applications*. John Wiley & Sons.
- Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612.
- Duijzer, L. E., van Jaarsveld, W., and Dekker, R. (2018). Literature review: The vaccine supply chain. *European Journal of Operational Research*, 268(1):174–192.
- Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., and Phillips, J. P. (2020). Fair allocation of scarce medical resources in the time of covid-19. *New England Journal of Medicine*, 382(21):2049–2055.
- Georgiadis, M. C., Tsiakis, P., Longinidis, P., and Sofioglou, M. K. (2011). Optimal design of supply chain networks under uncertain transient demand variations. *Omega*, 39(3):254–272.
- Golan, M. S., Trump, B. D., Cegan, J. C., and Linkov, I. (2020). The vaccine supply chain: A call for resilience analytics to support Covid-19 vaccine production and distribution. *arXiv preprint arXiv:2011.14231*.
- Gupta, S., Starr, M. K., Farahani, R. Z., and Matinrad, N. (2016). Disaster management from a pom perspective: Mapping a new domain. *Production and Operations Management*, 25(10):1611–1637.
- Huang, H.-C., Singh, B., Morton, D., Johnson, G., Clements, B., and Meyers, L. (2017). Equalizing access to pandemic influenza vaccines through optimal allocation to public health distribution points. *PLoS ONE*, 12(8):e0182720.
- Jiang, R. and Guan, Y. (2016). Data-driven chance constrained stochastic program. *Mathematical Programming*, 158(1):291–327.
- Kleywegt, A. J., Shapiro, A., and Homem-de Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502.

- Lampariello, L. and Sagratella, S. (2021). Effectively managing diagnostic tests to monitor the COVID-19 outbreak in Italy. *Operations Research for Health Care*, 28:100287.
- Liu, K., Li, Q., and Zhang, Z.-H. (2019). Distributionally robust optimization of an emergency medical service station location and sizing problem with joint chance constraints. *Transportation Research Part B: Methodological*, 119:79 – 101.
- Lu, M., Ran, L., and Shen, Z.-J. M. (2015). Reliable facility location design under uncertain correlated disruptions. *Manufacturing & Service Operations Management*, 17(4):445–455.
- Malik, A. A., McFadden, S. M., Elharake, J., and Omer, S. B. (2020). Determinants of covid-19 vaccine acceptance in the us. *EClinicalMedicine*, 26:100495.
- Mehrotra, S., Rahimian, H., Barah, M., Luo, F., and Schantz, K. (2020). A model of supply-chain decisions for resource sharing with an application to ventilator allocation to combat COVID-19. *Naval Research Logistics (NRL)*, 67(5):303–320.
- Moderna, Inc. (2021). Storage and Handling of Moderna Vaccines. <https://www.modernatx.com/covid19vaccine-eua/providers/storage-handling>.
- Mohajerin Esfahani, P. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166.
- Parker, F., Sawczuk, H., Ganjkanloo, F., Ahmadi, F., and Ghobadi, K. (2020). Optimal resource and demand redistribution for healthcare systems under stress from covid-19. *arXiv preprint arXiv:2011.03528*.
- PATH, World Health Organization (WHO) (2013). Delivering Vaccines: A cost comparison of in-country vaccine transport container options. https://path.azureedge.net/media/documents/TS_opt_in_country_transport_rpt.pdf.
- Pei, S. and Shaman, J. (2020). Initial simulation of sars-cov2 spread and intervention effects in the continental us. *MedRxiv*.
- Santini, A. (2021). Optimising the assignment of swabs and reagent for pcr testing during a viral epidemic. *OMEGA*, 102:102341.
- Schütz, P., Tomasgard, A., and Ahmed, S. (2009). Supply chain design under uncertainty using sample average approximation and dual decomposition. *European Journal of Operational Research*, 199(2):409–419.
- Shen, Z.-J. M., Zhan, R. L., and Zhang, J. (2011). The reliable facility location problem: Formulations, heuristics, and approximation algorithms. *INFORMS Journal on Computing*, 23(3):470–482.

