

Data-driven Fair Resource Allocation For Novel Emerging Epidemics: A COVID-19 Convalescent Plasma Case Study

Maryam Akbari-Moghaddam^{1,*}, Na Li^{1,2,3}, Douglas G. Down¹, Donald M. Arnold^{3,4}, Jeannie Callum^{5,6}, Philippe Bégin^{7,8}, Nancy M. Heddle^{3,4}.

1 Department of Computing and Software, McMaster University, Hamilton, Ontario L8S 4L7, Canada

2 Community Health Sciences, University of Calgary, Calgary, Alberta T2N 1N4, Canada

3 McMaster Centre for Transfusion Research, Department of Medicine, McMaster University, Hamilton, Ontario L8N 3Z5, Canada

4 Department of Medicine, McMaster University, Hamilton, Ontario L8N 3Z5, Canada

5 Department of Pathology and Molecular Medicine, Kingston Health Sciences Centre, Kingston, Ontario K7L 2V7, Canada

6 Department of Pathology and Molecular Medicine, Queen's University, Kingston, Ontario K7L 3N6, Canada

7 Section of Allergy, Immunology and Rheumatology, Department of Pediatrics, CHU Sainte-Justine, Montreal, Quebec H3T 1C5, Canada

8 Department of Medicine, CHUM, Université de Montréal, Montreal H3T 1J4, Quebec, Canada

* Corresponding author

Email: akbarimm@mcmaster.ca

ABSTRACT

Epidemics are a serious public health threat, and the resources for mitigating their effects are typically limited. Decision-makers face challenges in forecasting the supply and demand for these resources as prior information about the disease is often not available, the behaviour of the disease can periodically change (either naturally or as a result of public health policies) and can differ by geographical region. Randomized controlled trials (RCTs) using scarce resources such as blood products as a randomized intervention are affected by epidemics. In this work, we discuss a model that is suitable for short-term real-time supply and demand forecasting during emerging outbreaks. We consider a case study of demand forecasting and allocating scarce quantities of COVID-19 Convalescent Plasma (CCP) in an international multi-site RCT involving multiple hospital hubs across Canada (excluding Québec). We propose a data-driven mixed-integer programming (MIP) resource allocation model that assigns available resources to maximize a notion of fairness among the resource-demanding entities. Numerical results from applying our MIP model to the case study suggest that our approach can help balance the supply and demand of limited products such as CCP and minimize the unmet demand ratios of the demand entities. We analyze the sensitivity of our model to different allocation settings and show that our model assigns equitable allocations across the entities.

Keywords Resource Allocation; Epidemics; COVID-19 Convalescent Plasma; Data-driven Optimization; Demand Forecasting

1 Introduction

Epidemics have impacted the world many times, and will occur again in the future. Emergency responses to these epidemics can either be pre-event or post-event, where here event refers to the initiation of an infectious disease (Dasaklis, Pappis, & Rachaniotis, 2012). Forecasting the potential dangers and planning the necessary steps to deal with an epidemic are considered as pre-event tasks. Post-event responses occur after the disease has started spreading and is still in progress. The corresponding actions at these points are associated with treatment and allocating the corresponding available resources. Our focus in this paper is on post-event response situations.

Resource allocation decisions during emerging epidemics are challenging due to several key reasons. First, limited knowledge and historical data about disease demographics make it difficult to predict the demand for particular resources. Second, because epidemics are usually unexpected and can spread rapidly, there are often limited health

This paper has been submitted to the INFOR: Information Systems and Operational Research journal.

care resources (vaccines, blood products, medical equipment, etc.) compared to the total number of entities requesting them. Determining how to fairly allocate the limited resources becomes a challenge. Third, the demand can vary significantly between geographically dispersed entities. Finally, the decisions must be made in a timely manner. These issues motivate the investigation of supply and demand forecasting models for scenarios where there are small amounts of available data and the data can exhibit fundamental changes in behaviour. It is of interest to incorporate these models into algorithms that yield fair allocations.

Clinical trials can have different durations depending on the underlying research, with many lasting as short as several months (Piantadosi, 2017). Some of these clinical trials may require rapid roll-out to different participating sites right after the planning phase, which might not allow for prior data collection by the sites. Randomized Controlled Trials (RCTs) can face additional limitations in data collection within the study duration as different experimental interventions can be associated with specific time marks. During uncertain situations like emerging epidemics, RCTs that use scarce resources such as blood products as a randomized intervention can be affected when there is a need for a timely and fair resource allocation (Stawicki et al., 2020). Convalescent plasma is a blood product which has been used as a potential treatment for a number of diseases such as Ebola (Kraft et al., 2015; Van Griensven et al., 2016), influenza (Hung et al., 2011; Zhou, Zhong, & Guan, 2007), and COVID-19 (Chen, Xiong, Bao, & Shi, 2020). COVID-19 Convalescent Plasma (CCP) contains antibodies, or special proteins, generated by the body’s immune system in response to the novel coronavirus (Chen et al., 2020). It has been considered as an experimental treatment for hospitalized COVID-19 patients in a number of RCTs worldwide (Bégin et al., 2021; Gharbharan et al., 2020; L. Li et al., 2020; Simonovich et al., 2021). Despite the importance of blood supply and demand management for scarce products in RCTs, there is a lack of published research that provides insights on the problem. This can be related to several challenges encountered when making decisions on CCP allocation in an RCT setting. When such a scarce blood product serves as a randomized intervention in an RCT, poor linkage across blood donor recruitment, inventory management and the randomization process can raise biases (N. Li et al., 2022).

In this work, we tackle the real-time allocation of scarce resources during epidemics to entities located in widespread geographical locations and with different resource requirements. Furthermore, we consider an RCT setting which requires rapid roll-out and has limited historical data due to its short study duration. We are interested in demand forecasting models that can predict short-term demand in real-time. The demand forecasts directly impact the resource allocation decisions as inaccurate forecasts may lead to inefficient and unfair use of limited (and valuable) resources. We note that there are different notions of fairness (balance) in terms of resource allocation in the literature (Kumar & Kleinberg, 2000). In our case, we define fairness as minimization of the entities’ unmet demand ratios, but one could also generalize fairness to other notions. For instance, Karsu and Erkan (2020) propose an approach where imbalance is defined as the deviation from a reference distribution determined by the decision-maker. They show that in resource allocation problems, it is possible to maintain a mixed-integer programming (MIP) structure even after generalizing the notion of fairness. In general, there are wide-ranging views of fairness and how fairness metrics can be incorporated into optimization problems (see Heitmann & Brüggemann, 2014; Karsu & Morton, 2014, 2015; Mestre, Oliveira, & Barbosa-Póvoa, 2012; Smith, Harper, & Potts, 2013). We evaluate our proposed model on a case study of demand forecasting and CCP distribution in a multi-site RCT, where (i) there were limited historical supply and demand data, (ii) the supply was limited and restricted by manufacturing policies and donor recruitment, (iii) the demand arising from the demand entities was heterogeneous, and (iv) specific clinical requirements were needed for administering CCP transfusion. More details of the RCT, including how the approach presented in this paper is embedded and the resulting observations from a clinical perspective, can be found in (N. Li et al., 2022).

We make the following contributions: First, we discuss real-time short-term forecasting of supply and demand of scarce resources in epidemics with limited historical data. We propose the use of a forecasting model that does not require indeterminate epidemiological parameters (such as location and time-specific parameters) and thus does not require periodically updating the parameters. Secondly, we address challenges that may arise in an online setting due to extrapolation and sparse data. Next, we propose a data-driven MIP model for real-time multi-location allocation of scarce resources regularly and fairly to entities, which have heterogeneous demand. This approach maximizes a notion of fairness among the resource-demanding entities. Finally, numerical results of applying our model in a CCP case study show that our approach yields fair allocations that are both close to the scenario where supply and demand are known (rather than forecast) and are preferable to what was used in practice.

The rest of the paper is organized as follows. Section 2 presents the existing literature on demand forecasting and resource allocation approaches during infectious disease outbreaks and our motivation for this work. We describe our data-driven resource allocation problem in Section 3.1. Section 3.2 discusses in depth the supply and demand forecasting methods that we use and we define our proposed MIP resource allocation model in Section 3.3. The CCP case study, the numerical results of applying our MIP model to the case study, and sensitivity analysis of the model are discussed in Section 4. We conclude this work and discuss how it may inform responses to future pandemics in Section 5.

2 Motivation and Related Work

There has been a variety of work in various fields that tackle the problem of resource allocation during infectious disease outbreaks. In what follows, we divide the body of work into six categories and discuss the motivation behind our proposed methodology.

Compartmental Models and Resource Allocation

Epidemiological compartmental models, consisting of a set of nonlinear ordinary differential equations, can help model the dynamics of different epidemiological variables during a pandemic (Brauer, 2008). These models can give insight into disease-related information such as spread rate, the duration of an epidemic, and the total number of infected and recovered patients. Decision-makers can employ compartmental models to derive demand for medical resources to guide resource allocation decisions. Focusing on the COVID-19 outbreak (Velavan & Meyer, 2020), there have been many applications and tools developed by different organizations worldwide to forecast infections, hospitalizations, and deaths using compartmental models (Chowdhury et al., 2021; Gong et al., 2020; Tomar & Gupta, 2020). For instance, CHIME (Weissman et al., 2020) is a tool based on a Susceptible-Infectious-Recovered (SIR) model that can be used for forecasting the number of daily hospitalized COVID-19 patients in the short-term (e.g., up to 30 days).

Demand Forecasting Models and Resource Allocation

Another approach that researchers have studied for forecasting healthcare resources is using time series models or machine learning methods. The references for this approach are extensive, thus we only discuss a few studies as examples. Ferstad et al. (2020) introduce a time series model to forecast the availability and utilization of intensive and acute care beds. Nikolopoulos, Punia, Schäfers, Tsinopoulos, and Vasilakis (2021) use epidemiological and deep learning models to forecast the excess demand for products and services considering auxiliary data and simulating governmental decisions, while N. Li, Chiang, Down, and Heddle (2021) combine ideas from statistical time series modelling and machine learning to develop a hybrid demand forecasting model for red blood cell components using clinical predictors.

