DEMAND FORECASTING AND INVENTORY MANAGEMENT OF PERISHABLE INVENTORY - WITH A FOCUS ON BLOOD PLATELET TRANSFUSIONS

DEMAND FORECASTING AND INVENTORY MANAGEMENT OF PERISHABLE INVENTORY - WITH A FOCUS ON BLOOD PLATELET TRANSFUSIONS

BY

MARYAM MOTAMEDI, M.Sc.

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF COMPUTING AND SOFTWARE

AND THE SCHOOL OF GRADUATE STUDIES

OF MCMASTER UNIVERSITY

IN PARTIAL FULFILMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Doctor of Philosophy (2023) McMaster University
(Department of Computing and Software) Hamilton, Ontario, Canada

TITLE: Demand Forecasting and Inventory Management of Perishable Inventory - with a Focus on Blood Platelet Transfusions

AUTHOR: Maryam Motamedi
M.Sc. (Information Technology),
University of Tehran, Tehran, Iran

SUPERVISOR: Dr. Douglas G. Down - Dr. Na Li

NUMBER OF PAGES: xiii, 167

Abstract

Inventory management of perishable products has seen extensive study over the years; the perishable nature capturing the real-world phenomena of expiration after a limited shelf life. Such problems are challenging as they involve balancing demand fulfillment with minimal wastage. An added dimension to such problems, given the rise of machine learning, is to estimate future demand. Demand forecasts can be helpful for decision making, in particular they can be used for finding the optimal ordering quantity for the products. The central thesis of this dissertation is that by forecasting the demand and utilizing it in the inventory management process, we can build a more robust inventory system that takes additional information into consideration when making decisions.

Firstly, five different demand forecasting methods, ARIMA (Auto Regressive Integrated Moving Average), Prophet, lasso regression (least absolute shrinkage and selection operator), random forest and LSTM (Long Short-Term Memory) networks are utilized and evaluated via a rolling window method. Subsequently, we study the structural properties of the optimal ordering policy for perishable products with fixed shelf lives in a periodic-review single-item inventory system over a finite horizon, where demand forecasts are available. The objective is to find the optimal ordering policy that minimizes the total expected cost, consisting of a linear ordering cost, inventory holding cost, wastage cost, and shortage cost, over a finite horizon. We show that the optimal policy is a state-dependent base-stock policy in which the base-stock values are a function of the system's state, the inventory level, a vector of current and previous demand forecasts, and previous demand values. Moreover, we explore the monotonicity properties of the optimal policy. The monotonicity properties motivate us to propose a heuristic in which the order quantity is an affine function of the inventory level and forecast-dependent target inventory levels. We evaluate the performance of the proposed heuristic on platelet transfusion data for hospitals in Hamilton, Ontario. Experimental results show that the proposed heuristic is effective in minimizing the total cost while maintaining low on-hand inventory levels.

Acknowledgements

Firstly, I would like to thank my supervisors Dr. Douglas G. Down and Dr. Na Li. I would like to thank Dr. Douglas G. Down for his unwavering support, patience, and guidance throughout my Ph.D. program. I greatly appreciate all of the time you have invested in understanding my work, no matter how intricate the details. Our engaging discussions, especially those involving collaborative problem-solving on the board, stand out as the highlights of my Ph.D. program. Additionally, I would like to thank you for giving me the opportunity to attend many conferences and engaging me in different projects. These opportunities not only expanded my academic horizons but also connected me with a broader academic community. I am grateful and honored for having the opportunity to work with and learn from you.

I would like to thank Dr. Na Li for her excellent support and guidance throughout this journey. Your openness to discussion and genuine interest in my ideas allowed me to seek guidance and share ideas. Thank you for believing in my abilities, caring about my future direction, and for always encouraging me to explore beyond my comfort zone. I am sincerely thankful for the considerable amount of time you dedicated to clarifying the challenging statistical and medical concepts and material for me. I am profoundly grateful for everything I have learned from you on both academic levels and personal levels.

Next, I would also like to express my gratitude to members of my thesis committee, Dr. Wenbo He, Dr. Fei Chiang, and Dr. Vahid Sarhangian for their insightful feedback and comments that helped me improve this dissertation.

I also would like to thank Jessica Dawson for her help during this work and for the enriching experience I gained while collaborating with her. I would also like to thank Dr. Nancy Heddle at the McMaster Centre for Transfusion Research (MCTR) for providing the data used in this study and the invaluable feedback on the practical implications of this study.

Finally, I would like to thank my family. I would like to thank my husband, Mohammad, for his endless love and support through my Ph.D. journey. Thank you for always being there for me and taking caring me. Last but not least, I must thank my parents for always supporting me throughout my academic life, without them this journey would not have been possible.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Perishable products include a wide range of items, from everyday grocery products to healthcare items that are critical to human health but less frequently used than everyday staples. Perishable products are subject to wastages owing to their limited shelf lives. In many cases, such as with healthcare products, the production cost of perishable products is high (Kaur et al., 2022; Stokes et al., 2018). Therefore, wastages not only result in wastage costs, but also lead to the loss of production costs for the wasted products, and the high wastage rates can impose a significant cost on the system. For instance, the global cost of food waste is projected to be around 1 trillion USD annually (Everitt et al., 2022). Furthermore, wastages have adverse effects on climate change in several ways. Firstly, food wastage decomposes and releases greenhouses gases which are approximately 28 times stronger than carbon dioxide (Everitt et al., 2022). Secondly, certain perishable products wastages such as the majority of healthcare product wastage cannot be recycled and reused (Jemai et al., 2020). Thirdly, the transportation of waste to landfills or incinerators can also contribute to climate change by emitting greenhouse gases from trucks and other vehicles. Shortages of perishable products can also have significant costs for the system, such as missed revenue opportunities, decreased customer satisfaction in retail, and limited access to or delays in essential treatments in healthcare that can lead to detrimental health outcomes.

The main underlying cause of wastages and shortages in a perishable inventory system is the unknown demand. This makes managing perishable inventory systems a challenging task from both the theoretical and computational perspectives. However, in today's data rich environment, demand history can be used to estimate future demand values, which can then be incorporated into the inventory system to determine ordering quantities. Demand estimates may be accurate, with a low forecast error, or less accurate with higher forecast error. In both cases, including them in the inventory model gives organizations access to added demand information. This dissertation is concerned with the problem of forecasting demand for perishable products and determining the optimal ordering quantities for a periodic-review single-item inventory system, with a focus on platelet products.

1.1 Practical Motivation: Demand Forecasting

In a supply chain, when demand fluctuations are significant, retailers often hold excess inventory to deal with high demand variation. However, holding surplus inventory makes demand forecasting even more challenging for distribution centres. This can result in the bullwhip effect, the increased variation in demand as a result of moving upstream in the supply chain (Croson and Donohue, 2006). This effect arises in a number of domains, including grocery supply chains (Dejonckheere et al., 2004) as well as in healthcare service-oriented supply chains (Samuel et al., 2010; Rutherford et al., 2016). Apart from this bullwhip effect, there are many uncertainties faced by suppliers due to the high demand variation. Since perishable products have a limited shelf life, high inventory levels result in excessive wastage, something that could be mitigated with better demand forecasting. On the other hand, low inventory levels increase the risk of shortages. Accordingly, accurately forecasting the demand for perishable products is a core requirement of a robust demand and supply management system.

This research was originally motivated by the platelet management problem confronted by Canadian Blood Services (CBS). CBS is responsible for providing blood products in Canada, excluding the province of Québec. Currently, there is a yearly wastage rate of about 9% for hospitals in Hamilton, Ontario (with an approximate cost of $400,000 per year) and about 15% for CBS with seasonal variation (Office of the Auditor General of Ontario, 2020). The current frequent same-day urgent orders, considered as shortages, are about 14% of the total orders in Hamilton, Ontario. Given the high wastage and shortage rates, forecasting short-term demand for platelets is of particular value.

In this dissertation, we forecast demand to overcome the mentioned challenges. In particular, we utilize multiple demand forecasting methods, including univariate analysis (time series methods) and multivariate analysis (regression and machine learning methods), and evaluate the performance of these models for platelet demand forecasting. The literature on platelet demand forecasting is limited, where forecasts are mostly based on demand history (Critchfield et al., 1985; Silva Filho et al., 2012, 2013; Kumari and Wijayanayake, 2016; Volken et al., 2018; Fanoodi et al., 2019). Platelet demand is determined by clinical characteristics and procedures such as product expiry guidelines and revisions in transfusion medicine protocols for platelet transfusion (Alcaina, 2020). These clinical characteristics and procedures may change over time. For example, there was an increase in platelet shelf life from five to seven days in September 2017 in Canada. Additional factors affecting demand can be captured by external predictors. Although the forecasting models in this dissertation are developed for platelet demand forecasting, we believe that they are generalizable for other perishable products when appropriate predictors are used.

1.2 Theoretical Motivation: Inventory Management with Forecasting

Consider a supply chain with one manufacturer and multiple retailers. The manufacturer produces a single perishable product with fixed shelf life. At the beginning of each day (period), an order is placed and received immediately. Then, demand is realized and it is satisfied. However, since demand is not known in advance, the inventory may not be sufficient to fulfill the demand. On the other hand, the inventory may be at a high level in comparison to realized demand, resulting in wastages of the product. Thus, as we can see inventory management is one of the major challenges in a perishable product supply chain since it directly affects the shortage and wastage rates. Consequently, it has received extensive attention in recent years, yet much remains to be done. For products with extremely short shelf lives and highly variable daily usage, managing the demand and supply is even more challenging. The primary challenge in perishable product inventory management is the unknown demand. Since demand is not known beforehand, the current system is very sensitive to demand fluctuations. High unexpected demand could result in several shortages, while consecutive low demands can result in wastage. As a result, the core of inventory management of perishable products is the ordering and issuing decisions to request the products and allocate them to customers.

A key issue when modelling (and analyzing) an inventory system is the choice of demand model. A large number of existing studies in the perishable inventory literature mostly assume that there is no information about the demand beyond its distribution. They are concerned with making ordering decisions assuming an independent and identically distributed (i.i.d.) sequence of demands, where the demand distribution could be estimated using historical demand data. Some include additional demand information, but in an indirect way. One approach is to estimate the distributional parameters and adjust the ordering policy to include the error in estimating the demand distribution (Prak et al., 2017; Trapero et al., 2019; Saoud et al., 2022). In practical settings, the i.i.d. assumption can be problematic (future demand may depend on previous demand) and there may be additional factors that may influence demand. Therefore, taking into account today's data-rich environment, it is possible to make forecasts about demand, based on historical demand and additional quantities that influence demand. Some recent studies consider data features in addition to previous demand values for forecasting the demand in the inventory model (Drackley et al., 2012; Khaldi et al., 2017; Guan et al., 2017; Li et al., 2021; Abouee-Mehrizi et al., 2022). In this dissertation, we use daily demand forecasts that are generated from a separate process and include them in the inventory model. The main goal of this dissertation is to design effective ordering policies for perishable products with short shelf life and highly variable demand by utilizing demand information in the inventory management process in a direct manner. From a research perspective, integration of forecasts in the inventory problem of perishable products is intriguing and can result in a richer framework for addressing the ordering problem. In this dissertation, we study the structural properties of the optimal ordering policy when demand forecasts are incorporated into ordering decisions. This is in contrast with recent work that includes factors that influence demand, but does not make explicit demand forecasts.

1.3 Summary of Main Contributions

A summary of the main results and contributions of this dissertation are:

- ***Demand forecasting for perishable products with a focus on platelet products:*** In Chapter 4, we study the platelet demand forecasting problem confronted by CBS. We progressively build five demand forecasting models (of increasing complexity) to forecast platelet demand. The proposed methods are applied to determine the influence of demand history as well as clinical predictors on demand forecasting. The first two models are univariate time series that only consider the demand history, while the remaining three methods, multivariate regression, random forest, and artificial neural networks, consider clinical predictors. These five methods are utilized to pursue the following goals: i) more precise platelet demand forecasting for the benefit of both CBS and hospitals, ii) reducing the bullwhip effect, as a consequence of effective demand forecasting; and iii) investigating the impact of predictors on the forecasting accuracy. More specifically:

  1. We analyze the time series of platelet transfusion data by decomposing it into trend, seasonality and residuals, and detect meaningful patterns such as weekday/weekend and holiday effects that should be considered in any platelet demand predictor.

  2. We utilize five different demand forecasting methods from univariate time series methods to multivariate methods including regression and machine learning. Since CBS has no access to recipients' demographic data, our first method, Autoregressive Integrated Moving Average (ARIMA), only considers demand history for forecasting, while the second model, Prophet, includes seasonalities,

trend changes and holiday effects. We found that these models have issues with respect to accuracy, in particular when a limited amount of data are available, accordingly we apply a lasso regression method to include clinical predictors for demand forecasting. Finally, random forests and Long Short-Term Memory (LSTM) networks are used for demand forecasting to explore the nonlinear dependencies among the clinical predictors and the demand.

3. We utilize predictors in the demand forecasting process, and select those that are most impactful by using lasso regression for structural variable selection and regularization. Results show that incorporating the predictors in demand forecasting enhances the forecasting accuracy while increases the interpretability of the models.

4. We investigate the effect of different amounts of data on the forecasting accuracy and model performance and provide a holistic evaluation and comparison for different forecasting methods to evaluate the effectiveness of these models for different data types, providing suggestions on using these robust demand forecasting strategies in different circumstances. Results show that when having a limited amount of data (two years in our case), multivariate models outperform the univariate models, whereas having a large amount of data (eight years in our case) results in the ARIMA model performing nearly as well as the multivariate methods.

To the best of our knowledge, this study is the first that utilizes and evaluates different demand forecasting methodologies from univariate time series to multivariate models for platelet products and explores the effect of the amount of available data on these approaches. By conducting a thorough data analysis and appropriately selecting model parameters, the suggested models can be extended to other perishable products, making them generalizable.

- ***Structural properties of the optimal ordering policy:*** In Chapter 5, we analyze the structural properties of optimal ordering policies in a periodic-review single-item inventory system over a finite horizon when demand forecasts are incorporated in the inventory model. The demand forecast in each period depends on the previous demand forecast value. The objective is to minimize the total cost which consists of per unit ordering cost, per unit holding cost, per unit shortage cost, and per unit wastage cost. We demonstrate that the optimal cost satisfies a structural property called $L^\natural$-convexity, meaning that it is convex and submodular. We use this property to show that the optimal ordering policy is a state-dependent base-stock policy that depends on the inventory levels, current and previous demand forecast values. To the best of our knowledge, this study is the first that shows the optimal ordering policy is a state-dependent base-stock policy by considering demand forecasts that have correlation with previous demand forecast values.

- ***A heuristic to approximate the optimal policy:*** Due to the complexity of the optimal policy, we propose a heuristic as a simpler alternative. The proposed heuristic is motivated by the structural results for the optimal ordering policy, in which the order quantity is an affine function of the inventory level and forecast-dependent order-up-to levels. Despite the simplicity of the proposed heuristic, our experimental results suggest that its performance is comparable to that of the optimal policy. Specifically, the heuristic yields low shortages and wastages, while keeping the on-hand inventory close to the actual demand. As a result, it is unnecessary to incorporate a large set of past forecast values into the ordering decisions. Moreover, in Chapter 6, we implement a data-driven version of the heuristic by considering just wastage and shortage costs. The proposed policy is simple and easy to interpret while allowing us to comment on the impact of the forecast accuracy. Moreover, a range of sensitivity analyses are presented to investigate the generalizability of the proposed heuristic.