- Snyder, L. V. (2006). Facility location under uncertainty: A review. *IIE Transactions*, 38(7):537–554.
- Tebbens, R. J. D. and Thompson, K. M. (2009). Priority shifting and the dynamics of managing eradicable infectious diseases. *Management Science*, 55(4):650–663.
- Tolooie, A., Maity, M., and Sinha, A. K. (2020). A two-stage stochastic mixed-integer program for reliable supply chain network design under uncertain disruptions and demand. *Computers & Industrial Engineering*, 148:106722.
- Wan, X., Wang, W., Liu, J., and Tong, T. (2014). Estimating the sample mean and standard deviation from the sample size, median, range and/or interquartile range. *BMC Medical Research Methodology*, 14(1):1–13.
- Wang, C. and Chen, S. (2020). A distributionally robust optimization for blood supply network considering disasters. *Transportation Research Part E: Logistics and Transportation Review*, 134:101840.
- Wang, W., Wu, Q., Yang, J., Dong, K., Chen, X., Bai, X., Chen, X., Chen, Z., Viboud, C., Ajelli, M., and Yu, H. (2020). Global, regional, and national estimates of target population sizes for covid-19 vaccination: descriptive study. *BMJ*, 371.
- Yin, X., Büyüktaktakın, İ. E., and Patel, B. P. (2021). Covid-19: Data-driven optimal allocation of ventilator supply under uncertainty and risk. *European Journal of Operational Research*.
- Zhang, Y., Jiang, R., and Shen, S. (2018a). Ambiguous chance-constrained binary programs under mean-covariance information. *SIAM Journal of Optimization*, 28(4):2922–2944.
- Zhang, Y., Shen, S., and Erdogan, S. A. (2018b). Solving 0–1 semidefinite programs for distributionally robust allocation of surgery blocks. *Optimization Letters*, 12(7):1503–1521.

Appendix: Parameter Sources and More Results

We first present the list of parameters and decision variables of the SMIP model (1), and sources of parameters in vaccine distribution and test kit allocation problems in Tables 12 and 13, respectively.

Next, we present the detailed results of cost breakdown in the national vaccine allocation and state-level test kit distribution problems under different phases, resource settings and approaches.

Table 12: List of sets, parameters and decision variables.

Name	Definition
Sets:	
$\mathcal{I}$	Sets of potential sites for DCs
$\mathcal{J}$	Sets of demand sites
$\mathcal{T}$	Sets of time periods in the planning horizon
Parameters:	
c_i^o	Cost of operating DC at location i
c_i^h	Cost of installing unit capacity for DC at location i
c_{ijt}^s	Cost of shipping an unit from DC at location i to demand site j at period t
c_{jt}^I	Penalty cost of unit unsatisfied demand at demand site j at period t
c_{jt}^u	Holding cost of an unit at demand site j at period t
B_t	Total resource capacity at period t
M_i	Capacity limit of DC at location i
d_{jt}	Demand of demand site j at period t
Decision variables:	
x_i	Binary variable denoting whether DC at location i is open
h_{it}	Capacity of DC at location i at period t
s_{ijt}	Amount of resources sent from DC at location i to demand site j at period t
I_{jt}	Amount of inventory at demand site j at period t
u_{jt}	Unsatisfied demand amount at demand site j at period t

Table 13: Parameter sources in vaccine distribution and test kit allocation

Parameter	Vaccine Distribution	Test Kit Allocation
c_i^o	Warehouse rental price	Warehouse rental price
c_i^h	COVID-19 vaccine cost	COVID-19 test kit cost
c_{ijt}^s	Vans travel cost + Vaccine dimensions + Vaccine deliver cost	Vans travel cost + Vans capacity
c_{jt}^I	Warehouse storage cost + Energy cost	Warehouse storage cost + COVID-19 test kit dimension
c_{jt}^u	100	Three penalty settings as discussed in the paper

Table 14: Cost breakdown of vaccine distribution under full resources, different phases and approaches