Segmented Regression Models and Resource Allocation

Piecewise Linear Regression (PLR), also known as segmented linear regression, forecasting models are a special case of a larger set of models known as spline functions (Suits, Mason, & Chan, 1978). Modelling the regression function in "pieces" can be helpful when dealing with sparse data because we can still use linear regression models for data that does not fit a single line. To be more specific, PLR is a simple model that makes understanding the data easier by solving several linear regressions. Points at which the behaviour changes are called breakpoints, which act as boundaries between each piece. There have been a few studies on finding the number of breakpoints and their locations. Rosen and Pardalos (1986) propose a method for finding the minimum number of equally spaced breakpoints within a given error tolerance, a sequential method is proposed in Strikholm (2006) for finding the number of breakpoints, and Yang, Liu, Tsoka, and Papageorgiou (2016) propose a discontinuous piecewise linear approximation and how to determine the optimal breakpoint locations.

Piecewise linear models have been used in different applications when modelling the structural shifts in data and forecasting based on the most recent behaviour in data is desired. We will discuss a PLR model, namely MARS, in more detail in Section 3.2.1, where we discuss the supply and demand forecasting model used in this work. The MARS model has been used in healthcare for medical diagnosis using classification approaches (see Butte et al., 2010; Chou, Lee, Shao, & Chen, 2004; Lasheras et al., 2020; Senthilkumar & Paulraj, 2013; Serrano, Sánchez, Lasheras, Iglesias-Rodríguez, & Valverde, 2020; Yao, Yang, & Zhan, 2013). In the area of time series forecasting, López-Lozano et al. (2019) evaluate the generalizability of MARS for identifying thresholds for antibiotic consumption and Katris (2021) studies MARS and other time series approaches for predicting the evolution of reported COVID-19 cases to track the outbreak in Greece.

Disaster Management and Resource Allocation

In terms of resource allocation, the literature on emergency logistics involves disaster operations performed before or after the impact of a disaster where short-term post-disaster planning begins immediately. This includes activities among which relief distribution is more related to our scope of the study. Relief distribution consists of providing medical supplies, shelters, manpower, and other associated resources, and thus, resource allocation is a significant component (Caunhye, Nie, & Pokharel, 2012). A number of resource allocation models for relief distribution that consider resource assignment without determining flow quantities among arcs focus on allocating equipment to amend

maritime disasters, such as cleaning up oil spills. Unlike other disasters, oil spills create a single demand point, and studies in this area only consider direct equipment assignment to the spill areas (see Psaraftis, Tharakan, & Ceder, 1986; Srinivasa & Wilhelm, 1997; Wilhelm & Srinivasa, 1997).

Optimization Models and Resource Allocation

Many studies have focused on developing allocation models for medical resources during infectious disease outbreaks. A dynamic linear programming model based on an epidemic diffusion model is introduced in Liu and Liang (2013) to allocate medical resources. Preciado, Zargham, Enyioha, Jadbabaie, and Pappas (2013) analyze a networked version of a susceptible-infected-susceptible (SIS) epidemic model when different susceptibility levels are present. They propose a convex optimization approach for distributing vaccination resources in a cost-optimal manner and test their approach in a real social network. Yarmand, Ivy, Denton, and Lloyd (2014) consider two-phase vaccine allocation to different geographical locations. They capture each region's epidemic dynamics for different vaccination phases by a two-stage stochastic linear program (2-SLP) model and show that their model helps to reduce vaccine production and administration costs. Yin and Büyüktaktakın (2021) define two new equity metrics for fair resource allocation and develop a multi-stage stochastic programming vaccine allocation model and apply it to an Ebola virus disease case study. Furthermore, two resource allocation problems during outbreaks are discussed in Preciado, Zargham, Enyioha, Jadbabaie, and Pappas (2014), where they use geometric programming to solve the problems. Following the work in Preciado et al. (2014), Han, Preciado, Nowzari, and Pappas (2015) propose a data-driven robust optimization framework based on conic geometric programming. Their model is used to determine an optimal allocation of medical resources such as vaccines and antidotes and can help control an SIS viral spreading process in a directed contact network with unknown contact rates.

Other Epidemiological Frameworks and Resource Allocation

A number of works have investigated resource allocation frameworks using outbreak case studies. The problem of scheduling limited available resources between multiple infected areas is discussed in N. P. Rachaniotis, Dasaklis, and Pappis (2012), and their proposed deterministic scheduling model is studied in a case study of mass vaccination against A(H1N1)v influenza. A real-time synchronous heuristic algorithm is proposed in N. Rachaniotis, Dasaklis, and Pappis (2017) and is tested on the same case study as in N. P. Rachaniotis et al. (2012). Sun, DePuy, and Evans (2014) focus on allocating patients and resources between hospitals located in a healthcare network and propose a multi-objective optimization model. They discuss the application of their model in an influenza outbreak case study. Finally, a large integer programming problem framework for optimally allocating a resource donation is introduced in Anparasan and Lejeune (2019), and results of applying the framework to a 2010 cholera outbreak case study are reported. Closely related to our work is Du, Sai, and Kong (2021) where they study a multi-period location-specific resource allocation problem for cholera outbreak intervention. They consider a rolling time horizon and periodically determine an optimal intervention resource allocation strategy with their data-driven optimization approach. Also similar in spirit to our approach is Bekker, uit het Broek, and Koole (2023), who propose a model for making daily short-term predictions of the number of occupied ICU and clinical beds in the Netherlands due to COVID-19.

Motivation Behind Proposed Methodology

With the aforementioned approaches in mind, we now discuss the motivation behind the prediction methodology that we consider for real-time short-term demand and supply forecasting for experimental products in clinical trials. We observed the following limitations with the application of the methodologies above to our problem setting:

1. When an epidemic is first emerging, the epidemic state is only partially observable, and the parameters that the epidemiological models require are often indeterminate, as the disease information can only be obtained over a period of time and after sufficient cases are reported and required data is collected. Moreover, different geographical locations can show various characteristics in terms of the disease spread pattern. The estimates that compartmental models make are sensitive to the model's structure (Sikal, Little, Barnes, & White, 2016). Even in the presence of reasonable parameter estimates and simple model structures, real-time resource allocation requires the parameters to be periodically updated based on the disease spread rate and the number of people involved, whether susceptible or infected. Therefore, CHIME-like models may lead to poor approximation of the actual demand when used in a real-time setting where obtaining the most recent updated parameters may not always be possible.
2. The aforementioned demand forecasting models work best when large datasets are available and the models can capture the trend and seasonality, which is not possible during emerging epidemics. Furthermore, real-time demand forecasting can be challenging with any forecasting model if the demand is affected by external factors, such as

population characteristics, geographical locations, operational procedures, guidelines, and governmental policies. As far as we are aware, none of the mentioned approaches are consistent with the challenges we introduced in Section 1. It is quite difficult to directly apply these approaches to a real-time setting where short-term supply and demand forecasting is desired, there is a limited supply of resources, and the available data is sparse and shows fundamental changes in behaviour during different periods. This motivates us to avoid models that are reliant on a large number of parameters (as for CHIME-like models), as they require continual updates and they may not always yield accurate forecasts for resource supply and demand. Nonetheless, we seek a model that makes reasonable predictions even when facing such challenges. We will revisit these challenges in the case study in Section 4.

3. More sophisticated demand forecasting models, such as ARIMA, can also be used for time series analysis, especially demand forecasting with univariate data. However, satisfaction of basic requirements for such models is necessary which might not be possible in our setting. For instance, ARIMA and ARMA work best in the presence of stationary data and the rule of thumb for these models suggests having at least 50 but preferably more than 100 observations (Box & Tiao, 1975). We cannot meet that requirement, and as we start forecasting with almost no historical data, any forecasting model might fail to exhibit its best performance. We found a simple forecasting model such as MARS to be a good choice considering our non-stationary and limited data.

4. Our study is different from the aforementioned disaster management literature in the sense that although there is the common objective of resource allocation of limited resources, the time scale is different. Epidemics tend to be much longer than disasters such as earthquakes, and there is more uncertainty in the quantity of the effective resources due to the unknown nature of the disease. Furthermore, the evolution of epidemics might involve unexpected changes in supply and demand patterns, requiring suitable forecasting models to be incorporated into the resource allocation model.

This work focuses on (i) how the methodologies above are combined to solve a real-time short-term supply and demand forecasting and resource allocation problem, and (ii) how the combined methodology is applied to a clinical problem setting with various human interventions and limited historical data. We demonstrate that using MARS to determine inputs (supply and demand forecasts) to a resource allocation problem is an effective combination in an emerging epidemic setting. The prediction model in Bekker et al. (2023) consists of a linear programming model inspired by smoothing splines for predicting the arrivals and methods stemming from queueing theory to convert arrivals into occupancy. The motivation for choosing their model is similar to ours in the sense that it works with little historical data, which is a consequence of an emerging epidemic setting. To the best of our knowledge, no other study has tackled the problem of real-time multi-location allocation of scarce resources with limited historical data and without relying on epidemiological models.

3 Data-driven Resource Allocation Model

3.1 Problem Description

Our work considers the patients of a COVID-19 case study (details in Section 4) where the distribution of supply and demand is similar in terms of the population; however, there is a limited and highly-varied pattern for both in the collected data. The supply for this study comes from the CCP of the patients who have recovered from COVID-19, and donated blood. The demand comes from the patients who are currently infected with COVID-19 and require CCP, but will be considered as supply once they recover and donate CCP.

Hub-and-spoke structures can help effective access to blood products in a public health network and clinical trials (Fernandez-Sojo et al., 2022; Roberts et al., 2018). A hub-and-spoke structure was established for our case study where CCP units were distributed from regional blood centers to hospital hubs for subsequent CCP allocation to patients. Therefore, our model also considers a hub-and-spoke structure where we have a centralized supplier (hub) that interacts with H customers (spokes) and is responsible for satisfying their demand for R types of resources. Figure 1 shows a flowchart of our data-driven resource allocation process.