1.4 Overview of Dissertation

This dissertation is organized as follows. Chapter 2 provides an overview of inventory management fundamentals, along with a review of the literature on demand forecasting, inventory management, and blood supply chain. In Chapter 3, we give a brief overview of the structure of the CBS supply chain and describe the data used for this study. Chapter 4 presents five forecasting models used for forecasting the demand and how they are applied and evaluated on platelet data. Chapter 5 provides structural analysis of an inventory model for a periodic-review single-item inventory system over a finite horizon when demand forecasts are incorporated. We explore the structural properties of the optimal policy for this system and develop an effective heuristic policy. In Chapter 6, we provide a holistic evaluation of the proposed heuristic and compare its performance with that of a base-stock policy. Finally, Chapter 7 concludes the dissertation and outlines directions for future work related to the inventory management of perishable products.

Chapter 2

Background and Literature

In this chapter, an overview of the literature related to the proposed research is presented. The review consists of four parts. Firstly, we review previous work on demand forecasting, with a focus on blood products. Secondly, we analyze the body of work on inventory management of perishable products. The literature on perishable inventory management is quite extensive and so our review is not intended to be exhaustive. We exclude mathematical details and proofs, but attempt to provide a high-level view of the relevant material. Thirdly, we provide an overview of work on blood supply chain and platelet inventory management. Lastly, we discuss the recent stream of research on data-driven inventory models. We conclude the chapter by presenting our view on how this work contributes to the current literature.

2.1 Forecasting Methods

Demand forecasting for perishable products has gained significant attention in recent years due to their limited shelf life, which leads to potential wastage or discarding of unused items. We review the underlying literature on perishable demand forecasting using both univariate models and multivariate models, and discuss our contributions in comparison to the existing literature.

2.1.1 Univariate vs. Multivariate Models

Current work on demand forecasting for perishable products can be grouped into two main streams. Many studies use univariate time series models to forecast future demand. In these studies, forecasts are based solely on previous demand values, without considering other features that may affect the demand. Auto Regressive Integrated Moving Average (ARIMA) is one of the most popular univariate models that is used for forecasting demand. Da Veiga et al. (2014) use ARIMA and Holt-Winters (HW) models for forecasting the demand for perishable food products. Huber et al. (2017) propose a hierarchy for forecasting demand for perishable products at different organizational levels, from company level to store level. For that, first a hierarchy is built based on organizational structure. Next, they cluster products' demands based on their features using k-means clustering. Finally, an ARIMA model (and an extended version) are used that allow for inclusion of variables in the forecasting model. Taylor and Letham (2018) propose a novel forecasting model, Prophet, which is designed to forecast events created on Facebook. The forecasts take into account typical characteristics of business time series: trends, seasonality, holiday effects, and outliers. Puchalsky et al. (2018) study time series demand forecasting of perishable food products in agribusiness. They analyze the performance of Wavelet Neural Networks (WNNs) in combination with five optimization methods. Mor et al. (2019) assess and compare the performance of moving average, regression, multiple regression, and Holt–Winters models for forecasting dairy products demand. Priyadarshi et al. (2019)

apply different forecasting methods, ARIMA, Gradient Boosting Regression (GBR), extreme GBR (XGBoost/XGBR), random forest regression, and a Long Short-Term Memory (LSTM) network to forecast daily demand for three vegetables. Based on their data and considering only previous demand values, an LSTM network and random forest regression produce the best results in comparison to other models. Huber and Stuckenschmidt (2020) propose a large-scale demand forecasting model for a bakery chain by focusing on special calendar days (special occasions). They use Artificial Neural Networks (ANN), Recurrent Neural Networks (RNN), and Gradient-Boosted Regression Trees (GBRTs) to forecast daily demand at the store level.

Several recent studies include additional features other than demand history for demand forecasting. Du et al. (2013) propose an algorithm based on Support Vector Machines (SVM) to forecast demand of perishable farm products. Their model considers historical sales data and uses fuzzy methods to generate inputs for the SVM model. The inputs of the model are historical demand data and weather information that impact the demand. Van Donselaar et al. (2016) analyze the effect of promotions on a product's sales, and develop and assess different regression models and a moving average model to forecast the demand for these perishable products during promotions. Yang and Sutrisno (2018) apply a regression model and ANN on perishable products in a franchise business to forecast the short-term demand. Müller et al. (2020) consider a multi-product newsvendor problem with unknown demand distributions. They forecast the demand using an Exponential Smoothing (ES) model, and solve the traditional multi-product newsvendor problem based on the forecasts generated by the forecasting model. Furthermore, they propose a data-driven solution by using an ANN to calculate the optimal order quantities from data, streamlining the process into a single step.

2.1.1.1 Blood Products

In this section, we focus on demand forecasting for blood products in terms of univariate and multivariate models. Most studies on demand forecasting for blood products focus on univariate models (Frankfurter et al., 1974; Fortsch and Khapalova, 2016; Lestari et al., 2017). In terms of platelet products, there is a limited literature on platelet demand forecasting; most investigate univariate time series methods. In these studies, forecasts are based solely on previous demand values, without considering other features that may affect the demand. Critchfield et al. (1985) develop models for forecasting platelet usage in a blood centre using several time series methods including Moving Average (MA), Winter's method and Exponential Smoothing (ES). Silva Filho et al. (2012) develop a Box-Jenkins Seasonal Autoregressive Integrated Moving Average (BJ-SARIMA) model to forecast weekly demand for blood components, including platelets, in hospitals. They later extend this work in (Silva Filho et al., 2013). Kumari and Wijayanayake (2016) propose a blood inventory management model for the daily supply of platelets focusing on reducing platelet shortages by applying three time series methods, MA, Weighted Moving Average (WMA) and ES. Volken et al. (2018) use generalized additive regression and time-series models with ES to predict future whole blood donation, including platelets, and Red Blood Cells (RBC) transfusion trends. Fanoodi et al. (2019) use artificial neural networks and ARIMA models to forecast platelet demand by considering daily demands for eight types of blood platelets. They consider different demand data lags, 1 day, 2 days, 3 days, 4 days, 5 days, 6 days, 1 week, 15 days, 30 days, 90 days, 120 days, and 365 days, as the input data for the artificial neural networks.

In terms of including additional features for demand forecasting of blood products, there are a few papers that include additional information to forecast the demand. Drackley et al. (2012) estimate long-term blood demand for Ontario, Canada using age and sex-specific patterns derived from the historical transfusion records of patients. They forecast blood supply and demand for Ontario by considering demand and supply patterns, and demographic forecasts, with the assumption of fixed patterns and rates over time. Khaldi et al. (2017) apply Artificial Neural Networks (ANNs) to forecast the monthly demand for three blood components, RBCs, platelets, and plasma, for a case study in Morocco. Guan et al. (2017); Li et al. (2021); Abouee-Mehrizi et al. (2022) forecast the demand for blood products but they focus mainly on inventory management of blood products, not the forecast itself. Guan et al. (2017) propose an optimization ordering strategy in which they forecast the platelet demand for several days into the future and build an optimal ordering policy based on the predicted demand, concentrating on minimizing the wastage. Their main focus is on an optimal ordering policy and they integrate their demand model in the inventory management problem, meaning that they do not try to precisely forecast the platelet demand. Li et al. (2021) develop a hybrid model consisting of seasonal and trend decomposition using Loess (STL) time series and eXtreme Gradient Boosting (XGBoost) for RBC demand forecasting and incorporate it in an inventory management problem. Abouee-Mehrizi et al. (2022) estimate future platelet order quantities based on demand features and by using a lasso regression model. Twumasi and Twumasi (2022) apply K-Nearest Neighbour regression (KNN), Generalised Regression Neural Network (GRNN), Neural Network Autoregressive (NNAR), Multi-Layer Perceptron (MLP), Extreme Learning Machine (ELM), and an LSTM network for forecasting and backcasting blood demand to predict future and lost past demand data respectively, by using a rolling-origin procedure.

Table 2.1 provides an overview of the reviewed literature on blood demand forecasting.

Table 2.1: Literature review on forecasting methods (Blood products: RBC(Red Blood Cell) - PLT (platelet). For donation: WB (Whole Blood))

Author	Blood product	Real data	Demand forecasting	Supply forecasting	Criteria for evaluating
Critchfield et al. (1985)	PLT	✓	✓		Min total cost
Drackley et al. (2012)	RBC	✓	✓	✓	-
Silva Filho et al. (2012)	RBC - PLT		✓		Min shortage cost
Silva Filho et al. (2013)	PLT		✓		Min wastage cost
Kumari and Wijayanayake (2016)	PLT	✓	✓		Min shortage level
Guan et al. (2017)	PLT	✓	✓		Min wastage level
Khaldi et al. (2017)	RBC - PLT - Plasma	✓	✓		Min total cost
Fortsch and Perera (2018)	WB	✓		✓	Min wastage and shortage level
Volken et al. (2018)	WB - Not Specified	✓	✓	✓	Min shortage level
Fanoodi et al. (2019)	PLT	✓	✓		Min wastage and production cost
Li et al. (2021)	RBC	✓	✓		Min total cost
Abouee-Mehrizi et al. (2022)	PLT	✓	✓		Min total cost
Twumasi and Twumasi (2022)	Not Specified	✓	✓		Min wastage and shortage level

2.1.2 Discussion

As we have seen, the literature related to perishable demand forecasting is extensive, but mainly focuses on demand history. Although univariate models like ARIMA generally perform well, when future demand patterns are likely to change or there is uncertainty in future demand patterns, these models may not perform well. Some studies consider data features that influence the demand in addition to demand history, but they do not perform a holistic data analysis to consider the most relevant features for demand forecasting. Most of the proposed models consider only a few additional features along with demand for previous weeks/months as predictors in the forecasting models. Moreover, various measures are used in different studies for evaluating the forecasting models. Some studies evaluate the forecasts based on minimizing the wastage or shortage or both. Others consider additional costs such as inventory costs or production costs. Also, different metrics have been used in papers for evaluating the forecasting accuracy. While employing a pooled error statistic proves intriguing for quantifying the comprehensive forecasting error during model comparisons, distributional exploration of errors helps with providing a comprehensive and nuanced forecast evaluation. In general, there is a lack of standardized or consistent metrics for reporting and comparing errors among forecasting models.

Our problem of interest is different from the literature on forecasting demand for perishable products in several aspects. Firstly, to our knowledge, it is the first research that utilizes different demand forecasting methods for forecasting platelet demand by including predictors in the demand forecasting process. Secondly, we investigate the effect of different amounts of data for training and the frequency of retraining for univariate and multivariate models. Thirdly, we provide holistic evaluation of the forecasting models and provide managerial insights for blood centres.

2.2 Perishable Inventory Management

In this section, we provide a foundation for the investigation of the integration of demand forecasting and inventory management by surveying the relevant literature. First, we provide an overview of inventory management fundamentals, followed by a review of the literature on the inventory management of perishable products. The literature on inventory management of perishable products is extensive, and so we have chosen to exclude certain mathematical details and proofs, but attempt to provide a high-level review of the relevant material.

2.2.1 Inventory Management Fundamentals

A supply chain as defined by Beamon (1998) is "a structured manufacturing process wherein raw materials are transformed into finished goods, then delivered to end customers". Figure 2.1 shows the general structure of a supply chain. As we can see in the figure, a supply chain has different levels, typically commencing with suppliers and ending with the consumers. Also, as we see in the figure, a retailer orders products from a manufacturer and so there may be an amount of time between when an order is placed to replenish products and when the order is received, called the "lead time". Inventory management is a critical element of the supply chain and is performed at various levels throughout the supply chain, encompassing activities that span from supplier to the retailer. According to the American Production and Inventory Society (APICS), inventory management is defined as "the branch of business management concerned with planning and controlling inventories" (Toomey, 2000). The main objective of inventory management is to ensure that the inventory (stock) level of a product is kept at a desired level. Inventory management is a challenging task since the system should plan based on customer demand and system costs. When the product is perishable, this task becomes even more challenging. Perishable products are products that have a limited shelf life and normally are wasted after their shelf life is passed. As a result, this should be considered in managing the inventory of perishable products and wastages must be avoided as much as possible.

Inventory management problems normally consist of minimizing a cost function (maximizing a profit function). In general, there are five cost types in a perishable inventory management problem:

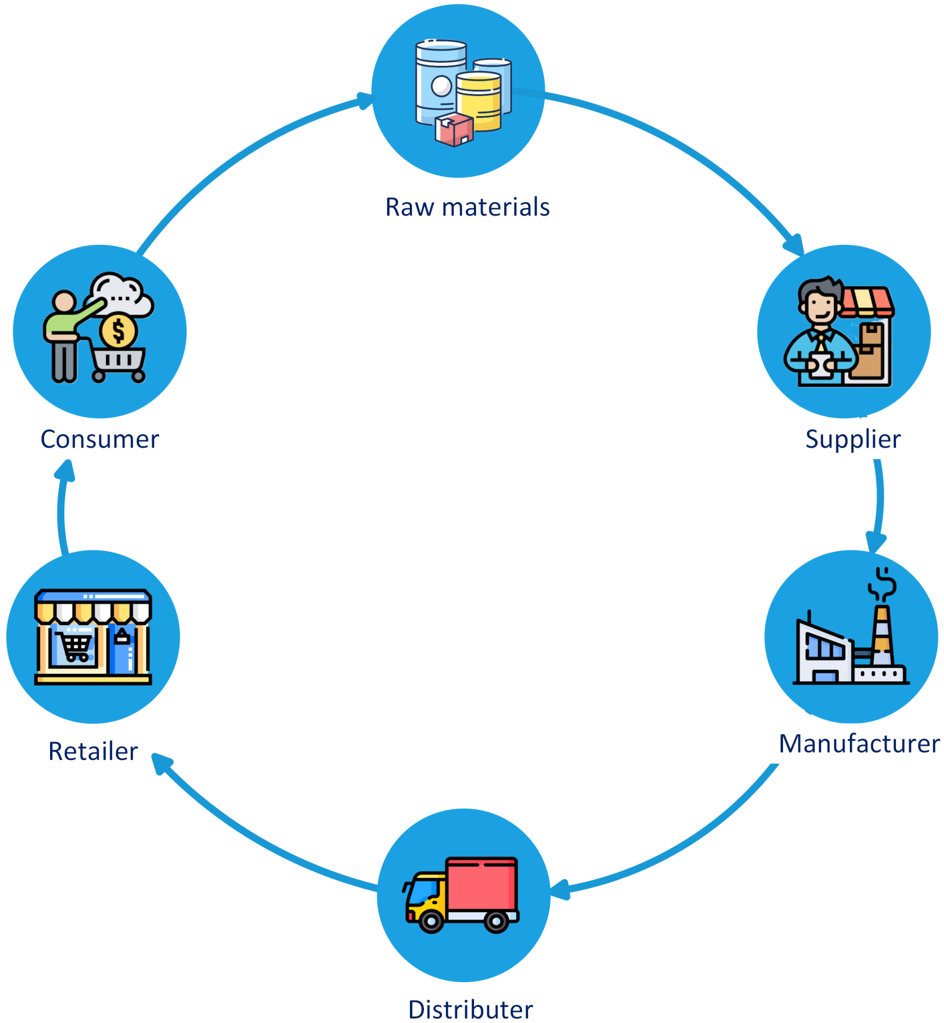

Figure 2.1: General structure of a supply chain

- ***purchase/production cost:*** the cost of producing or purchasing an item. When purchasing multiple items, a typical assumption is that this is the cost per unit multiplied by the order quantity.

- ***ordering cost:*** a fixed ordering cost related to the order. These costs are normally independent of size of the order, and are mainly related to transportation costs.