Phases	Approaches	Operating	Capacity	Shipping	Inventory	Penalty	Overall Cost
Phase 1	DT	\$97K	\$486M	\$39K	\$34	\$37M	\$524M
	SP	\$63K	\$496M	\$48K	\$155	\$6M	\$503M
	DRO	\$76K	\$515M	\$55K	\$80	\$292K	\$516M
Phase 2	DT	\$194K	\$3B	\$145K	\$259	\$232M	\$4B
	SP	\$175K	\$3B	\$125K	\$752	\$17M	\$3B
	DRO	\$107K	\$3B	\$545K	\$519	\$9M	\$3B
Phase 3	DT	\$290K	\$4B	\$180K	\$569	\$262M	\$4B
	SP	\$263K	\$4B	\$152K	\$2K	\$31M	\$4B
	DRO	\$259K	\$4B	\$411K	\$935	\$45M	\$4B

Table 15: Cost breakdown of vaccine distribution under scarce resources, different phases and approaches

Phases	Approaches	Operating	Capacity	Shipping	Inventory	Penalty	Overall Cost
Phase 1	DT	\$97K	\$486M	\$39K	\$34	\$37M	\$524M
	SP	\$63K	\$496M	\$52K	\$99	\$7M	\$503M
	DRO	\$54K	\$512M	\$58K	\$64	\$687K	\$513M
Phase 2	DT	\$105K	\$2B	\$48K	\$0	\$17B	\$19B
	SP	\$96K	\$2B	\$42K	\$0	\$17B	\$19B
	DRO	\$105K	\$2B	\$41K	\$0	\$17B	\$19B
Phase 3	DT	\$256K	\$3B	\$122K	\$0	\$13B	\$16B
	SP	\$175K	\$3B	\$134K	\$0	\$13B	\$16B
	DRO	\$209K	\$3B	\$221K	\$0	\$13B	\$16B

Table 16: Cost breakdown of test kit distribution under full resources, different phases and approaches

Phases	Approaches	Operating	Capacity	Shipping	Inventory	Penalty	Overall Cost
Apr. 2020	DT	\$2K	\$13M	\$4K	\$3	\$16M	\$29M
	SP	\$4K	\$17M	\$6K	\$9	\$395K	\$17M
	DRO	\$4K	\$20M	\$5K	\$2	\$2K	\$20M
Jun. 2020	DT	\$2K	\$1M	\$5K	\$2	\$638K	\$2M
	SP	\$2K	\$1M	\$5K	\$9	\$76K	\$1M
	DRO	\$2K	\$2M	\$7K	\$0	\$1K	\$2M
Dec. 2020	DT	\$4K	\$18M	\$5K	\$2	\$4M	\$22M
	SP	\$4K	\$20M	\$5K	\$6	\$535K	\$20M
	DRO	\$4K	\$22M	\$5K	\$0	\$12K	\$22M
Sept. 2021	DT	\$4K	\$17M	\$5K	\$2	\$3M	\$21M
	SP	\$4K	\$18M	\$5K	\$7	\$811K	\$19M
	DRO	\$5K	\$21M	\$5K	\$2	\$0	\$21M

Table 17: Cost breakdown of test kit distribution under scarce resources, different phases and approaches

Phases	Approaches	Operating	Capacity	Shipping	Inventory	Penalty	Overall Cost
Apr. 2020	DT	\$2K	\$13M	\$4K	\$3	\$16M	\$29M
	SP	\$4K	\$17M	\$6K	\$9	\$395K	\$17M
	DRO	\$5k	\$17M	\$4K	\$12	\$282K	\$17M
Jun. 2020	DT	\$2K	\$1M	\$5K	\$2	\$638K	\$2M
	SP	\$2K	\$1M	\$5K	\$9	\$76K	\$1M
	DRO	\$2K	\$2M	\$7K	\$0	\$1K	\$2M
Dec. 2020	DT	\$4K	\$17M	\$3K	\$0	\$17M	\$34M
	SP	\$4K	\$17M	\$3K	\$0	\$17M	\$34M
	DRO	\$5K	\$17M	\$3K	\$0	\$17M	\$34M
Sept. 2021	DT	\$4K	\$17M	\$4K	\$4	\$5M	\$22M
	SP	\$4K	\$17M	\$4K	\$4	\$5M	\$22M
	DRO	\$5K	\$17M	\$4K	\$4	\$5M	\$22M