We are interested in a setting where limited resources must be allocated on a regular basis (e.g. every week) to the entities requesting them, the demand for the resources can be heterogeneous and arises from geographically dispersed locations. We forecast both supply and demand for the following reasons: (i) Common source: The same population defines both supply and demand, but at different points in time. (ii) Dependence between supply and demand: Changes in the behaviour of the disease naturally lead to correlated changes in supply and demand. (iii) Uncertainty in supply and demand: A classic resource allocation problem considers known resource quantities. However, we do not know the number of resources that are available to be allocated at every point in time and we face uncertainty in both supply and

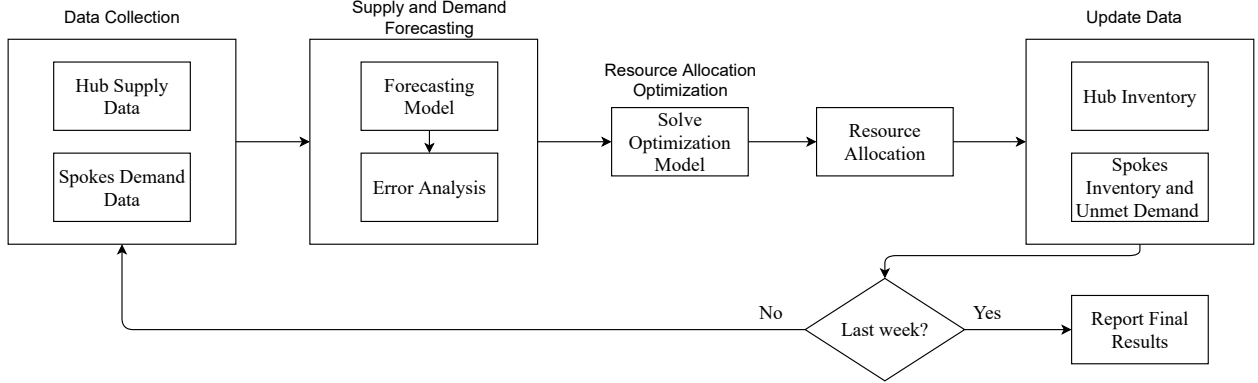


Figure 1. Data-driven Resource Allocation Process

demand due to the unknown nature of the disease. The uncertainty in inventory collection also indicates the impact of CCP collection in our study, resulting in a resource collection and allocation model.

We choose to work with the cumulative supply or demand as a cumulative sum of data samples over time is helpful in situations where one needs to smooth heterogeneous and sparse data and still make quantitative predictions for future supply and demand without altering the original data. Ellaway (1978) investigates the application of the cumulative sum technique in a neurophysiology study. The cumulative sum is shown to be a powerful technique for finding periods of change in data, as well as reducing real-time decision-making uncertainty. In our case, we forecast weekly supply and demand based on historical cumulative supply and demand data for a particular resource prior to resource allocations. In particular, we would like to fit a forecasting model weekly to our data where for each week, only the most recent week's observation is added to the dataset, and we cannot modify the predictions the model previously made (as real-time allocations are made based on the forecasts). We then perform error analysis of the forecasts, and allocate the available resources to the customers in a fair manner by solving an appropriate optimization problem. We assume that both the supplier and customers can hold inventories of the resources and can use them to satisfy future demand. Once the supplier allocates resources to customers, the decision is final, and the resources cannot be reassigned. Thus, we need to update the inventories at the end of each week, and consider them when solving the optimization model in the next week.

We discuss a PLR supply and demand forecasting model and our resource allocation optimization model in more detail in Section 3.2 and Section 3.3, respectively.

3.2 Supply and Demand Forecasting

We would like to forecast the supply and demand for a particular week $t + 1$ where we only have data available up to week t . Consider a dataset consisting of t observations where X is an array of elements x_i ($i = 1, 2, \dots, t$) representing the input variables (in our case, $X = 1, 2, 3, \dots, t$). Consider y to be a vector of observed (supply or demand) data samples y_i ($i = 1, \dots, t$). We use Multivariate Adaptive Regression Splines (MARS), a PLR model, to forecast the supply and demand.

3.2.1 Multivariate Adaptive Regression Splines (MARS)

Multivariate Adaptive Regression Spline (MARS) is a nonparametric regression approach that was introduced by Friedman (1991). The MARS model consists of a collection of simple linear models that can capture patterns and trends related to interactions and nonlinearities. MARS uses a series of piecewise linear pieces (splines) of different gradients. These pieces, also known as basis functions (BFs), are connected at positions called knots which allow thresholds, bends, and other departures from linear functions. A MARS model is specified as follows:

$$\hat{y}' = \beta_0 + \sum_{p=1}^P \beta_p \lambda_p(x),$$

where P is the number of BFs and each $\lambda_p(x)$ is a BF, which can be a spline function or the product of two or more spline functions. The parameters β_0 and β_p , $p = 1, \dots, P$ are estimated using the least-squares method. The basis functions are described as:

$$BF(x) = \{\max(0, x - c_p), \max(0, c_p - x)\},$$

where c_p is the knot of the spline (threshold value).

A forward stage and a backward stage are considered in the MARS algorithm. In the forward stage, BF functions and their potential knots are chosen, which may result in a complicated and over parameterized model. In the backward stage, to prevent overfitting, the algorithm considers deleting the BFs in increasing order of the amount that they reduce the training error (Friedman, 1991).

3.2.2 Error Analysis and Model Enhancement

Based on the slope of the last piece, the forecast supply and demand for week $t + 1$ under MARS is simply:

$$\hat{y}'_{t+1} = \hat{y}'_t + \frac{\hat{y}'_t - \hat{y}'_{t-1}}{x_t - x_{t-1}}.$$

We can improve the forecast value for week $t + 1$ (\hat{y}'_{t+1}) by calculating its forecast error (Jarrett & Khumuwala, 1987; Lu & Chen, 2020). We first fit an autoregressive (AR) model of order l using Conditional Maximum Likelihood to the residual error ($\epsilon_i = y_i - \hat{y}'_i$) data up to week t and use it to forecast the error for week $t + 1$:

$$\hat{\epsilon}_{t+1} = a_0 + a_1 \epsilon_t + a_2 \epsilon_{t-1} + \dots + a_l \epsilon_{t-l+1}, \quad (1)$$

where a_0, a_1, \dots, a_l are the coefficients obtained from the AR model and $\hat{\epsilon}_{t+1}$ is the forecast error for week $t + 1$. We finally calculate $\hat{y}_{t+1} = \hat{y}'_{t+1} + \hat{\epsilon}_{t+1}$, i.e., the improved supply or demand forecast value for week $t + 1$, and use it as our final forecast value for that week.

3.2.3 Challenges

We now discuss general challenges that may be faced when MARS is employed in an online manner:

- Predictions made using MARS may degrade if there is a sudden transitory change in data. For instance, we found that holidays may affect the amount of data collected in a particular week but the data follows its previous pattern after the holidays have passed. We discuss this issue in greater detail in the case study discussed in Section 4.
- MARS is fitting piecewise linear models, and thus the slope of the segment that ends with the most recent observation has a significant effect on the forecast for the next week (it may lead to a large underestimation/overestimation).

Although working with the cumulative sums for forecasting future supply and demand can help address the issue of data sparsity, a particular challenge arises when cumulative sums are employed in an online setting. Consider a situation where the demand before week t has a steep slope resulting in a (relatively) high forecast for week t . However, upon observing the demand for week t , one finds out that the actual demand was considerably lower. Thus, the slope that affects week $t + 1$'s forecast demand will be less steep than what it was when only data up to week $t - 1$ was available. This may result in the forecast cumulative demand for week $t + 1$ being lower than the previously forecast cumulative demand for week t , which is not possible. For instance, Figure 2a and Figure 2b demonstrate the cumulative forecast demand obtained by MARS for week 21 and week 22, respectively. The cumulative forecast demand for week 21 and week 22 are 22 and 19, respectively, under MARS, which cannot happen in practice. This is one consequence of considering an online setting where modifying the predictions is not possible as real-time decisions are made. The sparser the data, the more this issue can affect the model's predictions. A possible solution for dealing with such situations is to assume that the cumulative forecast value for week $t + 1$ is equal to the cumulative forecast value for week t . We will observe this issue in our case study discussed in Section 4 and deal with it in the suggested manner.

3.3 Resource Allocation Optimization

We allow for demand for resource r to be satisfied by resource r' . This can be represented with a matrix C of size $R \times R$, where an element (r, r') of C is 1 if demand for resource r can be satisfied by resource r' and is 0 otherwise. Furthermore, the supply used in our MIP model for resource r for week t is constrained as follows:

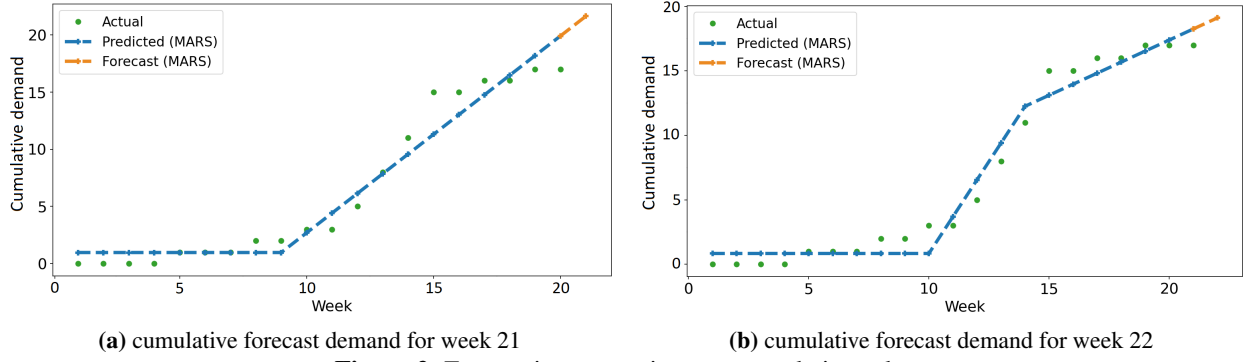


Figure 2. Forecasting a negative non-cumulative value

$$c_{r,t} = \begin{cases} \hat{s}_{r,t}, & \text{if } \hat{s}_{r,t} \leq s_{r,t} \\ s_{r,t}, & \text{otherwise,} \end{cases} \quad (2)$$

where $s_{r,t}$ and $\hat{s}_{r,t}$ are the actual and forecast supply for resource r on week t , respectively. The second case in (2) indicates that we cannot allocate resources beyond the actual available supply. We will formulate our resource allocation as a Mixed Integer Program (MIP). To do so, we require the following notation:

Indices

- t index of time periods, $t = 1, \dots, T$
- h index of customers, $h = 1, \dots, H$
- r, r' index of resource, $r, r' = 1, \dots, R$

Data

- $c_{r,t}$ the amount of resource r available at the supplier for assignment at time t (see (2) above)
- $i_{r,h,t-1}$ the inventory of resource r stored at customer h at time $t - 1$
- $\hat{d}_{r,h,t}$ the estimated demand for resource r by customer h at time t

Decision variables

- $v_{r,r',h,t}$ the number of units of resource r' assigned to customer h to satisfy demand for resource r at time t . We only consider the set of $v_{r,r',h,t}$ that correspond to $C(r, r') = 1$.