- ***holding cost:*** the cost of keeping an item in the inventory. It consists of the storage costs and material handling costs needed for products with special storage requirements. A typical assumption is that this is a per-unit cost.

- ***shortage cost:*** arises only when an item is unavailable upon demand. In this case, there are two possible outcomes. One possibility is that the order is backlogged, which incurs additional costs to the system. This is called a backlogging/backorder inventory system. If backlogging is not possible (e.g. the customer cannot wait), the order is lost which results in the loss of profit and customer dissatisfaction. The latter is called a lost-sales inventory system.

- ***wastage cost:*** arises only when an item's shelf life is passed. This cost is only related to perishable products and usually is considered equal to the product's production/purchase cost. If disposal of the item needs proper management, additional disposal cost of the wasted item would be added to the system.

In the literature, different ordering policies have been proposed for perishable (and nonperishable) products. A well-known ordering policy is called the (s,S) model, in which the inventory level is continuously reviewed, and an order is placed when the inventory level reaches a reorder level, s. The order quantity brings the inventory level back up to the order-up-to level, S. Figure 2.2 is a representation of the (s,S) inventory model. Some studies consider an (r,Q) model, which is similar to the (s,S) model. In an (r,Q) model, when the inventory position reaches the reorder level r, a consistent order quantity Q is placed. A special case of the (s,S) model is the base-stock policy in which the inventory level is continuously reviewed so that an order is placed when the inventory level drops to below S. The order quantity brings the inventory level back up to S.

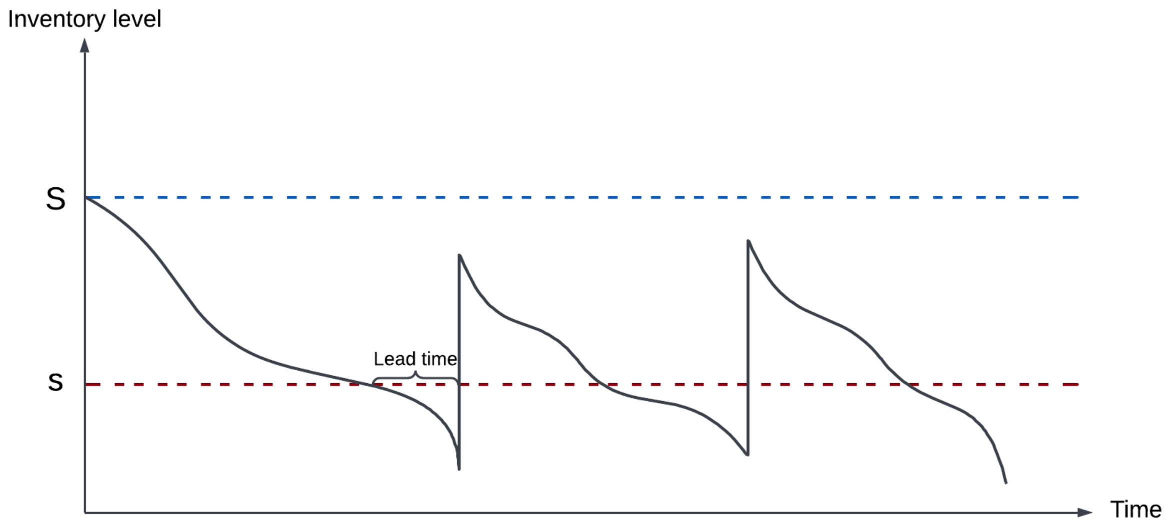

Figure 2.2: The (s, S) inventory model

The inventory ordering systems examined in the literature can be categorized from different views. For instance, in the literature, a differentiation is made between single-period models and multi-period models. A single-period model, also known as the newsvendor problem, is a periodic-review model in which inventory planning is performed for one period. At the end of the period, remaining products are usually disposed. This is a typical model for seasonal products, bakery items, and newspapers (this is of course only a partial list). It is also widely used as an approximation for multi-period models. A multi-period inventory problem refers to the management of inventory over multiple time periods, taking into account various factors such as demand, order lead times, and inventory costs. Moreover, unlike the single-period problem, remaining inventory is carried over from one period to the next, which makes the inventory management task more complicated.

In the remainder of this chapter, we explore the literature on inventory management of perishable products from different perspectives. For a more detailed review of the literature on inventory management we refer the reader to Karaesmen et al. (2011); Nahmias (2011); Perera and Sethi (2023a,b).

2.2.2 Periodic-review vs. Continuous-review Models

The literature on inventory management problems classifies them into two main categories: periodic-review systems and continuous-review systems. In a continuous-review system, the inventory is continuously tracked and its status is always known (Silver et al., 2016). However, in practice, it is not always necessary to continuously monitor the inventory, and the inventory is tracked at specific periods of time such as daily, upon demand, upon shipment, etc. This approach is referred to as a periodic-review inventory model.

2.2.2.1 Periodic-review Models

The literature on periodic-review perishable inventory systems is quite extensive. We review periodic-review models for perishable products based on a fixed or a random shelf life. Most studies focus on products with deterministic shelf life; such models are called fixed shelf life models. On the other hand, models in which the shelf life is not deterministic are called random shelf life models. Unlike products with a fixed shelf life, in models considering a random shelf life, the shelf life of product is not known before arrival, which makes the inventory management of such products more challenging.

Fixed Shelf Life

Inventory management and structural analysis of optimal policies for fixed shelf life perishable products for periodic-review models were pioneered by Fries (1975) and Nahmias (1975), and were revisited later by Nandakumar and Morton (1993). Minner and Transchel (2010) study a periodic-review inventory model for perishable products by considering different service-level constraints. They consider a lost-sales system under both First In First Out (FIFO) and Last In First Out (LIFO) issuing policies. Haijema (2011) proposes a stochastic model for calculating the optimal issuing policy for perishable products with a short shelf life. The model is formulated and solved as a Markov Decision Problem (MDP). Chen et al. (2014) study a joint pricing and inventory control problem for a perishable product with a fixed shelf life for both backlogging and lost-sales cases by considering linear ordering, backlogging or lost-sales penalty, holding, and disposal costs over a finite horizon. They show monotonicity properties of the optimal policies and calculated bounds for the optimal ordering policy. They also propose a heuristic as an alternative for the optimal policy. Muriana (2016) studies the Economic Order Quantity (EOQ) model and proposes a stochastic model for perishable food products by considering wastage and shortage costs along with the incorporation of the shelf life as a parameter to detect the number of wasted units in the inventory. Fu et al. (2019) focus on a periodic-review inventory model for newly manufactured or remanufactured perishable products and develop policies to find the optimal ordering amounts for both products. They show that as the age of units in the inventory increases, the order quantity converges to a fixed value.

Random Shelf Life

The literature on periodic-review inventory models with random shelf life is limited. Chen et al. (2014) extend their work on fixed shelf life periodic-review inventory (explained in the previous subsection) to random shelf life, under the condition that the inventories become wasted in the same sequence as they enter the system. Kouki et al. (2014) examine a periodic-review inventory problem for perishable products under a (T,S) inventory policy and consider both lost-sales and backlogging cases. In a (T,S) model, orders are scheduled at regular intervals of T periods, with each order quantity setting the inventory to the order-up-to-level, S. They assume that demand follows a Poisson process, and the products' shelf lives are exponentially distributed. They model their problem as a Markov process and optimize the expected cost for the system. In a later work, Kouki et al. (2015) analyze a base-stock inventory system for perishable products with Markovian demand and general shelf life and lead time distributions. They show monotonicity properties of the optimal cost and calculate the optimal base stock for the problem. Ketzenberg et al. (2018) consider a periodic-review inventory problem with random shelf life and determine the optimal expiration date, derived from the shelf life distribution, which affects the optimal order quantities. Abouee-Mehrizi et al. (2022) study a periodic-review perishable inventory problem with a random shelf life and zero lead time over a finite horizon. They propose two different models, one with a fixed shelf life and one that considers the random shelf life. For the random shelf life model, they examine the worst-case scenario when variability in the remaining shelf life is not included the model. Clarkson et al. (2023) consider a periodic-review model of a single product with a random and age-dependent shelf life over a finite horizon. They model the problem as an MDP and show that convexity holds in the penultimate period.

2.2.2.2 Continuous-review Models

The literature on continuous-review inventory models is extensive and diverse, most of which analyze inventory models for perishable products with random shelf life. In this section, similar to the periodic-review models, we review the continuous-review models based on product's shelf life being fixed or random.

Fixed Shelf Life

Continuous-review inventory models emerged as extensions of periodic-review models. There is a substantial body of literature on continuous-review inventory dedicated to the (s,S) inventory model. Weiss (1980) presents an extension of the continuous-review (s,S) policy for a perishable inventory system with fixed self life and zero lead time. They assume that demand follows a Poisson process. Subsequently, researchers extended their model to various settings. For instance, Liu and Lian (1999a) and Liu and Lian (1999b) present models with general renewal demand processes. Ravichandran (1995) investigate a continuous review $(S-1,S)$ inventory model with lost sales. They consider a random lead time, fixed shelf life, and a Poisson demand process and explore the optimal ordering policy. Perry and Posner (1998) study the continuous review $(S-1,S)$ inventory model for a system with fixed shelf life and constant lead times, where demand is modelled as a Poisson process. Olsson (2010) studies three variations of the continuous review $(S-1,S)$ model, penalizing backorders by quantity, penalizing backorders by both quantity and duration, and imposing a service level constraint. The optimal S value is obtained for the first and third model, while an approximation is provided for that for the second model. Baron et al. (2010) study a continuous review (s,S) inventory model of perishable products arriving in batches. They consider deterministic and exponentially distributed shelf lives, and two types of compound Poisson processes for demand. Using the Queueing and Markov Chain Decomposition methodology, Barron and Baron (2020) investigate a continuous review (s,S) policy. Their analysis consists of stochastic lead time, perishability, and state-dependent Poisson demand process. In a related study, Barron (2019) expand the work of Barron and Baron (2020) by incorporating demand uncertainty and stochastic batch demands into the model.

Some studies explore variations of the classic (s,S) model. Berk and Gürler (2008) analyze the (r,Q) ordering policy for a product with constant shelf life and lead time. They consider a lost-sales system and model the system's dynamics under an (r,Q) policy using an embedded Markov process. Kouki et al. (2016) consider multiple perishable items with random lifetimes under a continuous review (s,c,S) policy. This policy is a variation of the classic (s,S) model that considers two reorder points, s as a "must order" reorder point and c as a "potential" reorder point. They model the system as a Markov process and calculate the optimal parameters for the model using a decomposition approach under zero lead time. Rajendran and Srinivas (2020) propose two variants of review policies for platelet inventory management that are based on an (s,S,Q) policy. By considering additional inventory thresholds, either an order quantity of Q is placed or the inventory is filled to S. The policies are determined using stochastic mixed integer programming. Current work employing stock-level policies consider static base stock levels. Such policies are easy to implement, but they can result in large on-hand inventory levels. Benjaafar et al. (2011) consider a single-product production-inventory system that uses advanced demand information in the form of customers' preannouncements of their orders. They model their system as a continuous-time MDP and prove that the optimal policy is a state-dependent base-stock policy. Moreover, they show the monotonicity properties of the ordering policy and build four heuristics as approximations of the optimal policy.

Random Shelf Life

Considering a random shelf life for continuous-review inventory management is very common in the literature and there is a large stream on continuous-review inventory models for perishable products with random shelf life. One of the main reasons is that some papers consider a Markovian shelf life which simplifies the problem. For instance, Kalpakam and

Sapna (1994), Liu and Shi (1999) study the optimal (s,S) ordering policy for products with exponentially distributed shelf life and order lead time, Poisson demand process, and lost sales. Gürler and Özkaya (2008) analyze a continuous-review perishable inventory system under an (s,S) ordering policy. Products are assumed to have a random shelf life with a general distribution and the lead time is zero. The work of Kouki et al. (2015) described under the periodic-review models also deals with continuous-review models. Baron et al. (2020) analyze an (s,S) continuous-review model for perishable inventory with exponential shelf life and order lead time. They consider a lost-sales model under two types of demand distribution, Poisson and compound Poisson with general sizes. Kouki et al. (2020) analyze a base-stock perishable inventory system with general shelf life and order lead time under a continuous-review model. They model the system by using a queueing network model and show the optimal cost's monotonicity properties.

2.2.3 Single-echelon vs. Multi-echelon Models

In the literature, inventory models usually either consider a single facility (level) within a supply chain network or consider multiple facilities (levels) within the supply chain network. If the inventory model is focused on determining the optimal inventory levels of a single facility, it is called a single-echelon inventory model. On the other hand, if the model considers the optimization of inventory levels across multiple facilities in a supply chain, it is called a multi-echelon inventory model.

Single-echelon models are generally simpler to implement and manage in comparison to multi-echelon models since they are focused on optimizing the inventory at a single location. They do not consider the communication between different levels of the supply chain. Moreover, focusing on one facility can be advantageous since it allows decision makers to only consider specific constraints and requirements of that particular facility. Most of the papers in this area only consider a single facility in a supply chain. Duong et al. (2018) provide a review for single-echelon continuous review inventory management models of perishable products.

Multi-echelon models allow for optimization of inventory levels across multiple locations within the supply chain. Such models are closer to real life scenarios since they consider coordination and collaboration between different units in a supply chain. However, modelling a multi-echelon inventory system is more complex since the decision for one facility affects the total system performance. Most research on multi-echelon perishable inventory models focuses on two-echelon models. Weraikat et al. (2016) study a two-echelon pharmaceutical reverse supply chain in which wasted products are collected by a third-party company. They model the coordination between a producer and the companies responsible for collecting wasted medications from customers. They introduce incentives to customers for returning leftover products. Mahmoodi et al. (2016) propose approximation models for optimization of base-stock values in a two-echelon continuous-review perishable products inventory system consisting of a warehouse and a retailer. Tiwari et al. (2018) consider a two-echelon supply chain with deteriorating items that have price-dependent and stock-dependent demand, and analyze integrated and non-integrated policies for the joint optimization model. Hamdan and Diabat (2019) study a two-stage stochastic problem for RBC products to minimize wastages, supply chain costs, and product delivery time by considering production, inventory, and location decisions. Their model considers four echelons, mobile blood facilities, local blood banks, regional blood banks, and hospitals. The problem is solved as a Mixed Integer Programming (MIP) problem for a real case study from

Jordan. Amiri et al. (2020) propose a model for calculating the optimal sales level of perishable products in a two-echelon supply chain with one vendor and multiple buyers. They optimize the sales profit by an exact model and three heuristics, Particle Swarm Optimization (PSO), Co-evolutionary Particle Swarm Optimization (CPSO) and Genetic Algorithm (GA). Moshtagh et al. (2022) optimize the operational interaction between blood centers and hospitals within a two-echelon blood supply chain. To that end, they study three channel structures, including centralized, decentralized, and coordinated systems, and compare their performance.

Some studies consider multi-echelon models with more than two echelons. Xu et al. (2019) combine simulation and an improved particle swarm algorithm to develop a simulation-based optimization model for a three-level agricultural products inventory system. Shen et al. (2020) study the collaborative inventory management of agricultural products in a three-echelon system: farmers' professional cooperatives, distribution centers, and supermarkets. They use an improved genetic algorithm to maximize the profit by adjusting supply and ordering decisions across different levels of the supply chain.