We formulate our resource allocation problem as follows:

Objective function

$$\min \max_{h: \sum_{r=1}^R \hat{d}_{r,h,t} > 0} \frac{\sum_{r=1}^R (\hat{d}_{r,h,t} - \sum_{r'=1}^R v_{r,r',h,t} - i_{r,h,t-1})}{\sum_{r=1}^R \hat{d}_{r,h,t}}. \quad (3)$$

Constraints

$$\sum_{h=1}^H \sum_{r=1}^R v_{r,r',h,t} \leq c_{r',t}, \quad \forall r', \quad (4)$$

$$v_{r,r',h,t} \geq 0 \text{ and integer valued}, \quad (5)$$

$$\sum_{r'=1}^R v_{r,r',h,t} \leq \hat{d}_{r,h,t}, \quad \forall r, \forall h. \quad (6)$$

Constraint (4) prevents the over-allocation of available resources. Constraint (5) ensures the integrality and non-negativity of the resource allocations and constraint (6) keeps the allocation to each customer below the corresponding estimated demand. If all of the estimated demand at time t can be met, any excess supply is held in inventory at the hub.

The objective function (3) captures our notion of fairness: minimizing the largest ratio of unmet demand over all customers. It could be modified to capture other notions of fairness. For example, one can penalize larger deviations between the demand and what is allocated by using a quadratic objective function, which would create a Mixed Integer Quadratic Program (MIQP). If the problem size becomes too large, one could also look at Linear Programming (LP) relaxations of the problem.

4 The CONCOR-1 trial: A case study for a proposed application of the resource allocation model

The Randomized, Open-Label Trial of CONvalescent Plasma for Hospitalized Adults With Acute COVID-19 Respiratory Illness (CONCOR-1) was an RCT involving 72 academic and community sites across Canada, the USA, and Brazil (Bégin et al., 2021). The randomization in this RCT was performed at a ratio of 2:1 allocation to receive CCP or standard of care for a planned study population of 1200 patients, stratified by age (< 60 and ≥ 60 years). The first CCP unit for the trial was collected on April 24, 2020, and the first patient was randomized on May 14, 2020. The trial ceased on January 29, 2021 with a total of 940 randomized patients. The objective of the trial was to assess whether transfusing CCP reduces deaths or the proportion of patients requiring intubation at day 30 compared to standard of care for hospitalized COVID-19 infected adult patients (Bégin et al., 2021). Canadian Blood Services (CBS) is responsible for collecting the CCP units and is the national blood supplier across all provinces in Canada except Québec. The supply and demand management of the CCP products for patients enrolled in the trial was conducted by CBS and the CONCOR-1 team at the McMaster Centre for Transfusion Research. In what follows, as we are evaluating our approach, we will take the viewpoint that the trial is in progress.

Canada's blood supply chain network is currently centralized and comprises two levels: regional CBS distribution sites and hospital blood banks. Nine CBS blood distribution sites are currently located across Canada, and each centre attempts to meet the CCP demand from the hospital blood banks in its network. We use all available data from the trial which comes from 18 hospital hubs, 30 hospital sites, and 8 CBS distribution sites.

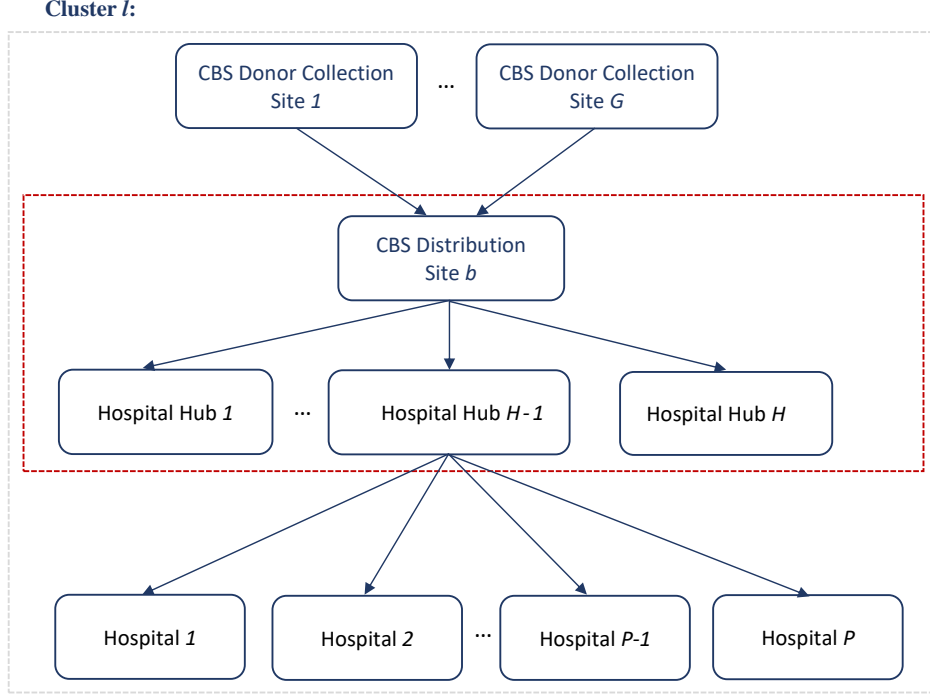
CCP is stored frozen and ideally must be transfused as soon as possible after being thawed, but at most within five days (Bloch et al., 2021). A patient who is randomized to the CCP arm in the trial requires a single dose of approximately 500 ml or two doses of 250 ml (from a single or two different donors) and the CCP unit is transfused to the patient within the first 24 hours after randomization (Bégin et al., 2021).

It is challenging for CBS to make decisions on CCP allocation, because (i) the CCP supply is limited and restricted by manufacturing policies and may not meet the total CCP demand, (ii) the trial involves hospitals from geographically dispersed locations and multiple blood distribution centres, (iii) the demand for CCP exhibits heterogeneity between different geographic regions, (iv) there is limited knowledge or historical data about the disease demographics making it difficult to forecast the supply and demand of CCP products, (v) there are specific clinical requirements for administering CCP transfusions, such as specific product dose, ABO blood group compatibility, and medical condition requirements, and (vi) the decisions must be made in real-time and once the CCP units are shipped to a hospital hub, redistribution to other hospital hubs is undesirable. Hence, the decision for every unit matters. The key observations of our work suggest that our proposed data-driven MIP model can help balance the supply and demand of CCP products and lead to a fair allocation of limited CCP products among hospital hubs.

The underlying network of our case study is shown in Figure 3. The corresponding definitions and notations are listed as follows:

- l : represents a cluster and is defined as a geographic region with only one CBS distribution site but one or more CBS donor collection sites, hospital hubs and individual hospital sites.
- g : represents an individual CBS donor collection site in cluster l , $g = 1, \dots, G$.
- b : represents the CBS distribution site in cluster l .
- h : represents an individual hospital hub under CBS distribution site b in cluster l , $h = 1, \dots, H$.
- p : represents an individual hospital site for hospital hub h in cluster l , $p = 1, \dots, P$.

The structure in Figure 3 is a hub-and-spoke structure (a CBS distribution site is a hub and its underlying hospital hubs are the spokes) and is consistent with our resource allocation MIP model proposed in Section 3.3. CBS distribution sites can centrally decide to reallocate the blood products if there is excess supply in a particular CBS region. Thus, we only consider a single CBS distribution site in our optimization problem. Based on the agreements between CBS and Héma-Québec, the blood supplier in Québec, less common blood groups (blood groups AB and B) can be shared (Bégin et al., 2021).

**Figure 3.** CCP allocation network

The CBS distribution site provides CCP to hospital hubs that in turn allocate units to other hospital sites in the area. We only focus on the allocations from the CBS distribution site to the hospital hubs.

As of May 2021, the distribution of ABO blood groups in Canada is O:46%, A:42%, B:9%, and AB:3% (Canadian Blood Services, 2021). In general, plasma for AB and O blood groups are universal donor and recipient, respectively. In practice, transfusing blood-specific units is prioritized. In the CONCOR-1 trial, due to the limited resources for B and AB plasma, patients with O and B blood groups can receive A and AB plasma, respectively, when an exact match is not available.

The dataset that we work with includes the CBS distribution site’s available CCP units after shipment at an aggregate level for each blood group on specific dates starting from September 1, 2020, up to January 25, 2021. It also contains data for the received CCP units from the CBS distribution site for each hospital hub, whether randomized to a patient or stored in inventory, from May 11, 2020, up to January 25, 2021. Furthermore, CCP-related information such as product dose, ABO group, and whether a unit was broken or leaking after being thawed for transfusion at a hospital hub are recorded. Table 1 provides a summary of the dataset’s attributes, their description, and their format.

We first create a cumulative weekly dataset of the number of new 500 ml units assigned to randomized patients at each hospital hub (considered as their CCP demand) and the total received units from CBS for each resource. We cumulatively sum hospital hubs’ weekly received CCP units and CBS weekly inventory difference to calculate CBS cumulative new weekly supply. To maintain and promote efficient CCP allocation, prior estimation of the hospital hub’s demand and CBS supply is necessary. One could either translate the CHIME model outputs to CCP supply or demand, or build a CHIME-like model that incorporates our variables. However, we found it more effective to work with the supply and demand directly and use PLR forecasting models. One reason is that the proportion of patients consenting to the trial also affects the total CCP demand. The consent rate can be affected by many factors such as a patient’s religion and other competing treatment strategies. So, the patient population may not be the same as the population considered in the CHIME or CHIME-like models. Furthermore, given the limited and highly-varied pattern of CCP supply and demand in our dataset, PLR models appear to be an effective tool for forecasting CCP units.

4.1 Model Assumptions

We have created an environment for weekly allocation of CCP units from the CBS distribution site to the hospital hubs participating in the CONCOR-1 trial. Our dataset contains CBS supply data from a later date than when the first patient was randomized to receive a CCP unit. So, we consider the week that CCP supply was first reported as the starting week

Table 1. Dataset description

Dataset	Attribute	Description	Format
CBS supply data	date	CBS distribution site’s aggregate available units after shipment on a date	Date
	A	Total number of blood group A CCP (250 ml) units	Integer
	AB	Total number of blood group AB CCP (250 ml) units	Integer
	B	Total number of blood group B CCP (250 ml) units	Integer
	O	Total number of blood group O CCP (250 ml) units	Integer
Hospital Hub data	Total	Total number of CCP (250 ml) units	Integer
	hospitalhub_ID	Unidentifiable unique ID of a hospital hub	String
	receiveddate	Hospital hubs received CCP units from the CBS distribution site on a date	Date
	DNL	The CCP unit’s de-identified ID	String
	productABOgroup	The ABO blood group of the CCP unit	String
	productdose	The CCP unit dose with 1 indicating 250 ml units and 2 indicating 500 ml units	Integer
	matched	Boolean variable with 1 indicating the CCP unit was matched with a randomized patient (indicating a demand), 0 if the unit was stored in the hospital hub’s inventory	Boolean
	thawed	Boolean variable with 1 indicating that the unit broke when thawing, 0 otherwise	Boolean

in our model (week of August 31, 2020). The last week in the dataset is the week of January 25, 2021, so the duration T of our model is 22 weeks. We combined hospital hubs that are very close in distance, under the same distribution network and with few COVID-19 patients. Small hospitals in distant areas with a very low number of hospitalized COVID-19 patients were removed from the study dataset.

Table 2 shows the actual supply and demand (500 ml units) reported for each resource and in total in this period. There are a total of 15 A, 14 O, 9.5 B, and 14 AB 500 ml CCP units received from Héma-Québec recorded in the dataset. The numbers in Table 2 are only associated with CBS supply (excluding any units received from Héma-Québec). We observe in the data that 4.5 A, 3 O, 2 B, and 0.5 AB 500 ml CCP units are unused at the hospital hubs because they were leaking or broken after being thawed for transfusion. These units are not included in our model as the wastage due to the thawed CCP units is negligible.

CCP Units (500 ml)	A	O	B	AB	Total
Supply	157	84.5	29	26	296.5
Demand	131.5	94.5	34.5	33.5	294

Table 2. CBS CCP supply and CCP demand (August 31, 2020 - January 25, 2021)

We assume that we have no information on the actual CCP supply and demand when deciding on how to allocate the CCP units. Thus, we perform weekly forecasts considering only the available data up to that week. The open source *scikit-learn-contrib py-earth* (Rudy, 2016) is used for applying the MARS model, where the maximum degree of interaction terms generated by the forward stage (the *max_degree* parameter) is set to 2 to better deal with the nonlinearities in the data, and the remaining parameters follow the default values. We are not performing any validations because our dataset is sparse and the performance of MARS cannot be significantly enhanced as using the cumulative sums has already smoothed the dataset. In general, a k -fold cross validation can be used with MARS to obtain a less biased estimate. We start making predictions four weeks from the start of available data to decrease the likelihood of severe underestimation or overestimation. Finally, a lag $l = 1$ is used in the autoregressive model of residual errors in (1). This choice is due to a week’s supply and demand tending to be closer to the amounts for the most recent week, as well as trying to avoid sudden changes that we might face when considering a larger lag.

We deal with the challenges that might occur in our forecasting process, as discussed in Section 3.2.3; if a negative non-cumulative forecast value on week $t + 1$ arises, we instead set the cumulative forecast value equal to the cumulative forecast value for week t . Furthermore, we account for expected sudden transitory changes in data, such as weeks containing holidays. We observe that the hospital hubs’ requested CCP units are relatively lower on weeks containing the Christmas, Boxing Day, and New Year’s Day statutory holidays. This issue is due to fewer working days or reduced staffing for these weeks. Thus, it is important to account for these weeks as they can be falsely detected as breakpoints. We have chosen the value of 0.7 as a reasonable adjustment factor and multiply the slope of forecast line going through the weeks containing Christmas, Boxing Day, and New Year’s Day holidays by this factor. One might need different adjustment factors depending on the extent the supply and demand are affected on a particular holiday.

From this point on, wherever we refer to MARS, we are considering the model's forecasts after accounting for the mentioned challenges and performing error analysis. Because CCP is stored in 250 ml units, we allocate 250 ml units from CBS to the hospital hubs in our model. The forecast values are decimal values; therefore, we round the supply and demand forecast values to the nearest integer so that all values correspond to 250 ml units. All the results reported in the figures and tables correspond to 500 ml units.

We are interested in making a fair allocation of different ABO blood group CCP units among the hospital hubs while minimizing their unmet CCP demand proportions. We consider four resources ($R = 4$) of A, O, B, and AB and examine two compatibility matrices C for assigning the CCP units in our MIP problem, as described in Section 3.3: (i) the identity compatibility matrix, and (ii) the ABO compatibility matrix used in the CONCOR-1 trial which allows the transfusion of A and AB plasma to patients with O and B blood groups, respectively, when the same blood group is not available (CONCOR-1 compatibility matrix). For our primary study, we use the forecast supply and demand for each of the resources in our model for each week. Furthermore, to evaluate different allocation scenarios, we also forecast the aggregate supply and demand for each week and consider the distribution of ABO blood groups in Canada as the probabilities for calculating the forecast CCP for each resource (Zietz, Zucker, & Tatonetti, 2020). For this sensitivity analysis, we use the following probabilities based on the distribution of Canadian blood groups on each run of our model for generating the supply and demand for each blood group: $w_{cA} = w_{hA} = 0.42$, $w_{cO} = w_{hO} = 0.46$, $w_{cB} = w_{hB} = 0.09$, and $w_{cAB} = w_{hAB} = 0.03$, where w_{c_r} and w_{h_r} are the supply and demand probabilities for blood group r , respectively. The goal of the sensitivity analysis is to determine the performance of the model in situations where it is only possible to forecast the aggregate amount of resources in terms of supply and demand. This analysis also helps gain insights as to the generalizability of the results of our primary study.

We assume that both CBS and hospital hubs can store the excess units at the end of each week to use them in later weeks. We solve our MIP model based on forecast supply and demand for each resource, and the actual inventories held at CBS and hospital hubs on the week under study. We assume that the hospital hubs can use a compatible CCP unit according to the CONCOR-1 compatibility matrix (if available) for a patient when the same blood group CCP is not available. We examine both identity and CONCOR-1 compatibility matrices for solving our MIP model; however, at the inventory level of the hospital hubs, only the CONCOR-1 compatibility matrix is used. We shall see later that this is a good combination for the situations where forecast supply and demand are used, and the forecasting errors are not too large. The identity compatibility matrix at the MIP level prevents issues in terms of greedy allocation of scarce resources, and the CONCOR-1 compatibility matrix at the inventory level of the hospital hubs allows for the efficient compensation of forecasting errors.

Because we know the actual demand for a particular week only after that week has passed and because the supply is limited, we might not fully meet all demands. In these cases, the unmet demand is carried over to the next week. We note that the optimization model ensures that the CCP units of a resource allocated to a hospital hub are never more than its forecast demand. Excess units stored in hospital hubs' inventories are due to the possible difference between the actual and forecast values.

A total of m runs are used to calculate the mean final unmet demand (\bar{u}_T) and the mean ratio of final unmet demand to total demand (\bar{z}_T) (for ease of reading, we will subsequently refer to these quantities as unmet demand and unmet demand ratio, respectively) for each hospital hub:

$$\bar{u}_T = \frac{\sum_{r=1}^R u_{T_r}}{m}$$

where the duration of our model is $T = 22$ weeks, and u_{T_r} is the hospital hub's final unmet demand for resource r and,


$$\bar{z}_T = \frac{\sum_{r=1}^R u_{T_r}}{m \sum_{t=1}^T \sum_{r=1}^R d_{t_r}}$$

where d_{t_r} is the hospital hub's actual newly-added demand on week t for resource r .


4.2 Results

In Figure 4, we show the (real-time) cumulative weekly supply and demand forecasts over time, respectively, in terms of total, A, O, B, and AB CCP units after using MARS for the dataset considering only the available data up to that week. The real-time demand forecasting for each week in Figure 4 uses the cumulative aggregate demand data over all hospital hubs up to that week. We observe that MARS performs well even with our limited dataset. Table 3 shows Root Mean Square Error (RMSE) and Mean Absolute Percentage Error (MAPE) of our supply and demand forecasts under MARS for total, A, O, B, and AB CCP units; lower values are better. We observe that MARS fits the data well after performing error analysis and accounting for the mentioned challenges discussed in Section 3.2.2 and Section 3.2.3. We note that we are making forecasts from week 4, which explains the severe supply overestimation for week 4 in Figure


4e where the data is sparse. In such situations, the cumulative value remains the same until a subsequent cumulative forecast is higher.




(a) Total units




(b) A units



(c) O units



(d) B units



(e) AB units

Figure 4. Model performance in forecasting CCP supply and demand

All the experiments are run on Spyder Python IDE on a computer with an Intel(R) Core(TM) i7-4510U processor running at 2.60 GHz using 94 MB of RAM, running Windows 10 version 21H2. The weekly supply and demand forecasting took on average approximately 0.04 and 0.28 seconds, respectively. Our MIP model has 43 variables and 81 constraints, and solving the weekly MIP model took on average approximately 0.02 seconds. Table 4 reports our CCP allocation model's performance in terms of unmet demand and unmet demand ratio for our primary study, i.e.,

Hub	RMSE					MAPE				
	Total	A	O	B	AB	Total (%)	A (%)	O (%)	B (%)	AB (%)
Hospital Hub 1	1.06	0.43	0.84	0.43	0.41	16.07	17.05	19.73	20.45	22.73
Hospital Hub 2	1.74	1.21	0.91	0.84	0.58	12.87	18.79	15.50	17.82	17.86
Hospital Hub 3	2.97	1.88	1.56	0.81	0.87	11.32	19.41	12.59	19.89	18.67
Hospital Hub 4	2.33	1.59	0.95	0.61	0.49	21.70	25.45	53.45	16.75	16.67
Hospital Hub 5	1.76	1.10	0.80	1.02	0.46	7.31	15.26	11.52	22.35	17.36
Hospital Hub 6	3.00	1.38	1.73	1.06	0.66	15.12	13.86	14.60	22.66	14.94
Hospital Hub 7	2.64	2.01	1.04	-	-	27.91	25.58	27.12	-	-
All Hospital Hubs	7.01	4.27	1.99	1.57	1.57	3.67	6.45	3.78	18.36	14.46
CBS Distribution Site	10.68	5.25	5.41	1.78	3.35	6.12	6.01	9.90	7.59	18.00

Table 3. RMSE and MAPE of supply and demand forecasts under MARS

allocating the resources based on the forecast supply and demand for each blood group. In Table 5, we compare the results to both when no forecasting is required, i.e., the actual values of supply and demand for each resource are known for each week and the actual allocations in the CONCOR-1 trial. In both tables, two different compatibility matrices (identity and CONCOR-1) are chosen for the MIP allocation model. Finally, in Table 6, we analyze the sensitivity of our model to different allocation settings by forecasting the aggregate supply and demand and using the distribution of Canadian blood groups to calculate the supply and demand for each resource. A total of $m = 300$ runs are considered for calculating unmet demand and unmet demand ratio and the standard error (SE) of their corresponding 95% confidence intervals is reported.