2.2.4 Theoretical Models vs. Ordering Heuristics

Inventory models can be studied by explicitly computing the optimal ordering policy, or failing that, to identify structural properties that the optimal policies satisfy. This stream of work, although intriguing and valuable, is often quite challenging. Specifically, considering dynamics of the inventory system along with the necessity to account for the age of perishable products results in a fairly complex problem. Some papers study the structure of the optimal policy. However, the complexity of the optimal policy may make it difficult to give the optimal policy explicitly. Even if that were possible, the policy may be so complex that it is problematic to implement. Thus, optimal policies are often approximated using insights on their structure. Our work is closely related to the stream of literature on the structure of optimal ordering decisions and the development of ordering heuristics in perishable inventory management.

2.2.4.1 Theoretical Models

Work on analysis of the optimal policy dates back to Fries (1975) and Nahmias (1975), where they characterize the structure of the optimal policy for perishable products with multiperiod shelf life and zero lead time. They show that the optimal order quantity is decreasing in the inventory levels of different ages. However, due to the complexity of their analysis, their results are difficult to generalize. Subsequent works include the work of Cohen (1976) that studies the complexity of the optimal policy and shows that finding the optimal policy is quite complex, even for a very simple problem of a two-period shelf life product. Williams and Patuwo (1999) determine the optimal order quantity for a single period, periodic-review, two-period shelf life product system with lost sales. Zipkin (2008) introduced a new method for analyzing the structural properties of the standard lost-sales inventory system, which is more convenient in comparison to the existing methods, and establishes new bounds for the optimal policy. Haijema (2014) explores optimal ordering, issuing, and disposal policies for perishable products by considering age in the inventory problem. They show that optimal ordering, issuing, and disposal policies are all stock-age dependent. Fang et al. (2021) propose a model for dynamic pricing and ordering decisions for multi-period perishable and substitutable products, and study the structural properties of the optimal decisions. Zhang et al. (2022) study platelet inventory management for a hospital in the US. They prove that a myopic transshipment policy is optimal for their case study, which exhibits short shelf life and asymmetric demand. Furthermore, they show that myopic transshipment is a lower bound for more general cases. For a more comprehensive review of the literature on analysis of optimal policies for perishable inventory models, we refer readers to (Perera and Sethi, 2023b,a; Goldberg et al., 2021).

One direction for characterizing the structure of optimal policies for discrete-time discrete-state periodic review problems is exploring $L^\natural$-convexity and multimodularity. Multimodularity was first introduced by Hajek (1985). $L^\natural$-convexity is closely related to the concept of multimodularity in discrete-time systems and was first developed by Murota (1998). It was used as one of the techniques to analyze the structural properties of optimal policies in discrete-time systems. Lu and Song (2005) use $L^\natural$-convexity to explore optimal policies in an assemble-to-order inventory system and to develop an exact algorithm for an order-based model. Gong and Chao (2013) use $L^\natural$-convexity to examine the optimal expected total discounted cost for periodic-review inventory systems with remanufacturing and finite capacities. They show that the optimal policies are a modified remanufacture-down-to policy and modified total-up-to policy for systems with remanufacturing and without remanufacturing, respectively. Chen et al. (2018) develop a transformation technique for converting a nonconvex minimization problem to its equivalent convex problem. This transformation allows for the preservation of some desirable structural properties like convexity, $L^\natural$-convexity, and submodularity. Li and Yu (2014) explore structural properties of optimal policies in perishable inventory systems by using multimodularity. Following their work, several recent studies have also employed the concept of multimodularity to derive structural properties of optimal decisions (Liu et al., 2019; Chen et al., 2019). Among the papers discussed in the previous paragraph, Zipkin (2008) and Chen et al. (2014) study structural properties of the optimal policy by proving the $L^\natural$-convexity of the cost function in their inventory problems. Using a similar approach to that used in these two papers, we explore the structural properties of the optimal policy for our problem. For a more detailed study of the properties and applications of $L^\natural$-convexity, we refer readers to (Topkis, 1998; Simchi-Levi et al., 2005; Chen, 2017).

2.2.4.2 Ordering Heuristics

Due to the complexity of finding the optimal policy for perishable inventory problems and the curse of dimensionality, some studies focus on heuristics as an alternative for the optimal policy. For example, Nahmias (1976); Deniz et al. (2010) focus on reducing state variable dimensions by aggregating state variables. On the other hand, several papers propose approximation approaches. A base-stock policy is a well-known widely-used heuristic, see Nahmias (1977); Nandakumar and Morton (1993); Cooper (2001). Gürler and Özkaya (2008) propose a heuristic for a continuous-review perishable inventory system under an (s,S) ordering policy and by considering a strictly positive lead time. The heuristic adjusts the reorder level, s, and the order-up-to level, S, based on the expected demand during the lead time. Deniz et al. (2010) develop a heuristic that considers substitution between age-dependent demand for issuing and ordering policies for perishable products. Li et al. (2016) study ordering and clearance sales decisions for perishable products selling under a LIFO issuing policy and develop two myopic heuristic policies that reduce state variable dimensions. Chao et al. (2015) develop approximation approaches based on worst-case performance by considering features that make demand nonstationary, such as seasonality, to characterize their heuristic algorithm. Chao et al. (2018) propose an easy-to-compute approximation algorithm for perishable inventory systems with positive lead times and finite ordering capacities, and demonstrate that it offers a theoretical worst-case performance guarantee. Chen et al. (2019) apply a lookahead heuristic method to obtain an approximation for optimal platelet collection, production, issuing, and disposal decisions. Chen et al. (2021) investigate inventory management of perishable products by considering different types of demand, various lost-sale costs, and freshness level requirements. Based on the structure of the optimal policy, they develop a novel approximation technique called the adaptive approximation approach, which is nearly optimal with an average optimality gap of 0.30% for the example that they consider. Clarkson et al. (2023) develop two new heuristic policies, a newsvendor heuristic and periodic-review heuristic, for a single product periodic-review model. Products are perishable with random age-dependent shelf lives. In a recent study, Bu et al. (2023) perform asymptotic analysis for a periodic-review perishable inventory system with zero lead time over an infinite horizon, and show that a base-stock policy is asymptotically optimal when using a FIFO issuing policy.

2.2.5 Age-dependent Models

In some applications, perishable products with different freshness levels are required. For instance, in a healthcare system, there may be different demands that require products with different ages. Some studies consider these different demand types in the supply chain and study age-differentiated demands that require products with different freshness levels.

There are several inventory models considering age-differentiated demands in the blood supply chain. Haijema (2013) proposes a novel stock-level ordering policy, (s,S,q,Q), which is a variant of the periodic review (s,S) policy that restricts the order quantity by a minimum, q, and a maximum, Q. Two types of issuing policies, LIFO and FIFO, are considered and the optimal ordering policy is calculated by formulating an MDP and approximating its solution using simulation. In a subsequent work, Haijema (2014) considers both the stock level and product's age for inventory management and uses stochastic dynamic programming for defining ordering, issuing, and disposal policies. Civelek et al. (2015) discuss the optimal ordering and allocation policies for a platelet supply chain by considering different demands requiring platelet units with different ages. In addition to classical inventory costs, they consider a mismatch cost for satisfying demand with a product of an age that differs from the required age. Gunpinar and Centeno (2015) propose an inventory management model for RBC and platelets to control shortage and wastage. In their paper, they consider two types of demands, blood items with different freshness. In addition, unlike most papers in the literature, they consider the crossmatch-to-transfusion ration (C/T ratio) and release period. Ensafian et al. (2017) develop a two-stage stochastic inventory management model to minimize shortage, wastage, and other platelet supply chain costs. They consider three types of demand based on the age of blood units and blood group compatibility rules. In addition, they consider apheresis as well as whole blood collection methods. However, they only focus on platelets although apheresis yields different products. In addition, they do not consider differences between these two methods in terms of availability and the amount of blood yielded. Kara and Dogan (2018) propose two inventory management policies for perishable products, a stock-age dependent policy and a quantity dependent policy. The policies are derived using reinforcement learning. Rajendran and Ravindran (2019) propose a platelet inventory management model by considering two types of demand, regular and emergency. They develop a stochastic integer programming model under an (s, S) ordering policy and use a variant of a genetic algorithm to solve it. By considering two demand types requiring platelet units with different freshness, Chen et al. (2019) develop a joint collection and production model for the platelet supply chain. They model their problem using stochastic dynamic programming and address their extended problem using a lookahead heuristic. By considering different types of patients, Ensafian and Yaghoubi (2017) develop a bi-objective platelet supply chain under both FIFO and LIFO policies to determine the best inventory and production decisions. Unlike most works in the literature, their model assumes the apheresis method as an alternative to whole blood decomposition for the production of platelets. They optimize the cost and freshness of platelets using a robust mixed integer program and demonstrate the applicability of their model in a case study. Larimi et al. (2019) discuss lateral transshipment in a platelet supply chain in which different collection methods including apheresis, platelet-rich plasma, and buffy coat methods are investigated for obtaining typical, irradiated and washed platelets for different demand types requiring platelets with different age. They discuss the risk of production under the different collection methods and develop a bi-objective robust optimization problem to explore inventory decisions in collection, test, production, and distribution processes. Chen et al. (2021) study periodic-review single-product inventory management with multiple demand classes, each having specific freshness requirements. They consider perishable products with fixed shelf lives under lost sales and assume zero lead time and fixed costs. They characterize several monotonicity properties of the optimal policy and propose an ordering heuristic.

2.3 Blood Supply Chain Management

A Blood Supply Chain (BSC) consists of the collection, testing, production, and distribution of blood from donors to patients. These patients can be either routine (regular) patients who require blood products as a part of their treatment or emergency patients such as patients with the need of a blood product for surgeries or emergency trauma treatment. Therefore, there is a vital need to have a sufficient supply of blood products; otherwise, loss of lives may result. The subject of BSC management has attracted a vast amount of attention from researchers in different areas. Different BSC configurations ranging from internal collection centres inside of hospitals to multiple mobile and permanent donation centres have been studied. Consequently, the structure of blood supply chains may differ from country to country. Generally, a blood supply chain consists of three main facilities (levels), including collection (donation) centres, blood centres, and hospitals. There are various aspects to a blood supply chain such as designing the system, decision-making about the process such as collection, transportation and inventory management, and forecasting.

The work of Jennings (1973) is one of the earliest works explicitly considering the structure of a blood supply chain. It was followed by several works considering different aspects of a blood supply chain. Or and Pierskalla (1979) study the transportation location-allocation problem of blood banks by determining the optimal number and location of blood banks, optimal allocation of hospitals to the blood banks and optimal routing of blood products with respect to the system costs. Cohen and Pierskalla (1979) propose an ordering policy for the red blood cell supply chain as a base-stock value that depends on daily demand, the average transfusion to cross match ratio, and the cross-match release period. Prastacos (1984) provides an overview of blood inventory management from both theoretical and practical perspectives. Pierskalla (2004) studies supply chain management practices and challenges of blood banks. The author discusses various aspects including inventory management, distribution, transportation, and quality control, and suggests strategies to improve the efficiency and effectiveness of blood supply chains. For more comprehensive literature reviews on blood supply chain management, see (Beliën and Forcé, 2012; Pirabán et al., 2019).

Some studies focus on the design issues of a BSC (Ensafian and Yaghoubi, 2017; Zahiri and Pishvaee, 2017; Ramezanian and Behboodi, 2017; Osorio et al., 2018; Eskandari and Sharifi, 2018; Samani et al., 2019; Hamdan and Diabat, 2020; Khalilpourazari and Hashemi Doulabi, 2022). There is a large number of studies in the literature addressing location, allocation and inventory management policies in a blood supply chain. Zahiri et al. (2015) study location and allocation issues in a blood supply chain with both mobile and fixed donation centres by using a robust mixed-integer programming method for modelling the problem. A case study is applied to illustrate their model in a real-world situation. Dillon et al. (2017) formulate a two-stage inventory model to specify periodic review policies for RBC (Red Blood Cell) inventory management. The uncertain demand is dealt with by forecasting the demand for the next 12 weeks by producing a sample of 100 scenarios using a Monte Carlo sampling approach. Attari et al. (2018) propose a stochastic bi-objective design problem to make decisions about location, allocation, and inventory in the collection and distribution phases of the blood supply chain, and apply Benders decomposition for solving their large-scale optimization problem. Bruno et al. (2019) explore a reorganization of a regional blood system in Italy by developing a facility location model to make the collection process more efficient. By considering both costs and maximum unsatisfied demand, Zahiri and Pishvaee (2017) develop a multi-product uncertain blood supply chain network design model. They consider different ABO blood types and ABO-substitution in their model. Robust optimization is used for handling the uncertainty of the model and solving the fuzzy mathematical programming problem.

Several studies develop Markov chain models for ordering and issuing problems in a blood supply chain. Pegels and Jelmert (1970) propose a discrete time Markov chain model and explore the impact of two issuing policies, modifications of FIFO and LIFO policies, on the age of the transfused units as well as on average inventory levels, and consequently on the shortage and wastage probabilities. Brodheim et al. (1975) model a number of inventory and distribution policies for blood products as discrete-time Markov chains. They assume that orders are for fixed quantities that are added to the inventory at regular intervals. Their results indicate that the FIFO policy has better performance compared to other issuing policies. Bar-Lev et al. (2017) propose a stochastic model for blood inventory management by considering supply and demand as Poisson processes. A Level Crossing Technique (LCT) is used for calculating the steady-state inventory level, but due to the complexity of LCTs for general underlying distributions, fluid and diffusion limits for the steady state inventory level are obtained. Sarhangian et al. (2018) study the performance of several issuing policies for RBC using a continuous-time Markov chain model. They propose a new threshold-based issuing policy to decrease the age of transfused RBCs which is a two-stage process and a combination of FIFO and LIFO policies. The model is an $M/M/1+D$ Markov queue in which blood units expire after D days. To formulate their threshold-based policy, they used the results of Parlar et al. (2011) to obtain the Laplace Transform of sojourn time of blood units in inventory, and ultimately the performance of their threshold-based policy compared with that of FIFO and LIFO in terms of average age of transfused units, and proportion of wastages and lost demand. The main critique of the papers that are based on Markov models is that they assume that order arrivals follow a Poisson process rather than being under control of the hospitals. However, in real-world problems, hospitals place orders for blood products, and as a result, the arrivals do not follow a Poisson process.

2.3.1 Platelet Inventory Management

There are a number of works focusing on platelet inventory management. Sirelson and Brodheim (1991) propose a predictive model for the inventory control of platelets by associating a base-stock level with the shortage and wastage rates using a linear regression method. Blake (2010) presents an overview of platelet inventory and ordering problems from the perspective of both producers and customers. He also gives a literature review of perishable inventory models related to blood products followed by a discussion of current challenges. In a separate work, Blake (2009) discusses work by Van Dijk et al. (2009) and the use of operational research for managing platelet inventory and ordering. Van Dijk et al. (2009) formulate a platelet inventory problem as a dynamic programming problem and find a nearly optimal order-up-to production policy. Zhou et al. (2011) develop a dynamic programming model to study the optimal order-up-to ordering policy for a platelet supply chain by considering regular and optional orders. Haijema (2014) considers the stock level and product age and uses stochastic dynamic programming for defining ordering, issuing, and disposal policies. They consider FIFO and LIFO policies and provide a comparison for the model under these two policies. Abdulwahab and Wahab (2014) study an inventory model for platelets with stochastic supply and demand, and deterministic lead time. They develop their model using dynamic programming, and the optimal ordering policy is approximated by a newsvendor model. Rajendran and Ravindran (2017) develop a stochastic integer programming problem to determine ordering policies for platelets focusing on minimizing the wastage. They propose three heuristic ordering strategies and compare their performances with a base-stock policy. Abouee-Mehrizi et al. (2023) propose an Approximate Dynamic Programming (ADP) model to study the optimal ordering quantity for platelet inventory management under endogenous uncertainty of shelf life. They evaluate performance of the proposed policy with real data for hospital in Hamilton, Ontario.