Hub	Total Demand	Total Forecast Demand	Identity		CONCOR-1	
			\bar{u}_T	\bar{z}_T (%)	\bar{u}_T	\bar{z}_T (%)
Hospital Hub 1	12.5	12	1.00	8.00	1.00	8.00
Hospital Hub 2	29	30.5	0.50	1.72	1.00	3.45
Hospital Hub 3	74	79	4.50	6.08	10.00	13.51
Hospital Hub 4	26	30.5	1.00	3.85	1.50	5.77
Hospital Hub 5	30	30	1.50	5.00	1.00	3.33
Hospital Hub 6	105	110.5	8.50	8.10	10.00	9.52
Hospital Hub 7	17.5	20.5	0	0	0	0

Table 4. Resource allocation results under forecast supply and demand for each resource

Hub	Total Demand	MIP Allocations					
		Identity		CONCOR-1		Allocations in CONCOR-1	
		\bar{u}_T	\bar{z}_T (%)	\bar{u}_T	\bar{z}_T (%)	\bar{u}_T	\bar{z}_T (%)
Hospital Hub 1	12.5	1.50	12.00	1.00	8.00	3.00	24.00
Hospital Hub 2	29	1.00	3.45	0.50	1.72	1.50	5.17
Hospital Hub 3	74	7.00	9.46	4.00	5.41	20.00	27.02
Hospital Hub 4	26	2.00	7.69	0.50	1.92	4.50	17.31
Hospital Hub 5	30	1.50	5.00	0.50	1.67	5.00	16.67
Hospital Hub 6	105	10.00	9.52	7.00	6.67	27.00	25.71
Hospital Hub 7	17.5	0	0	0	0	1.00	5.71

Table 5. MIP model allocations under actual supply and demand versus actual allocations in CONCOR-1

The key observations obtained from using our data-driven resource allocation model in the CONCOR-1 case study are as follows:

1. The role of hospital hubs' proportion of CCP demand — *A hospital hub's proportion of CCP demand for a resource can lead to its demand not being fully met if the available supply of the resource on a particular week is limited compared to its total demand for the same week. Our approach can help minimize the unmet CCP demand ratios and lead to balanced and fair CCP allocation decisions.*

We note that there is a shortage of 10 O, 5.5 B, and 7.5 AB CCP units due to the limitation in supply of these blood groups, as reported in Table 2. Thus, the CCP demand proportion of a hospital hub for a particular resource and the variance in the total demand between the hospital hubs in each week can affect the unmet demand and unmet demand ratio. For instance, if the available supply on week t can meet all the demand on week t , a hospital hub's proportion of CCP demand for a resource is not an issue. On the other hand, if the available supply on week t is limited compared to

Hub	Total Demand	Total Forecast Demand	Identity		CONCOR-1	
			$\bar{u}_T \pm SE$	$\bar{z}_T \pm SE$ (%)	$\bar{u}_T \pm SE$	$\bar{z}_T \pm SE$ (%)
Hospital Hub 1	12.5	10.5	1.00 ± 0.03	8.04 ± 0.26	0.93 ± 0.03	7.47 ± 0.27
Hospital Hub 2	29	30	1.62 ± 0.10	5.57 ± 0.36	1.65 ± 0.10	5.70 ± 0.34
Hospital Hub 3	74	77	3.66 ± 0.15	4.94 ± 0.20	5.97 ± 0.20	8.07 ± 0.27
Hospital Hub 4	26	28.5	1.29 ± 0.08	4.97 ± 0.31	1.52 ± 0.08	5.85 ± 0.30
Hospital Hub 5	30	29.5	1.26 ± 0.06	4.20 ± 0.19	1.21 ± 0.05	4.03 ± 0.16
Hospital Hub 6	105	110	8.57 ± 0.17	8.16 ± 0.16	7.63 ± 0.19	7.27 ± 0.18
Hospital Hub 7	17.5	20.5	0.02 ± 0.01	0.14 ± 0.07	0.42 ± 0.03	2.38 ± 0.15

Table 6. Resource allocation results under aggregate forecast supply and demand (Sensitivity Analysis)

the total demand for the same week, a hospital hub's high demand proportion for a limited resource CCP can lead to its demand not being fully met. This issue is unavoidable in our setting where we make predictions and allocations on a weekly basis and the optimization model finds a solution based only on the current situation.

We observe in Table 4 that Hospital Hub 1, Hospital Hub 3, and Hospital Hub 6 have larger unmet demand ratios under both compatibility matrices. Hospital Hub 1's unmet demand ratio is high because of its unmet CCP demand proportions. However, its unmet demand is not high and is due to the demand that was not met on the last week in our model (and may have been satisfied if we continued for additional weeks). In fact, the supply and demand for the last week can highly affect the final unmet demands of all hospital hubs. Hospital Hub 3 and Hospital Hub 6 require a higher proportion of B and AB CCP units compared to the other hospital hubs on weeks when the supply of these units is limited. A hospital hub's unmet demand on a particular week is moved to the next week and will have an effect on its future allocations. We observe that the rest of the hospital hubs have similar unmet demand and unmet demand ratio values. This suggests that our proposed data-driven MIP model leads to a reasonable and fair balance of limited CCP products between the hospital hubs under the MARS forecasting model.

2. The role of the compatibility matrix

2.1. *When the actual supply and demand is unbalanced, not known before allocation, and hence error due to forecasting is present, using the identity compatibility matrix in the MIP level is preferred as it prevents the allocation of limited resources to demand for a more abundant resource.*

We notice in Table 4 that the total unmet demand for the identity compatibility matrix is lower than for the CONCOR-1 compatibility matrix. While the CONCOR-1 compatibility matrix satisfies as much demand as possible in the current week, its use in this real-time setting can cause issues for future demand, in particular by allocating resources with lower long term supply to demands that can (eventually) be satisfied by more abundant resources. We found that under the CONCOR-1 compatibility matrix, the shortage of O units was compensated by excess A units, and excess B demands were met by AB units. Because the supply of A units was the highest during the trial, and the supply of AB units was the lowest, using AB units to meet demand from other blood groups led to shortages of AB units in future weeks. Furthermore, because the supply and demand forecasts are used instead of the actual values, we might overestimate the demand for a rare resource, such as the AB blood group CCP that is also compatible with another resource, so its allocation increases the final unmet demand of some hospital hubs. However, under the identity compatibility matrix, greedy allocations at the MIP level are prevented. If the forecasting error is not too large, the combination of the identity compatibility matrix at the MIP level and the CONCOR-1 compatibility matrix at the inventory level of hospital hubs is preferred under forecast supply and demand as the identity compatibility matrix at the MIP level prevents greedy allocations that cause issues for future demand, but the forecasting error results in the allocation of some compatible resources that combined with the CONCOR-1 compatibility matrix at the hospital hubs' inventory level can help meet more demand. Therefore, for our primary study, we recommend this combination of compatibility matrices that prevents future misallocation and absorbs the forecasting errors in an effective manner.

2.2. *The CONCOR-1 compatibility matrix results in lower unmet demand when the actual supply and demand for each week is known. The total unmet demand under this compatibility matrix is close to the lowest achievable value, the actual shortage in supply.*

While using the identity compatibility matrix in Table 4 where supply and demand forecasts are used exhibits better results, the opposite is true when solving the MIP model based on the assumption of knowing the actual supply and demand for each week. This can be observed by comparing the total unmet demand in Table 5 under "MIP Allocations" for both compatibility matrices (23 versus 13.5). When the actual supply and demand are used, over-allocation of limited resources will not happen and the final unmet demand is close to the actual supply shortage. We note again that demand for 5.5 B and 7.5 AB CCP units cannot be met due to supply shortage which is almost exactly the total unmet demand result of our MIP model under the CONCOR-1 compatibility matrix (13.5). This observation also reinforces the notion in the previous observation that forecasting errors (combined with unbalanced supply and demand)

are what drive the recommendation of the identity compatibility matrix under forecast supply and demand. The identity compatibility matrix does not allow cross-transfusion, so 10 O units cannot be satisfied by the excess A units due to supply shortage. The total unmet demand under the identity compatibility matrix (23) further validates the performance of our model in efficiently allocating the available resources as it is equal to the actual shortage in supply, as shown in Table 2.

2.3. The identity compatibility matrix results in lower unmet demand under forecast supply and demand compared to when it is used under actual supply and demand, as reasonable forecast error can better help in meeting the demand.

Comparing the unmet demand ratio values in Table 4 and Table 5 under the CONCOR-1 compatibility matrix for each hospital hub, we observe that the unmet demand ratio for all hospital hubs is improved (or is the same) in Table 5, significantly so for Hospital Hub 3 and Hospital Hub 6. Furthermore, the unmet demand ratio values in Table 5 are closer to each other under the CONCOR-1 compatibility matrix than those reported in Table 4, where the results are affected by the forecasting errors. However, we notice that all unmet demand ratio values in Table 4 under the identity compatibility matrix are lower than (or the same as) those in Table 5. The reason is that using the identity compatibility matrix at the MIP level under actual supply and demand is not a good choice considering that the MIP model never assigns more CCP units than the actual demand. This choice leads to the CONCOR-1 compatibility matrix at the inventory level of the hospital hubs not helping at all in meeting the demand. However, under forecast supply and demand, although the identity compatibility matrix prevents the allocation of a resource to a demand for another resource, the forecasts induce some compatible assignments. Hence, when the identity compatibility matrix at the MIP level is combined with the CONCOR-1 compatibility matrix at the inventory level of the hospital hubs, and the supply and demand forecasts do not have large errors, there is more chance to meet the demand, as what is observed for Hospital Hub 3 and Hospital Hub 6.

2.4. The MIP model's results after using the aggregate forecast supply and demand (sensitivity analysis) are close to the primary study when the identity compatibility matrix is used, and improved under the CONCOR-1 compatibility matrix. However, the identity compatibility matrix is a better choice than the CONCOR-1 compatibility matrix when using the aggregate forecast supply and demand and in the presence of forecast errors, consistent with what is observed in the primary study.