While platelet demand is determined by clinical factors that are not included in an i.i.d. demand sequence, the discussed works mainly consider an i.i.d. demand sequence and do not include additional demand information in the inventory management process.

2.4 Data-driven Models

There is a recent stream of studies that incorporate additional demand information in the inventory management process. These studies can be categorized into two main groups: (i) prediction and optimization as a single step, (ii) predict then optimize. The second group is close to the work of this dissertation presented in Chapter 5.

In the first group, demand forecasts are included in the inventory optimization problem, rather than being a separate process. These models include additional demand information in the inventory model indirectly. In other words, there is no separate process for forecasting the demand, and demand is predicted inside the inventory model. Guan et al. (2017) propose a convex optimization problem in which they forecast the platelet demand for several days into the future and build an optimal ordering policy based on the predicted demand, concentrating on minimizing the wastage while maintaining a minimum inventory level. Ban and Rudin (2019) study a data-driven newsvendor model that considers different observations of features that influence the demand. They propose algorithms based on Empirical Risk Minimization (ERM) and Kernel-weights Optimization (KO) approaches. Closely related to the work of Guan et al. (2017), Abouee-Mehrizi et al. (2022) consider a periodic review, perishable inventory control problem over a finite horizon, with zero lead-time and propose two models, a fixed age model and a robust model, for platelet inventory management. The objective is to determine daily ordering quantities while minimizing the wastage and shortage costs over this finite horizon. Demand is satisfied according to the Oldest-Unit-First-Out (OUFO) allocation policy, and unsatisfied demand is lost. The demand forecasts are included in the inventory optimization problem, rather than being a separate process. These models include additional demand information in the inventory model indirectly. In other words, there is no separate process for forecasting the demand, and demand is predicted inside the inventory model.

The second group that utilize data in the inventory model follow a two step process of first forecasting the demand and next using demand forecasts for optimizing inventory decisions. A classical example of this approach is presented in Elmachtoub and Grigas (2022). Li et al. (2021) propose a data-driven multi-period inventory problem for RBC products that includes RBC demand predictions. They forecast the RBC demand and incorporate the forecasts in the inventory model. Since forecast errors exist, they introduce two extra decision variables, target inventory and reorder constraints, to control these errors in the ordering policies. Their model is a variation of the classical (s,S) policy in which they define S to compensate for demand overestimations, and define s to compensate for demand underestimations. Both s and S are calculated based on the data and the predicted demand. Closely related to this stream, we follow a two-step process of first forecasting the demand for perishable products, and then utilizing demand forecasts in the inventory management process. We believe that by considering demand forecasting as a separate process, suppliers can strategically make decisions in various parts of the supply chain, such as production planning and resource allocation.

2.5 Summary

Most studies on perishable inventory consider an i.i.d. demand sequence assumption for their models. Many of the papers we reviewed in this dissertation assume that the demand follows a Poisson process. Some studies criticize the i.i.d. demand sequence assumption in inventory models. In classical inventory management theory, safety stocks are established based on the standard deviation of past demand forecasting errors over a lead time. Some recent studies provide solutions for adjusting the ordering policy to include the error in estimating the demand distribution. Prak et al. (2017) show when the i.i.d. assumption is violated, it may lead to significantly underestimated safety stock levels. They present corrected lead time demand variance expressions and reorder levels for correcting the lead time forecast errors when the lead time is constant. Trapero et al. (2019) use a non-parametric kernel density estimation and parametric GARCH (Generalised AutoRegressive Conditional Heteroscedastic) models to estimate the distribution of lead time forecast errors, and use them to generate inventory safety stock levels. Saoud et al. (2022) suggest estimating the lead time variance of forecast errors, rather than estimating the point forecast error variance, and extending it over the lead time interval. This new research stream suggests an estimation of demand distribution. Our work is in a similar direction as we do not consider an i.i.d. demand sequence since such distributions cannot capture the underlying factors that affect the demand.

Some recent studies include additional demand information in the inventory problem. Abouee-Mehrizi et al. (2022); Guan et al. (2017) incorporate additional factors that influence demand in the inventory problem, but do not explicitly forecast the demand. Demand forecasting and inventory management are considered as a single process. In (Benjaafar et al., 2011) additional demand information is in the form of additional information provided by the customers. Levi et al. (2015); Zheng et al. (2016); Huber et al. (2019); Keskin et al. (2021) present data-driven solutions for the newsvendor problem by considering the empirical forecast error distribution rather than a demand distribution assumption. Given the abundance of data available today, it is possible to construct demand forecasts that can be incorporated into the inventory system as additional demand information for determining ordering quantities. Many organizations perform demand forecasting as part of their decision-making processes, and incorporating these forecasts into the inventory model can be beneficial. Thus, when demand data are available, one can benefit from including additional information in the inventory management process. From the practical point of view, accurate demand forecasting itself is important for supply chain management purposes. Accurate demand forecasting can be used for decision making in many parts of the supply chain such as production planning and resources and staff management. Since there may be some fundamental limit to how accurate the forecasts can be, one important challenge would be how to use the demand forecasts in an effective manner that can take into account how good the forecasts are.

In this dissertation, we study the problem of inventory management of perishable products in the presence of previous demand information. We forecast the future demand based on demand history and explore the structural properties of the optimal ordering policy when demand forecasts are incorporated into ordering decisions. Unlike previous works, in this research demand forecasting and inventory management are considered as two different processes. In practice, demand forecasts can be generated from time series estimators such as ARIMA models or LSTM networks, in which forecasts are generated based on the correlation to previous demand and demand forecast values. Thus, in this dissertation, we consider these correlations in the inventory model for studying the structural properties of the optimal policy. We show that there exists an optimal ordering policy that is a base-stock policy that depends on the state, i.e. inventory levels, current and previous forecast values, and (indirectly) previous demand values. Moreover, by using the structural results, we propose a heuristic that integrates demand forecasts in the ordering policy in a simple and intuitive manner.

Chapter 3

Data Description and Analysis

In this chapter, we provide a description of the data used for this study, constructed by processing Canadian Blood Services (CBS) shipping data and the TRUST (Transfusion Research for Utilization, Surveillance and Tracking) database at the McMaster Centre for Transfusion Research (MCTR) for platelet transfusion in hospitals in the city of Hamilton, Ontario. The study is approved by the Canadian Blood Services Research Ethics Board and the Hamilton Integrated Research Ethics Board (HiREB number 7293).

Two organizations, CBS and Héma-Québec, are responsible for providing blood products and services in transfusion and transplantation for Canadian patients. The former operates within all Canadian provinces and territories excluding Québec, while the latter is in charge of the province of Québec. The current blood supply chain for CBS is an integrated network consisting of a regional CBS distribution centre and several hospitals, as illustrated in Figure 3.1. Currently, there are nine regional blood centres operating for CBS, each covering the demand for several hospitals (Hospital Liaison Specialists, 2020). Hospitals request blood products from the regional blood centres for the next day, yet the regional blood centres are not aware of the actual demands as each hospital has its own blood bank.

Furthermore, recipient demographics and hospital inventory management systems are not disclosed to CBS or the regional blood centres.

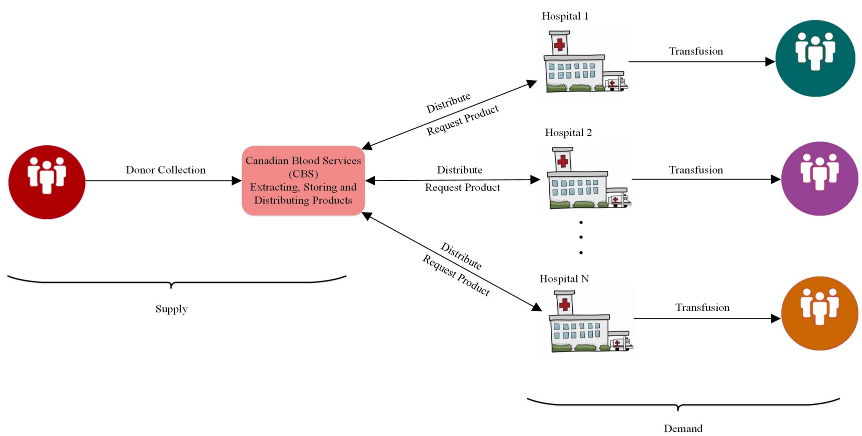

Figure 3.1: CBS blood supply chain with one regional blood centre and multiple hospitals

We consider a blood supply system consisting of CBS and four major hospitals operating in Hamilton, namely, Hamilton General Hospital, Juravinski Hospital, McMaster University Medical Centre (MUMC), and St. Joseph's (STJ) Healthcare Hamilton. These hospital blood banks operate with a Transfusion Medicine (TM) laboratory team to manage blood product transfusions to patients. We study platelet transfusions in the described blood supply system.

Platelet products are a vital component of patient treatment for bleeding problems, cancer, AIDS, hepatitis, kidney or liver diseases, traumatology and in surgeries such as cardiovascular surgery and organ transplants (Kumar et al., 2015). In addition, miscellaneous platelet usage and supply are associated with several factors such as patients with severe bleeding, trauma patients, aging population and emergence of a pandemic like COVID-19 (Stanworth et al., 2020). The first two factors affect the uncertain demand pattern, while the latter two factors result in donor reduction. Platelet products have a shelf life of five to seven days before considering test and screening processes that typically last two days (Fontaine et al., 2009), so the remaining shelf life of the platelets that arrive at the hospitals is typically between three to five days. The extremely short shelf life along with the highly variable daily platelet usage makes platelet demand and supply management a highly challenging task, invoking a robust blood product demand and supply system. Platelet products are collected and produced at CBS and after testing for viruses and bacteria (a process which lasts two days), platelets are ready to be shipped to hospitals and transfused to patients.

As a result of internal inventory management practices, these four hospitals are considered as one entity. At the beginning of the day, hospitals receive platelet products that were ordered on the previous day, from CBS. In the case of shortages, hospitals can place expedited (same-day) orders at a higher cost. Prior to September 2017, platelets had five days of shelf life, while after this date, the shelf life of platelets was increased to seven days. After exceeding the shelf life, platelet products are expired and discarded.

3.1 Data Description

We study a large clinical database with 61377 platelet transfusions for 47496 patients in hospitals in Hamilton, Ontario from 2010 to 2018. We analyze the database to extract trends and patterns, and find relations between the demand and clinical predictors. The data are high dimensional, with more than 100 variables that can be divided into four main groups: 1. the blood inventory data such as product name and type, received date, expiry date, 2. patient characteristics such as age, gender, patient ABO Rh blood type, 3. the transfusion location such as intensive care, cardiovascular surgery, hematology, and 4. available laboratory test results for each patient such as platelet count, hemoglobin level, creatinine level, and red cell distribution width. The laboratory tests are prescribed by physicians based on clinical needs and can help to decide whether a patient needs platelet transfusion. In this research, the laboratory test results are processed and used along with other information to forecast future platelet demand. A summary of the data collection and cleaning process is presented in Appendix A.

Additionally, we add new calculated predictors such as the number of platelet transfusions in the previous day and previous week, the number of received units in the previous day, and day of the week. Table 3.1 gives the set of predictors that are used in this study along with their descriptions. These predictors are selected by a lasso regression model (Tibshirani, 1996) which is explained in detail in Section 4.1.2.1. As we can see from Table 3.1, predictors have different ranges, and hence are standardized by z-score normalization. All data processing and analysis and model implementations are carried out using the Python 3.7 programming language.

Table 3.1: Data variable definition and description

Name	Description
abnormal_ALP	Number of patients with abnormal alkaline phosphatase
abnormal_MPV	Number of patients with abnormal mean platelet volume
abnormal_hematocrit	Number of patients with abnormal hematocrit
abnormal_PO2	Number of patients with abnormal partial pressure of oxygen
abnormal_creatinine	Number of patients with abnormal creatinine
abnormal_INR	Number of patients with abnormal international normalized ratio
abnormal_MCHb	Number of patients with abnormal mean corpuscular hemoglobin
abnormal_MCHb_conc	Number of patients with abnormal mean corpuscular hemoglobin concentration
abnormal_hb	Number of patients with abnormal hemoglobin
abnormal_mcv	Number of patients with abnormal mean corpuscular volume
abnormal_plt	Number of patients with abnormal platelet count
abnormal_redcellwidth	Number of patients with abnormal red cell distribution width
abnormal_wbc	Number of patients with abnormal white cell count
abnormal_ALC	Number of patients with abnormal absolute lymphocyte count
location_GeneralMedicine	Number of patients in general medicine
location_Hematology	Number of patients in hematology
location_IntensiveCare	Number of patients in intensive care
location_CardiovascularSurgery	Number of patients in cardiovascular surgery
location_Pediatric	Number of patients in pediatrics
Monday	Indicating the day of the week
Tuesday	Indicating the day of the week
Wednesday	Indicating the day of the week
Thursday	Indicating the day of the week
Friday	Indicating the day of the week
Saturday	Indicating the day of the week
Sunday	Indicating the day of the week
lastWeek_Usage	Number of units transfused in the previous week
yesterday_Usage	Number of platelet units transfused in the previous day
yesterday_ReceivedUnits	Number of units received by the hospital in the previous day

3.2 Exploratory Analysis for Trends, Seasonality and Holiday Patterns

In order to propose a short-term demand forecasting model, we first explore the data for identifying temporal (daily/monthly) patterns that can inform our demand forecasting techniques. In particular, we investigate seasonality, day of the week, non-stationarity effects, and correlations among the predictors. The data analysis ranges from 2010/01/01 to 2018/12/31. An initial observation is that the demand is highly variable, with a transfused daily average of 17.90 units and a standard deviation of 7.05 units.

Observations for non-stationarity: The Augmented Dickey-Fuller (ADF) test (Cheung and Lai, 1995) is applied on the time series data to examine the stationarity. The results of the ADF test show that the data is not stationary (P value = 0.085) before 2016, but it becomes stationary from 2016 onwards (P value <0.001).

Observations for seasonality: We apply the Seasonal and Trend decomposition using Loess (STL) model to decompose the time series data into trend, seasonality, and residuals. We also apply the one-way Analysis of Variance (ANOVA) test to compare the means of the transfused units in different months, and the means of the transfused units during weekdays and weekends. Moreover, we explore the trend, holidays, weekly seasonality, and yearly seasonality using the Prophet model, explained in detail in Section 4.1.1 (such detailed understanding is not required at this point).

Figure 3.2 shows the time series data decomposition using the STL model. As we can see in the seasonal part, there are recurring temporal patterns in the data. The results of the one-way ANOVA test also show that there is a significant difference between the means of the daily transfusions during weekdays and weekends (F = 5.13, P value <0.001) and the means of daily transfusions in different months (F = 3.94, P value <0.001), which provide strong evidence in favour of the presence of weekly and monthly seasonalities.