We notice that although the results under the identity compatibility matrix remain almost unchanged, the total unmet demand and unmet demand ratio values under the CONCOR-1 compatibility matrix in Table 6 are lower than those in Table 4. This appears to arise due to the fact that including randomness for generating the supply and demand for each blood group based on the aggregate supply and demand may slightly lower the forecasting error. Furthermore, performing sufficient runs helps to smooth the effect of scenarios where resources with limited supply, such as the AB blood group CCP, are allocated in a greedy manner and cause issues for future demand. We note that the results under the identity compatibility matrix are better than those under the CONCOR-1 compatibility matrix in Table 6 as the identity compatibility matrix prevents any cross-allocation of units and lowers the effect of forecast errors, consistent with what we observed in Table 4. Comparing the unmet demand and unmet demand ratio values under the CONCOR-1 compatibility matrix in Table 5 and Table 6 under "MIP Allocations", it is important to note that while having exact knowledge of supply and demand leads to a fairer allocation, the degradation in performance is not unreasonable if we use the forecast values instead, further supporting the efficacy of our approach. However, as previously discussed, using the identity compatibility matrix in the MIP model under actual supply and demand leads to more unmet demand (as compared to using the CONCOR-1 compatibility matrix) due to not allocating excess units of compatible resources, which prevents the CONCOR-1 compatibility matrix at the inventory level of the hospital hubs to help meet the demand. Similar to our primary study, if the forecast errors are not too large, using the identity compatibility matrix at the MIP level under forecast supply and demand results in making some compatible assignments that can better help meet the demand when combined with the CONCOR-1 compatibility matrix at the hospital hubs' inventory level.

For this case study and based on all the above observations, we conclude that our primary study when the identity compatibility matrix is used at the MIP level is the most promising choice for our real-time setting. The reason is that as long as the forecast errors for the individual blood groups are not too large, using the combination of the identity compatibility matrix at the MIP level and the CONCOR-1 compatibility matrix at the inventory level helps in both preventing greedy allocations and compensating for forecasting errors. Furthermore, no randomness due to supply and demand probabilities is included in this choice for calculating the supply and demand for each blood group (as what is assumed in our sensitivity analysis), and the assignments of the model are not affected by any ABO blood distributions.

3. The role of the MIP model — *The results obtained from our MIP model both under actual and forecast supply are preferable to what was used in practice.*

We compare the unmet demand and unmet demand ratio for both compatibility matrices using the MIP model allocations under actual supply and demand, and the actual allocations in the CONCOR-1 trial. We observe that the results under

"MIP Allocations" are more fair than those under "Allocations in CONCOR-1" (the unmet demand ratio values are closer to each other). Furthermore, the total unmet demand under "MIP Allocations" is notably lower (23 and 13.5 versus 62) which further supports the use of our optimization model, as it accounts for the hospital hubs' unmet demand and inventories.

5 Conclusion and Future Work

Decision-makers often face challenges in terms of (i) allocating limited resources, such as vaccines, blood products, and medical equipment, and (ii) forecasting the supply and demand for these resources during epidemics as there is limited knowledge and historical data about disease demographics. This work has considered the problem of real-time short-term supply and demand forecasting and fair allocation of limited resources during epidemics by using PLR forecasting models and introducing a data-driven MIP resource allocation model. We have studied the application of our proposed MIP model in a CCP clinical trial case study with the objective of minimizing each hospital hub's unmet ratio of CCP demand. We showed that as long as a hospital hub does not have a high demand proportion for a limited blood group CCP on a particular week (in which case a fair allocation is not possible), our MIP model leads to a balanced and fair final unmet ratio of CCP demand between the hospital hubs. We also showed that allocating a compatible resource to satisfy the demand for a resource helps in situations where the actual supply and demand is known, but might be problematic when we are forecasting the supply and demand and the actual supply of the compatible resource is limited. It would be interesting to investigate the range of the forecasting errors within which a particular compatibility matrix is preferred.

Examining PLR forecasting models on larger datasets and comparing their performance with more advanced machine learning and time-series forecasting models could be investigated in future work. We have addressed several challenges that arise when dealing with sparse data in a real-time setting. We are interested in evaluating other forecasting models that can adapt to these challenges while yielding reasonable forecasts. It would be worthwhile to investigate our MIP model's performance when other supply and demand forecasting models are used.

Finally, multiple objective functions and different notions of fairness can be considered in a single resource allocation problem. Using more than one objective function and focusing on other notions of fairness such as minimizing the aggregate unmet demand over all hospital hubs or minimizing the transportation costs of shipping CCP units with respect to the location of the hospital hubs are examples of problems of interest. It would also be interesting to see our MIP model's performance when applied to other allocation settings with limited supply. Our MIP model only considers a centralized supplier, thus we leave the model's evaluation with the presence of multiple suppliers to future investigation.

Acknowledgments

The authors would like to thank Julie Carruthers, Erin Jamula, and Melanie St John at the McMaster Centre for Transfusion Research for their administrative support.

Data Availability Statement

Due to the ethically, legally, and commercially sensitive nature of the research, participants of this study did not agree for their data to be shared publicly, so supporting data is not available.

Funding

The authors acknowledge the support and funding provided by a Mitacs Research Training Award (Award IT22358), the McMaster Centre for Transfusion Research, and the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant Program (RGPIN-2016-04518 and RGPIN-2022-02999).

Disclosure statement

No potential competing interest was reported by the authors.

References

- Anparasan, A., & Lejeune, M. (2019). Resource deployment and donation allocation for epidemic outbreaks. *Annals of Operations Research*, 283(1), 9–32.
- Bégin, P., Callum, J., Heddle, N., Cook, R., Zeller, M. P., Tinmouth, A., ... Arnold, D. M. (2021). Convalescent plasma for adults with acute COVID-19 respiratory illness (CONCOR-1): Study protocol for an international, multicenter, randomized, open-label trial. *Trials*, 22(323).
- Bekker, R., uit het Broek, M., & Koole, G. (2023). Modeling COVID-19 hospital admissions and occupancy in the Netherlands. *European Journal of Operational Research*, 304(1), 207–218.
- Bloch, E. M., Goel, R., Wendel, S., Burnouf, T., Al-Riyami, A. Z., Ang, A. L., ... So-Osman, C. (2021). Guidance for the procurement of COVID-19 convalescent plasma: Differences between high- and low-middle-income countries. *Vox Sanguinis*, 116(1), 18–35.
- Box, G. E., & Tiao, G. C. (1975). Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association*, 70(349), 70–79.
- Brauer, F. (2008). Compartmental models in epidemiology. In *Mathematical Epidemiology* (pp. 19–79). Springer.
- Butte, N. F., Wong, W. W., Adolph, A. L., Puyau, M. R., Vohra, F. A., & Zakeri, I. F. (2010). Validation of cross-sectional time series and multivariate adaptive regression splines models for the prediction of energy expenditure in children and adolescents using doubly labeled water. *The Journal of Nutrition*, 140(8), 1516–1523.
- Canadian Blood Services. (2021). *What's my blood type?* <http://www.blood.ca/en/blood/donating-blood/whats-my-blood-type>. ([Online; accessed 20-October-2022])
- Caunhye, A. M., Nie, X., & Pokharel, S. (2012). Optimization models in emergency logistics: A literature review. *Socio-economic Planning Sciences*, 46(1), 4–13.
- Chen, L., Xiong, J., Bao, L., & Shi, Y. (2020). Convalescent plasma as a potential therapy for COVID-19. *The Lancet Infectious Diseases*, 20(4), 398–400.
- Chou, S.-M., Lee, T.-S., Shao, Y. E., & Chen, I.-F. (2004). Mining the breast cancer pattern using artificial neural networks and multivariate adaptive regression splines. *Expert Systems with Applications*, 27(1), 133–142.
- Chowdhury, M. E., Rahman, T., Khandakar, A., Al-Madeed, S., Zughaier, S. M., Doi, S. A. R., ... Islam, M. T. (2021). An early warning tool for predicting mortality risk of COVID-19 patients using machine learning. *Cognitive Computation*, 1–16.
- Dasaklis, T. K., Pappis, C. P., & Rachaniotis, N. P. (2012). Epidemics control and logistics operations: A review. *International Journal of Production Economics*, 139(2), 393–410.
- Du, M., Sai, A., & Kong, N. (2021). A data-driven optimization approach for multi-period resource allocation in cholera outbreak control. *European Journal of Operational Research*, 291(3), 1106–1116.
- Ellaway, P. (1978). Cumulative sum technique and its application to the analysis of peristimulus time histograms. *Electroencephalography and Clinical Neurophysiology*, 45(2), 302–304.
- Fernandez-Sojo, J., Delgadillo, J., Vives, J., Rodriguez, L., Mendoza, A., Azqueta, C., ... Querol, S. (2022). A hub-and-spoke model to deliver effective access to chimeric antigen receptor T-cell therapy in a public health network: The Catalan Blood and Tissue Bank experience. *Cytotherapy*.
- Ferstad, J. O., Gu, A. J., Lee, R. Y., Thapa, I., Shin, A. Y., Salomon, J. A., ... Scheinker, D. (2020). A model to forecast regional demand for COVID-19 related hospital beds. *medRxiv*.
- Friedman, J. H. (1991). Multivariate adaptive regression splines. *The Annals of Statistics*, 1–67.
- Gharbharan, A., Jordans, C. C., Geurtsvankessel, C., den Hollander, J. G., Karim, F., Mollema, F. P. N., ... Rijnders, B. J. (2020). Convalescent plasma for COVID-19. A randomized clinical trial. *MEDRxiv*.
- Gong, J., Ou, J., Qiu, X., Jie, Y., Chen, Y., Yuan, L., ... Hu, B. (2020). A tool for early prediction of severe coronavirus disease 2019 (COVID-19): A multicenter study using the risk nomogram in Wuhan and Guangdong, China. *Clinical Infectious Diseases*, 71(15), 833–840.
- Han, S., Preciado, V. M., Nowzari, C., & Pappas, G. J. (2015). Data-driven network resource allocation for controlling spreading processes. *IEEE Transactions on Network Science and Engineering*, 2(4), 127–138.
- Heitmann, H., & Brüggemann, W. (2014). Preference-based assignment of university students to multiple teaching groups. *OR Spectrum*, 36(3), 607–629.
- Hung, I. F., To, K. K., Lee, C.-K., Lee, K.-L., Chan, K., Yan, W.-W., ... Yuen, K.-Y. (2011). Convalescent plasma treatment reduced mortality in patients with severe pandemic influenza A (H1N1) 2009 virus infection. *Clinical Infectious Diseases*, 52(4), 447–456.
- Jarrett, J. E., & Khumuwalla, S. B. (1987). A study of forecast error and covariant time series to improve forecasting for