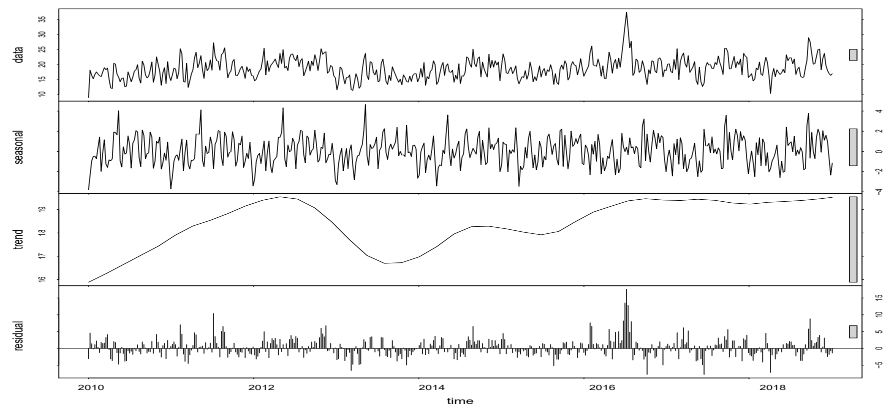

Figure 3.2: Time series decomposition using STL method

Since the data becomes stationary from 2016 onwards, we also explore the trend, holidays, weekly seasonality, and yearly seasonality (seasonality within a year) starting from 2016 using the Prophet model. As we can see from Figure 3.3, there is a downward trend from the beginning of 2016 to July 2017 and an upward trend from July 2017 to the end of 2018. Almost all holidays have a negative effect on the model, except for July 1st. This means that the demand is lower than regular weekdays for almost all of the holidays, except for July 1st.

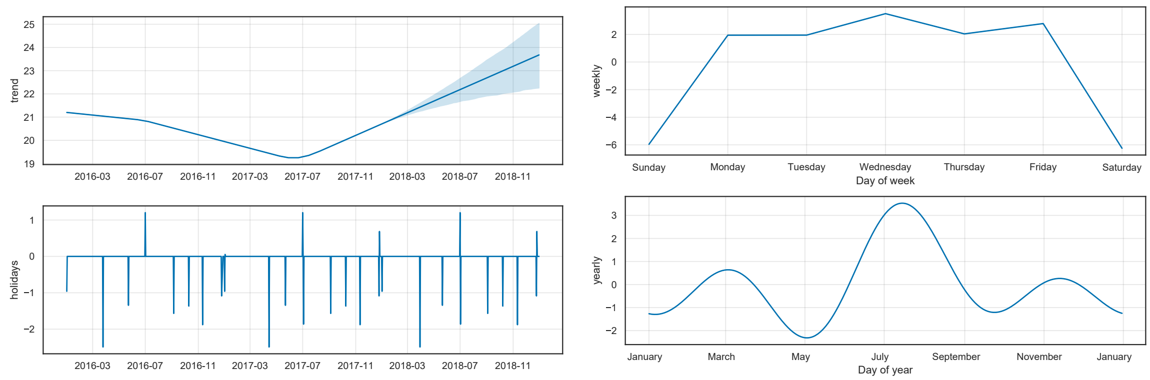

Figure 3.3: Prophet model for exploring trends, holiday effects, weekly and yearly seasonality - since these components are combined through a generalized additive model, the values of y-axes in the plots represent the quantity to be added to or subtracted on a specific day

We can also see that there is weekly seasonality in which Wednesdays have the highest demand while the weekends have the lowest demand. Moreover, the yearly seasonality, captured by Fourier series in the Prophet model, depicts three cycles: 1. January to May in which March has the highest demand while May has the lowest demand; 2. May to September in which the demand is highly variable. July has the highest demand in this cycle and the highest demand of all months while May has the lowest demand in the cycle and also the lowest demand of all months; 3. September to January with a slight variation in demand - November with the highest and January with the lowest demands.

Observations for day of the week effect: We also compare the mean daily units transfused based on day of the week by plotting the mean against day of the week, and also by applying the t-test to compare the mean daily units transfused during weekdays and weekends.

Figure 3.4 compares the mean daily units transfused and the mean daily units received based on day of the week, month, and year. Reviewing the figures, it can be concluded that generally, the range of the daily units received is greater than the range of total transfused units in a given, day, month, or year which suggests higher variability for the total number of received units (standard deviation of 9.33) compared to transfused units (standard deviation of 7.04). Indeed, this has its roots in the bullwhip effect, that is hospitals tend to order more than their actual demand. We see an opportunity to better coordinate supply (number of units received) with demand (number of units transfused) through the development of a daily demand predictor.

Figure 3.4(a) shows that the number of received units is more than the transfused units over the years. However, there does not appear to be any pattern between the number of received and transfused units, suggesting that there is no yearly impact on platelet demand. Figure 3.4(b) indicates that there is a near-uniform pattern for the number of received and transfused units by month.

Finally, in Figure 3.4(c), it can be noted that on Mondays the number of received units tends to be significantly larger than for the weekend. It is also noticeable that on weekends the number of received and transfused units clearly differ from the weekdays due to lower staffing levels over the weekends. On Saturday, on average, the number of transfused units is lower than the number of received units, but on Sunday, the total number of received units drops considerably, resulting in a large gap between the number of received and transfused units. This appears to be compensated for by the total number of received units on Monday. This is an additional effect that could be mitigated by better coordination between supply and demand. Moreover, there is a significant difference in the mean daily platelet usage when comparing weekdays to weekends (weekday = mean [sd]: 21.20 [6.22], weekend = mean [sd]: 12.37 [4.60], t-test: 95% confidence interval for the difference in means: (7.97, 10.34), P value <0.001). Consequently, there is a clear weekday/weekend effect, in agreement with Figure 3.3, which appears to be caused by various reasons including lower staffing levels and operating hours over the weekends and prophylactic platelet transfusions to cancer patients on Fridays to ensure that their platelet counts remain sufficiently high over the weekend.

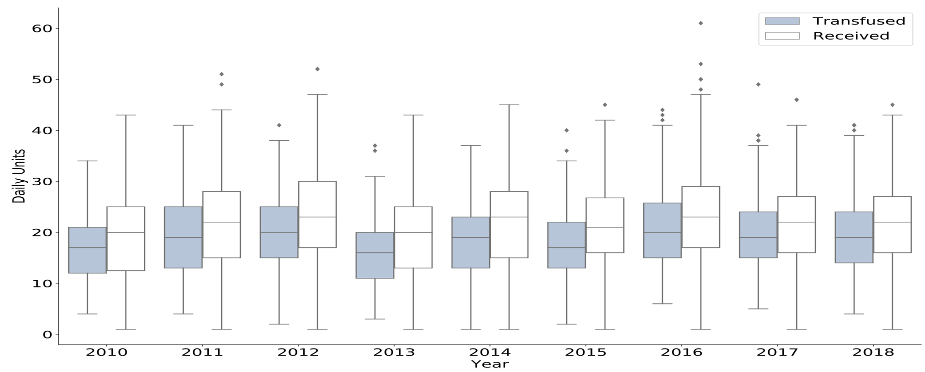

(a) Mean daily units transfused vs. mean daily units received (year)

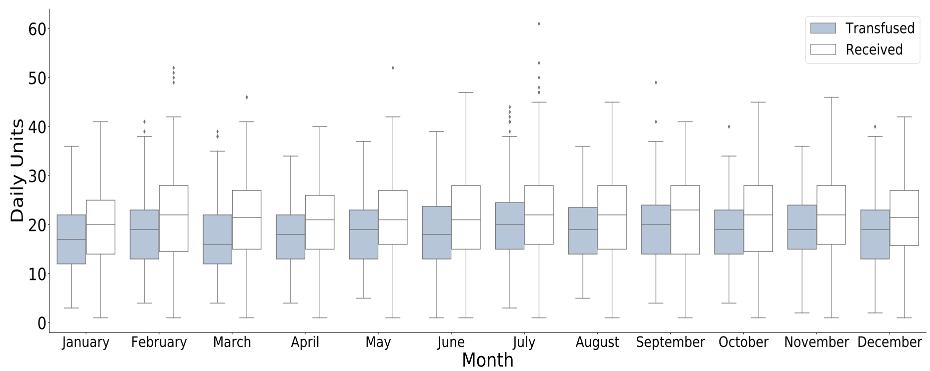

(b) Mean daily units transfused vs. mean daily units received (month)

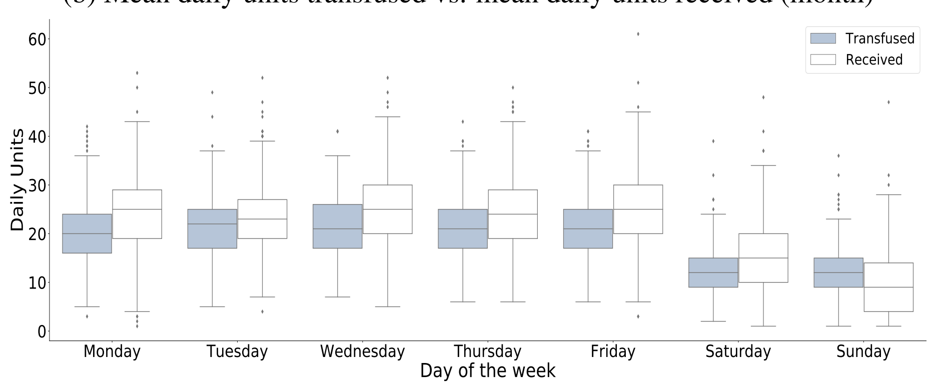

(c) Mean daily units transfused vs. mean daily units received (day-of-the-week)

Figure 3.4: Mean daily units transfused vs. mean daily units received

Observations for highly correlated predictors: One of the data characteristics is that clinical predictors are highly correlated. These high correlations can affect the performance of a regression model, mainly because of the violation of model assumptions. Moreover, it can reduce the interpretablity of the model, a small change in the model or data may result in unexpected changes in the predictors' coefficients.

We calculate the Pearson correlation between the selected predictors. The Pearson correlation measures the linear relationship between two variables, ranging from -1 to 1, where -1 corresponds to a perfect negative correlation and 1 corresponds to a perfect positive correlation. As shown in Figure 3.5, the predictors, in particular the daily numbers of patients with abnormal laboratory test results, are highly correlated. These high correlations give rise to some challenges when the predictors are considered in the demand forecasting process, as discussed in Table B.1 of Appendix B.

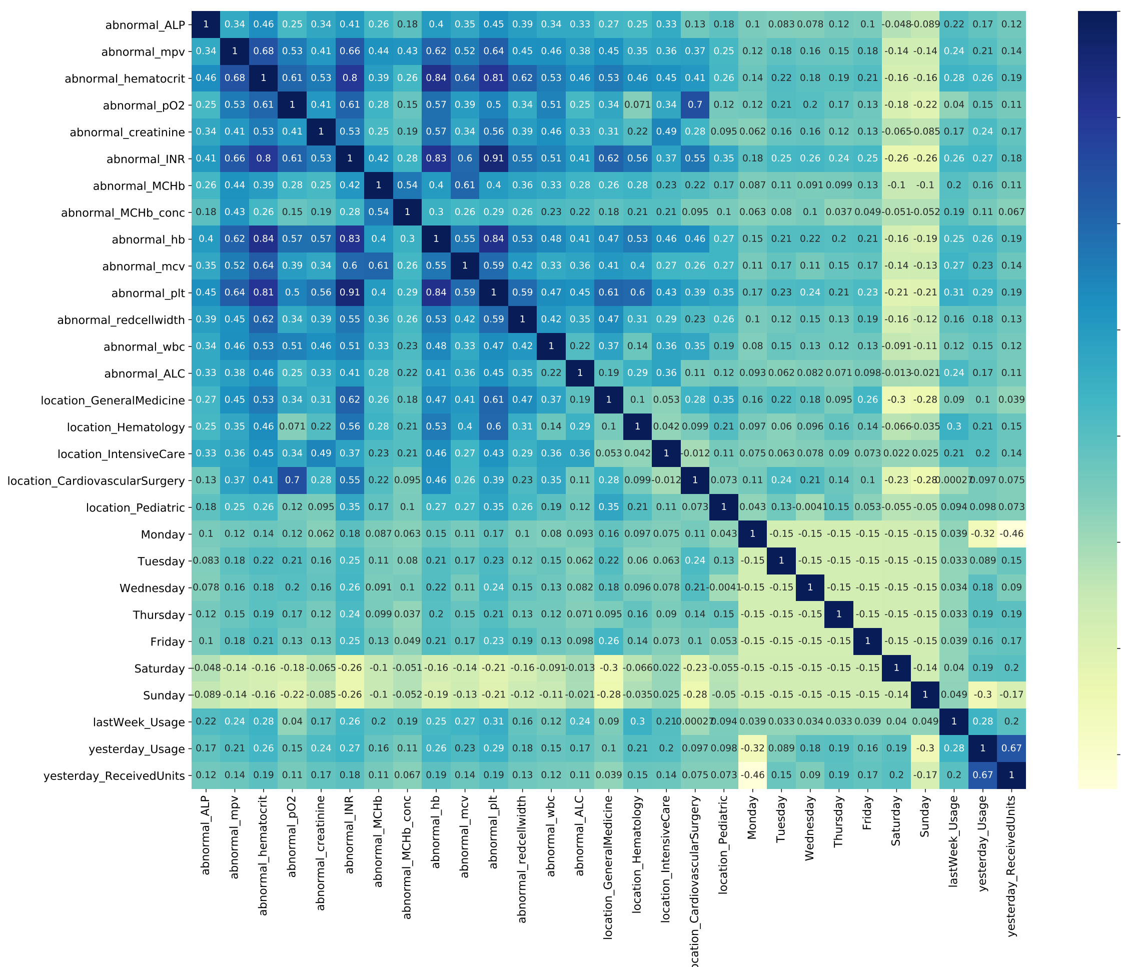

Figure 3.5: Pearson correlation among variables

Chapter 4

Demand Forecasting: from Univariate Time Series to Multivariate Models

In this chapter, we provide a description of multiple demand forecasting methods, including univariate analysis and multivariate analysis for forecasting the demand in general, and evaluate the performance of these models for platelet demand forecasting by using a rolling window analysis for retraining the models. First, we discuss the forecasting models used for forecasting the demand. After that, we conduct a comprehensive evaluation of these models when applied to different amounts of data. Lastly, we compare and analyze the models based on different metrics, and offer overall recommendations.

4.1 Demand Forecasting Models

This section explains the five forecasting models used for forecasting the platelet demand in Hamilton hospitals. The Autoregressive Integrated Moving Average (ARIMA) and Prophet models are univariate models that forecast the demand based on demand history. Univariate models can be used to forecast the demand for other blood products since they only use previous demand values. Lasso regression (least absolute shrinkage and selection operator), random forest, and Long Short-Term Memory (LSTM) networks are multivariate models that consider various predictors in addition to demand history for forecasting the demand. These models can also be extended to other blood products since the variable selection is carried out through an automated process. The training of these models is not contingent on a specific blood product.

4.1.1 Univariate Models

The univariate models, ARIMA and Prophet, forecast the demand solely based on the previous demand values. The ARIMA model does not consider seasonality in data and is considered as a baseline model. The Prophet model, on the other hand, considers trend, seasonality, and holidays for forecasting the demand.

4.1.1.1 ARIMA Model

An autoregressive integrated moving average model consists of three components, an autoregressive (AR) component that considers a linear combination of lagged values as the predictors, a moving average (MA) component of past forecast errors (white noise), and an integrated component where differencing is applied on the data to make it stationary. Let $y_1, y_2, \ldots, y_t$ be the demand values over time period t; the time series data can be written as:

$$y_t = f(y_{t-1}, y_{t-2}, y_{t-3}, \ldots, y_{t-n}) + \varepsilon_t \tag{4.1.1}$$

An ARIMA model assumes that the value of demand is a linear function of a number of previous past demand values and previous error values. Thus, the ARIMA model can be written as:

$$\hat{y}_t = \mu + \vartheta_1 y_{t-1} + \vartheta_2 y_{t-2} + \vartheta_3 y_{t-3} + \cdots + \vartheta_p y_{t-p} + \varepsilon_t - \phi_1 \varepsilon_{t-1} - \phi_2 \varepsilon_{t-2} - \phi_3 \varepsilon_{t-3} - \cdots - \phi_q \varepsilon_{t-q} \quad (4.1.2)$$

where $\hat{y}_t$ is the response variable (the predicted demand), μ is a constant, ϑ_i and ϕ_j are model parameters in which $i = 1, 2, \ldots, p$ and $j = 0, 1, 2, \ldots, q$, p and q are the model orders and define the number of autoregressive terms and moving average terms, respectively.