- financial decision making. *Managerial Finance*, 13(2), 20–24.
- Karsu, Ö., & Erkan, H. (2020). Balance in resource allocation problems: A changing reference approach. *OR Spectrum*, 42(1), 297–326.
- Karsu, Ö., & Morton, A. (2014). Incorporating balance concerns in resource allocation decisions: A bi-criteria modelling approach. *Omega*, 44, 70–82.
- Karsu, Ö., & Morton, A. (2015). Inequity averse optimization in operational research. *European Journal of Operational Research*, 245(2), 343–359.
- Katris, C. (2021). A time series-based statistical approach for outbreak spread forecasting: Application of COVID-19 in Greece. *Expert Systems with Applications*, 166, 114077.
- Kraft, C. S., Hewlett, A. L., Koepsell, S., Winkler, A. M., Kratochvil, C. J., Larson, L., ... Ribner, B. S. (2015). The use of TKM-100802 and convalescent plasma in 2 patients with Ebola virus disease in the United States. *Clinical Infectious Diseases*, 61(4), 496–502.
- Kumar, A., & Kleinberg, J. (2000). Fairness measures for resource allocation. In *Proceedings 41st Annual Symposium on Foundations of Computer Science* (pp. 75–85).
- Lasheras, J. E. S., Donquiles, C. G., Nieto, P. J. G., Moleon, J. J. J., Salas, D., Gómez, S. L. S., ... de Cos Juez, F. J. (2020). A methodology for detecting relevant single nucleotide polymorphism in prostate cancer with multivariate adaptive regression splines and backpropagation artificial neural networks. *Neural Computing and Applications*, 32(5), 1231–1238.
- Li, L., Zhang, W., Hu, Y., Tong, X., Zheng, S., Yang, J., ... Liu, Z. (2020). Effect of convalescent plasma therapy on time to clinical improvement in patients with severe and life-threatening COVID-19: a randomized clinical trial. *Jama*, 324(5), 460–470.
- Li, N., Chiang, F., Down, D. G., & Heddle, N. M. (2021). A decision integration strategy for short-term demand forecasting and ordering for red blood cell components. *Operations Research for Health Care*, 100290.
- Li, N., Zeller, M. P., Shih, A. W., Heddle, N. M., St. John, M., Bégin, P., ... Tinmouth, A. (2022). A data-informed system to manage scarce blood product allocation in a randomized controlled trial of convalescent plasma. *Transfusion*. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1111/trf.17151>
- Liu, M., & Liang, J. (2013). Dynamic optimization model for allocating medical resources in epidemic controlling. *Journal of Industrial Engineering and Management (JIEM)*, 6(1), 73–88.
- López-Lozano, J.-M., Lawes, T., Nebot, C., Beyaert, A., Bertrand, X., Hocquet, D., ... study group, T. (2019). A nonlinear time-series analysis approach to identify thresholds in associations between population antibiotic use and rates of resistance. *Nature Microbiology*, 4(7), 1160–1172.
- Lu, M., & Chen, Y. (2020). Improved estimation and forecasting through residual-based model error quantification. *SPE Journal*, 25(02), 951–968.
- Mestre, A. M., Oliveira, M. D., & Barbosa-Póvoa, A. (2012). Organizing hospitals into networks: A hierarchical and multiservice model to define location, supply and referrals in planned hospital systems. *OR Spectrum*, 34(2), 319–348.
- Nikolopoulos, K., Punia, S., Schäfers, A., Tsinopoulos, C., & Vasilakis, C. (2021). Forecasting and planning during a pandemic: COVID-19 growth rates, supply chain disruptions, and governmental decisions. *European Journal of Operational Research*, 290(1), 99–115.
- Piantadosi, S. (2017). *Clinical trials: A methodologic perspective*. John Wiley & Sons.
- Preciado, V. M., Zargham, M., Enyioha, C., Jadbabaie, A., & Pappas, G. (2013). Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks. In *52nd IEEE Conference on Decision and Control* (pp. 7486–7491).
- Preciado, V. M., Zargham, M., Enyioha, C., Jadbabaie, A., & Pappas, G. J. (2014). Optimal resource allocation for network protection against spreading processes. *IEEE Transactions on Control of Network Systems*, 1(1), 99–108.
- Psaraftis, H. N., Tharakan, G. G., & Ceder, A. (1986). Optimal response to oil spills: The strategic decision case. *Operations Research*, 34(2), 203–217.
- Rachaniotis, N., Dasaklis, T. K., & Pappis, C. (2017). Controlling infectious disease outbreaks: A deterministic allocation-scheduling model with multiple discrete resources. *Journal of Systems Science and Systems Engineering*, 26(2), 219–239.
- Rachaniotis, N. P., Dasaklis, T. K., & Pappis, C. P. (2012). A deterministic resource scheduling model in epidemic control: A case study. *European Journal of Operational Research*, 216(1), 225–231.
- Roberts, H. W., Wagh, V. K., Mullens, I. J., Borsci, S., Ni, M. Z., & O'Brart, D. P. (2018). Evaluation of a hub-and-spoke model for the delivery of femtosecond laser-assisted cataract surgery within the context of a large randomised controlled trial. *British Journal of Ophthalmology*, 102(11), 1556–1563.

- Rosen, J. B., & Pardalos, P. M. (1986). Global minimization of large-scale constrained concave quadratic problems by separable programming. *Mathematical Programming*, 34(2), 163–174.
- Rudy, J. (2016). *py-earth: A Python implementation of Jerome Friedman's multivariate adaptive regression splines*. <http://www.github.com/scikit-learn-contrib/py-earth>. ([Online; accessed 20-October-2022])
- Senthilkumar, D., & Paulraj, S. (2013). Diabetes disease diagnosis using multivariate adaptive regression splines. *AGE*, 768, 52.
- Serrano, N. B., Sánchez, A. S., Lasheras, F. S., Iglesias-Rodríguez, F. J., & Valverde, G. F. (2020). Identification of gender differences in the factors influencing shoulders, neck and upper limb MSD by means of multivariate adaptive regression splines (MARS). *Applied Ergonomics*, 82, 102981.
- Silal, S. P., Little, F., Barnes, K. I., & White, L. J. (2016). Sensitivity to model structure: A comparison of compartmental models in epidemiology. *Health Systems*, 5(3), 178–191.
- Simonovich, V. A., Burgos Pratx, L. D., Scibona, P., Beruto, M. V., Vallone, M. G., Vázquez, C., ... Belloso, W. H. (2021). A randomized trial of convalescent plasma in Covid-19 severe pneumonia. *New England Journal of Medicine*, 384(7), 619–629.
- Smith, H. K., Harper, P. R., & Potts, C. N. (2013). Bicriteria efficiency/equity hierarchical location models for public service application. *Journal of the Operational Research Society*, 64(4), 500–512.
- Srinivasa, A. V., & Wilhelm, W. E. (1997). A procedure for optimizing tactical response in oil spill clean up operations. *European Journal of Operational Research*, 102(3), 554–574.
- Stawicki, S. P., Jeanmonod, R., Miller, A. C., Paladino, L., Gaieski, D. F., Yaffee, A. Q., ... Garg, M. (2020). The 2019–2020 novel coronavirus (severe acute respiratory syndrome coronavirus 2) pandemic: A joint American college of academic international medicine-world academic council of emergency medicine multidisciplinary COVID-19 working group consensus paper. *Journal of Global Infectious Diseases*, 12(2), 47.
- Strikholm, B. (2006). *Determining the number of breaks in a piecewise linear regression model* (SSE/EFI Working Paper Series in Economics and Finance No. 648). Stockholm School of Economics. Retrieved from <https://EconPapers.repec.org/RePEc:hhs:hastef:0648>
- Suits, D. B., Mason, A., & Chan, L. (1978). Spline functions fitted by standard regression methods. *The Review of Economics and Statistics*, 132–139.
- Sun, L., DePuy, G. W., & Evans, G. W. (2014). Multi-objective optimization models for patient allocation during a pandemic influenza outbreak. *Computers & Operations Research*, 51, 350–359.
- Tomar, A., & Gupta, N. (2020). Prediction for the spread of COVID-19 in India and effectiveness of preventive measures. *Science of the Total Environment*, 728, 138762.
- Van Griensven, J., Edwards, T., de Lamballerie, X., Semple, M. G., Gallian, P., Baize, S., ... Haba, N. Y. (2016). Evaluation of convalescent plasma for Ebola virus disease in Guinea. *New England Journal of Medicine*, 374(1), 33–42.
- Velavan, T. P., & Meyer, C. G. (2020). The COVID-19 epidemic. *Tropical Medicine & International Health*, 25(3), 278.
- Weissman, G. E., Crane-Droesch, A., Chivers, C., Luong, T., Hanish, A., Levy, M. Z., ... Halpern, S. D. (2020). Locally informed simulation to predict hospital capacity needs during the COVID-19 pandemic. *Annals of Internal Medicine*, 173(1), 21–28.
- Wilhelm, W. E., & Srinivasa, A. V. (1997). Prescribing tactical response for oil spill clean up operations. *Management Science*, 43(3), 386–402.
- Yang, L., Liu, S., Tsoka, S., & Papageorgiou, L. G. (2016). Mathematical programming for piecewise linear regression analysis. *Expert Systems with Applications*, 44, 156–167.
- Yao, D., Yang, J., & Zhan, X. (2013). A novel method for disease prediction: Hybrid of random forest and multivariate adaptive regression splines. *Journal of Computers*, 8(1), 170–177.
- Yarmand, H., Ivy, J. S., Denton, B., & Lloyd, A. L. (2014). Optimal two-phase vaccine allocation to geographically different regions under uncertainty. *European Journal of Operational Research*, 233(1), 208–219.
- Yin, X., & Büyüktaktakın, İ. E. (2021). A multi-stage stochastic programming approach to epidemic resource allocation with equity considerations. *Health Care Management Science*, 24(3), 597–622.
- Zhou, B., Zhong, N., & Guan, Y. (2007). Treatment with convalescent plasma for influenza A (H5N1) infection. *New England Journal of Medicine*, 357(14), 1450–1451.
- Zietz, M., Zucker, J., & Tatonetti, N. P. (2020). Associations between blood type and COVID-19 infection, intubation, and death. *Nature Communications*, 11(1), 1–6.