In order to fit an ARIMA model, first the ADF test is applied on the time series data to examine the stationarity, and the standard auto_arima() function in Python is used for hyperparameter tuning and determining the optimal model order. A function is developed in Python to implement the ARIMA model via a rolling-origin strategy.

4.1.1.2 Prophet Model

Prophet is a time series model introduced by Taylor and Letham (2018) that considers common features of business time series: trends, seasonality, holiday effects and outliers. The Prophet model was developed for forecasting events created on Facebook and is implemented as an open source software package in both Python and R. Let g_t be the time series trend function which shows the long-term pattern of data, s_t be the seasonality which captures the periodic fluctuations in data such as weekly, monthly or yearly patterns, and finally h_t be the non-periodic holiday effect. These features are combined through a generalized additive model (GAM) (Hastie and Tibshirani, 1987), and the Prophet time series model can be written as:

$$\hat{y}_t = g_t + s_t + h_t + \varepsilon_t \quad (4.1.3)$$

The normally distributed error ε_t is added to model the residuals. We use the Prophet library in Python for implementing the Prophet model and develop a function for implementation via a rolling-origin strategy.

4.1.2 Multivariate Models

In order to explore the effect of including clinical predictors in the forecasting process, in the next step we introduce three multivariate models that incorporate clinical predictors: lasso regression, random forest, and LSTM networks. These machine learning models are implemented to forecast the demand based on demand history and multiple predictors. Lasso regression is used as a forecasting model and a variable selection method to select the most relevant predictors for the multivariate models.

4.1.2.1 Lasso Regression

We use lasso regression (Tibshirani, 1996) since it allows predictors to be included in the demand forecasting process. Lasso uses an L1 penalty, which tends to push some coefficients towards exactly zero, hence it performs variable selection by reducing the impact of irrelevant or less important predictors. This leads to a reduction in model complexity while improving the prediction accuracy. By considering the actual demand on day t ($t = 1,2,...,N$) as y_t and the predicted demand on day t as the product of the clinical predictors (z_{tj}) and their corresponding coefficients β_j, where $j = 1,2,...,M$ specifies the clinical predictor, the lasso regression model is the solution to the following optimization problem:

$$\arg\min \sum_{t=1}^{N} (y_t - \sum_j \beta_j z_{tj})^2 + \lambda \sum_{j=1}^{M} |\beta_j| \tag{4.1.4}$$

$$\text{subject to } \sum_{j=1}^{M} |\beta_j| \leq \tau. \tag{4.1.5}$$

The optimization problem defined in (4.1.4)-(4.1.5) chooses the coefficients, β, that minimize the sum of squares of the errors between the actual values (y) and the response variable values, with a sparsity penalty (λ) on the sum of the absolute values of the model coefficients. Constraint (4.1.5) forces some of the coefficients (that have a minor contribution to the estimate) to be zero. Predictors that have non-zero coefficients are selected in the model, and the response variable is calculated as follows:

$$\hat{y}_t = \beta z_t \tag{4.1.6}$$

In this study, lasso regression is used as a variable selection method to find important predictors for platelet demand. Subsequently, this information is used for demand forecasting. We use the LassoCV function from the sklearn package in Python to implement the lasso regression. The penalty coefficient λ is chosen through five-fold cross-validation. A function is developed to implement the lasso regression via a rolling-origin strategy.

4.1.2.2 Random Forest

Random forests, first proposed in Ho (1995), are ensemble methods that use decision trees. We chose to explore random forests as they can capture nonlinear relationships between predictors while also being interpretable, as what a decision tree does can be understood by simply looking at it. Decision trees in a forest are trained using bootstrapped samples and are only allowed to consider a subset of the predictors when choosing splits. Considering the actual demand on day t as y_t, and the set of days in the bootstrap samples as D, a tree starts with a root node that has an attached value μ:

$$\mu = \frac{1}{|D|} \sum_{t \in D} y_t \tag{4.1.7}$$

This node creates two child nodes that separate data points based on a clinical predictor, u, where one node gets data with the value of u on day t (z_{tu}) less than a value v and the other node gets data with z_{tu} greater than or equal to v. These child nodes have attached values calculated in the same way as the root, $\mu_1 = \frac{1}{|\{t|z_{tu}<v\}|} \sum_{t:z_{tu}<v} y_t$ and $\mu_2 = \frac{1}{|\{t|z_{tu}\geq v\}|} \sum_{t:z_{tu}\geq v} y_t$.

The split measures, u and v, are chosen by minimizing the variance of the model. A random forest grows a number of these trees, K, and produces a prediction for a set of clinical predictors, z_t, by averaging together the predictions of each of the trees:

$$\hat{y}_t = \sum_{i=1}^{K} T_i(z_t) \tag{4.1.8}$$

where each tree T_i takes a set of clinical predictors and traverses the nodes of tree i using the splits found as described above. Forecasting problems can have linear or nonlinear relationships among the model predictors. Random forests can work on both linear and nonlinear data, and are able to capture nonlinear dependencies among the predictors. We use the RandomForestRegressor function from the scikit-learn package in Python to implement the random forest. Hyperparameter tuning is achieved by using grid search on the number of trees, maximum tree depth, and the number of features to consider when looking for the best split. The best split in a tree is chosen by minimizing MSE (Mean Square Error) and five-fold cross-validation is used to reduce overfitting. We developed a function in Python to implement the random forest model via a rolling-origin strategy.

4.1.2.3 LSTM Network

LSTM networks are a class of Recurrent Neural Networks (RNN) that were introduced by (Hochreiter and Schmidhuber, 1997) and are capable of learning long-term dependencies in sequential data. In theory, RNNs should be capable of learning long-term dependencies, however they suffer from the so-called vanishing gradient problem. Consequently, LSTM networks are designed to resolve this issue. An LSTM network maps a set of input neurons (also called units) to a set of output neurons through a hidden layer. A neuron or unit in an LSTM network consists of an input gate (i_t), a forget gate (f_t), a cell state (c_t), and an output gate (o_t).

The hidden layer output can be written as a function of the gates, the model input (here the clinical predictors (z_t)), and the previous output of the hidden layer:

$$h_t = \sigma_h(i_t, f_t, c_t, o_t, z_t, h_{t-1}) \tag{4.1.9}$$

The output of the LSTM network, here the demand forecasts, is calculated as a weighted sum of the hidden layer outputs plus a bias, b:

$$\hat{y}_t = wh_t + b \tag{4.1.10}$$

Like random forests, LSTM networks are able to capture nonlinear dependencies among the predictors. We implement the LSTM network using the TensorFlow package (Abadi et al., 2016). The LSTM network is trained by using the ADAM optimizer (Kingma and Ba, 2014), and MSE is used as the loss function for this optimizer. For hyperparameter tuning, grid search is performed to find the best model parameters (including the number of epochs, batch size, and number of hidden layers) toward the minimum MSE. Moreover, 10-fold cross-validation is used to reduce overfitting. A function is developed in Python to implement the LSTM network model via a rolling-origin strategy.

4.2 Rolling Window Evaluation

We fit the forecasting models multiple times in order to collect multiple out-of-sample one-step ahead forecast errors by using a rolling window. The rolling window is used as part of the demand forecasting process to periodically retrain the models and use more recent data. The flowchart of the proposed demand forecasting system is given in Figure 4.1. We retrain each model periodically, according to two parameters, the training window and the retraining period. When we retrain a model, we use a training window of the most recent data. For evaluation, we consider a rolling-origin evaluation, similar to the one presented in (Tashman, 2000). Many studies consider a fixed-origin evaluation, but we consider a rolling-origin evaluation to improve the efficiency and reliability of out-of-sample tests (Tashman, 2000). In a rolling-origin evaluation, the forecasting origin is successively updated and new forecasts are produced from each new origin. We set the forecasting window and rolling steps to be the same as the retraining period.

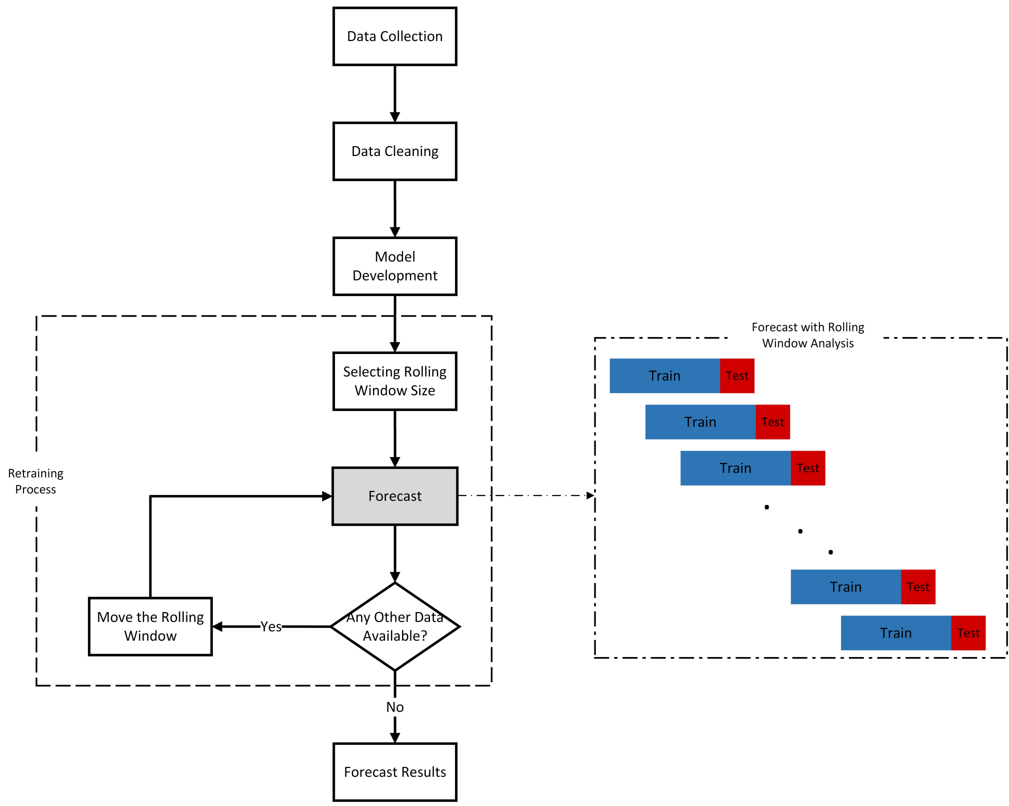

Figure 4.1: Proposed forecasting process

Here we consider two training windows, two years (starting from 2016) and eight years (starting from 2010), to explore the impact of data volume. The forecasting horizon is one year (2018) in which next day forecasts are generated for each of the retraining periods. We consider retraining periods of 1, 7, 30, and 90 days, to examine the trade-off between the accuracy and the overhead of retraining. The forecasting accuracy is computed by averaging the Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Symmetric Mean Absolute Percentage Error (SMAPE) over the forecasting horizon for each rolling origin.

4.3 Results

This section presents demand forecasting comparisons for univariate and multivariate models, and the forecasting performance of the models trained with training window sizes of two and eight years and retraining periods of 1, 7, 30, and 90 days. We implement the models to forecast the daily demand aggregated over four hospitals for one day ahead via a rolling-origin strategy. We consider the aggregated demand as a result of internal inventory management practices. The hospitals work together to share inventory, so we considered all four hospitals as one entity. We periodically retrain our models based on the rolling window analysis explained in Section 4.2.

4.3.1 Demand Forecasting Comparisons for Univariate Models

Figure 4.2 compares the forecasts generated by the univariate models, with a training window of two years and by retraining every day, and the actual demand. The actual demand has a large variance (mean [sd]: 19.28 [7.36]). The ARIMA model's forecasts have significantly lower variance (mean [sd]: 18.89 [3.09]) in comparison to the actual demand, meaning that the forecasts cannot capture the wide range of the actual demand. Despite having a larger variance than the ARIMA model, Prophet shows a similar behavior (mean [sd]: 19.35 [4.40]).

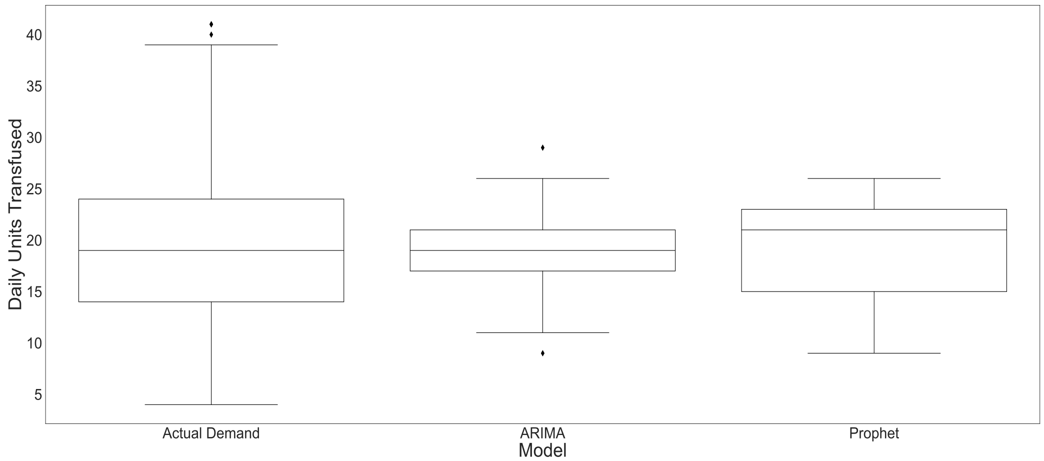

Figure 4.2: Comparison of the actual demand and the predicted demand from univariate models

Next, we examine the univariate models' residuals via the ACF (Autocorrelation Function). Figure 4.3 gives the coefficients of correlation between a value and its lag for ARIMA and Prophet. As we can see in Figure 4.3(a), there is an autocorrelation at time seven (and multiples of seven) due to weekly seasonality that is not incorporated in the model. Since seasonality is one of the primary features of our time series data, we include seasonality directly in the forecasting process by using the Prophet model. As we can see in Figure 4.3(b), there is no repeated autocorrelation pattern for Prophet residuals.

We also perform a pairwise t-test to compare the univariate models' residuals. The results show a statistically significant difference between the ARIMA residuals (mean [sd]: 0.39 [6.80]) and Prophet residuals (mean [sd]: -0.07 [5.90], t-test: 95% confidence interval for the difference in means: (0.08, 0.85), *P* value = 0.018), indicating higher residuals in the ARIMA model.

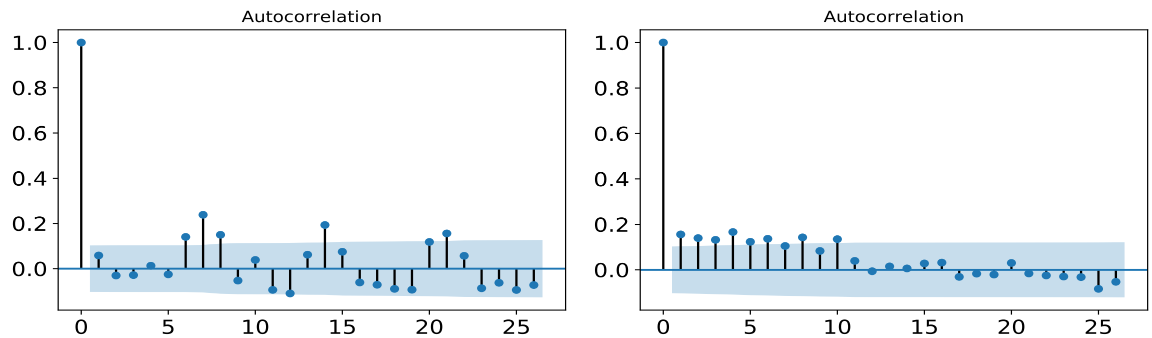

(a) ACF for ARIMA residuals (b) ACF for Prophet residuals

Figure 4.3: ACF plots for ARIMA and Prophet residuals with a training window of two years and by retraining every day

4.3.2 Demand Forecasting Comparisons for Multivariate Models

We begin this section with an examination of selecting the clinical predictors for the multivariate models. Next, we compare the forecasts generated by the multivariate models and the actual demand.

4.3.2.1 Selecting the Predictors Using Lasso Regression

As discussed in Section 3.1, the data has more than 100 features, and we select predictors via lasso regression. The 29 clinical predictors that are introduced in Section 3.1 are selected by lasso regression and used for training the multivariate models. More specifically, we consider clinical indicators, consisting of laboratory test results, patient characteristics and hospital census data as well as operational related indicators, including the previous week's platelet usage and previous day's received units with the aim of accurate demand forecasting.

We calculate the confidence intervals for these clinical predictors (also referred to as the model predictors). There are multiple methods for calculating a confidence interval for the predictors; one of the most popular is the bootstrap method (Efron and Tibshirani, 1994). The bootstrap method is used in the experiments for calculating the confidence intervals for the predictors used in the multivariate models. As shown in Figure 4.4, the predictors' coefficients have a wide range, so we see high values (abnormal_plt = 0.23) as well as low values (Friday = -0.39) for the lab tests and day of the week. Overall, we can see that the range of the predictors' coefficients for the 95% confidence interval is narrow. Detailed information about the predictors and their corresponding coefficients are given in Table B.1 of Appendix B.

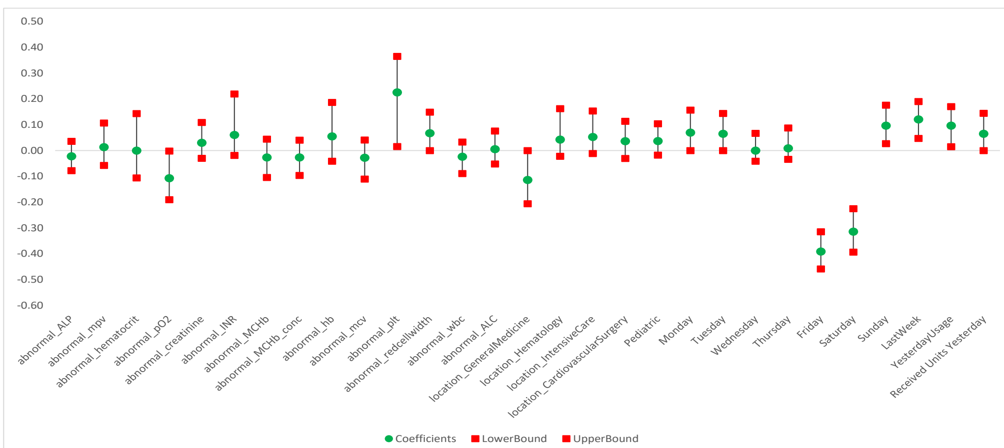

Figure 4.4: Confidence interval for predictors' coefficients - Lasso regression

4.3.2.2 Comparisons of Multivariate Models Forecasts

Figure 4.5 shows the actual daily units transfused and the forecasts generated by the multivariate models, lasso regression, random forest and LSTM network, with a training window of two years and by retraining every day. The forecast means of lasso regression (mean [sd]: 19.12 [3.62]) and random forest (mean [sd]: 19.72 [4.28]) are very close to the actual mean demand, but forecast standard deviations are much lower than the actual demand standard variation. LSTM network forecasts have a slightly lower mean (mean [sd]: 18.01 [3.55]) but significantly lower standard deviation than the actual demand.

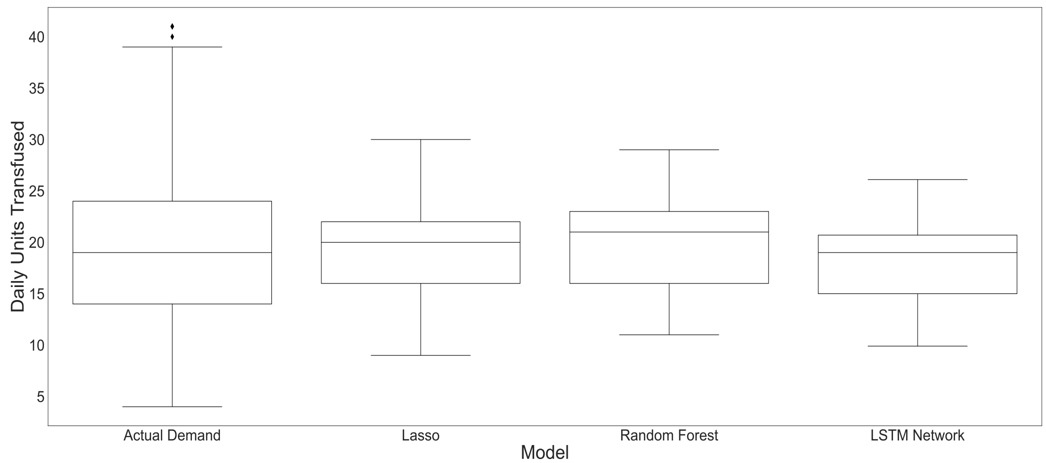

Figure 4.5: Comparison of the actual demand and the predicted demand from multivariate models

Next, a repeated measures ANOVA test is performed for comparing the multivariate models' residuals with each other. The results of the test show a statistically significant difference between the lasso regression, random forest, and LSTM network residuals (F = 35.86, P value <0.001). To show which models' residuals are significantly different, we perform pairwise comparisons by using a pairwise t-test. Table 4.1 gives the results of the pairwise t-test for the models' residuals, showing that they are significantly different from each other. The P values are adjusted using the Bonferroni multiple testing correction method.

Table 4.1: Comparison of multivariate models residuals using a pairwise t-test

Model	**Descriptive statistics**		**T-test**		
	Mean	**Standard Deviation**	**Model**	**95% confidence interval for the difference in means**	***P* value**
Lasso Regression	0.16	6.39	Random Forest	(0.09, 1.12)	0.020
Random Forest	-0.44	8.77	LSTM Network	(1.60, 1.83)	<0.001
LSTM Network	1.27	8.34	Lasso Regression	(-1.57, -0.65)	<0.001

4.3.3 Performance Comparisons

The performance of the forecasting models is computed based on a rolling-origin evaluation and by four error measures, RMSE, MAE, MAPE, and SMAPE. The first two error measures, RMSE and MAE, are absolute measures while the remaining ones, MAPE and SMAPE, are relative measures. The errors are measured for each rolling origin for the test data and reported in Figures 4.6-4.9 and Table 4.2. Table 4.2 gives the mean and standard deviation of the errors for different training window sizes and retraining periods.

Figures 4.6 and 4.7 compare the RMSE and MAE of the models trained with different training window sizes and retraining periods. As we can see in these figures and in Table 4.2, increasing the size of the training window mostly affects the univariate models, ARIMA and Prophet. ARIMA's performance improves when moving from two years to eight years of data. Since ARIMA's forecasts are only based on the previous demands, and the seasonality in data has not changed significantly during the eight years, the model parameters, p and q, are more robust for longer time series data (including 5 lagged values and a moving average of 2), resulting in more accurate forecasts. In general, when a limited amount of data are available, the ARIMA model has a high forecast error not only because its forecasts are solely based on the previous demands, but also due to the fact that it cannot capture the seasonality in the data. Prophet's accuracy is also improved as the amount of data increases. However, unlike ARIMA, the forecast errors are similar for different retraining periods. The results for lasso regression and the LSTM network indicate that there is not much difference for these methods when there is a large amount of data for training, or when different retraining periods are considered. Random forest does see a slight improvement with eight years of data, and it is the only multivariate model to see this improvement. Its forecast errors are very close for different retraining periods.

In terms of the retraining periods, retraining the models less frequently reduces the variability of the error. If we compare Figure 4.6(a) with Figure 4.6(g), we see that the RMSE error is less variable in Figure 4.6(g) for all the models, similarly for MAE in Figure 4.7. This can also be verified from the results in Table 4.2, where we see lower standard deviations as we move down to retraining every 90 days.

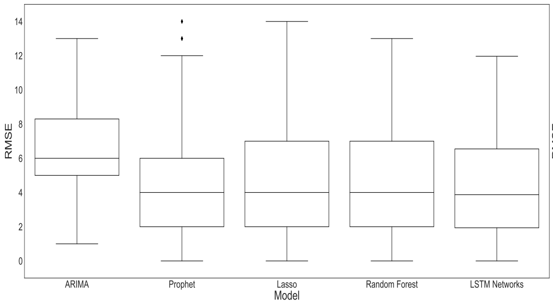

(a) 2 years rolling window, retraining every day

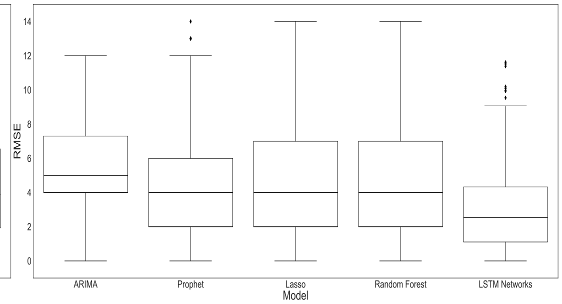

(b) 8 years rolling window, retraining every day

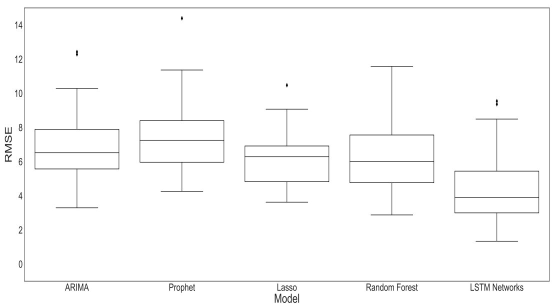

(c) 2 years rolling window, retraining every 7 days

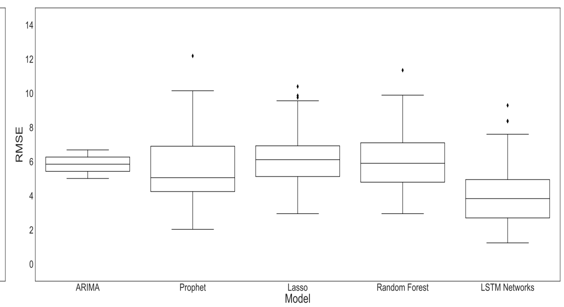

(d) 8 years rolling window, retraining every 7 days

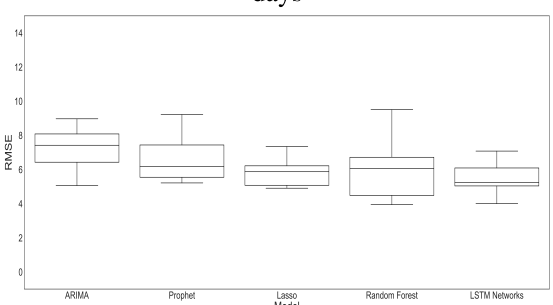

(e) 2 years rolling window, retraining every 30 days

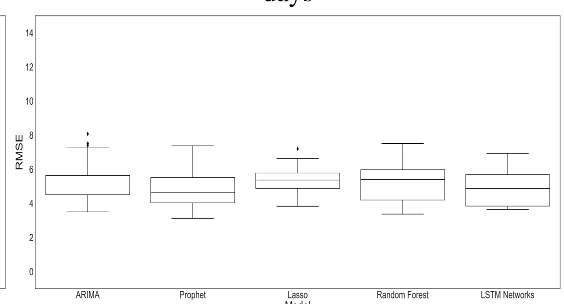

(f) 8 years rolling window, retraining every 30 days

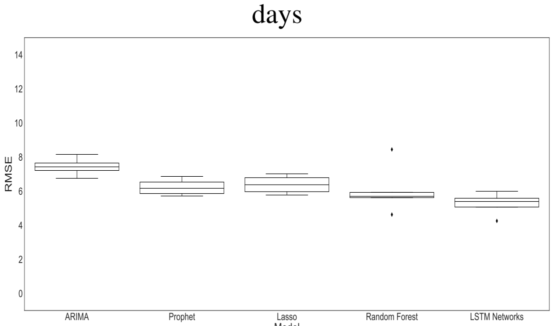

(g) 2 years rolling window, retraining every 90 days

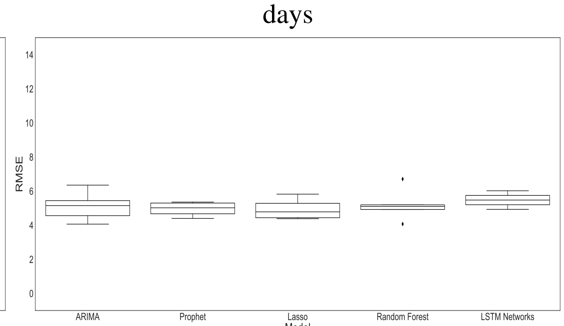

(h) 8 years rolling window, retraining every 90 days

Figure 4.6: RMSE with different training window sizes and retraining periods

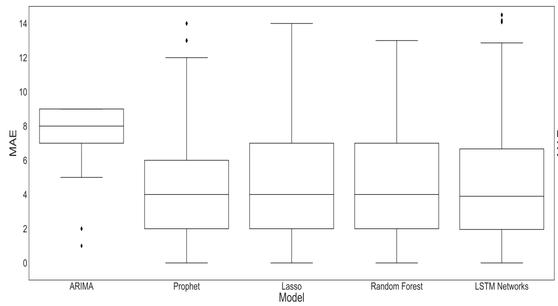

(a) 2 years rolling window, retraining every day

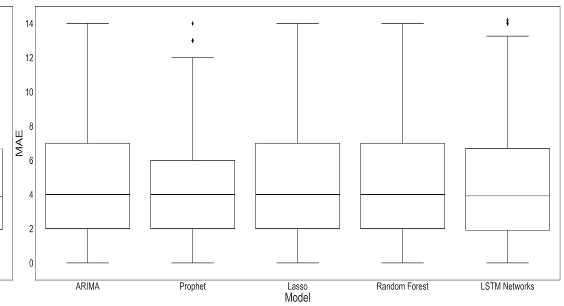

(b) 8 years rolling window, retraining every day

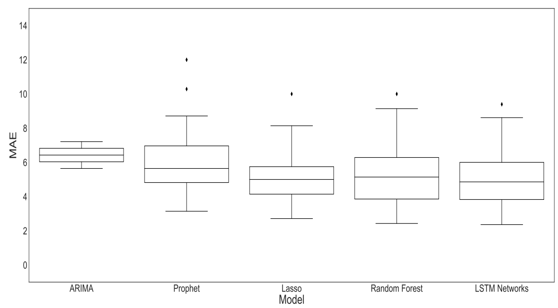

(c) 2 years rolling window, retraining every 7 days

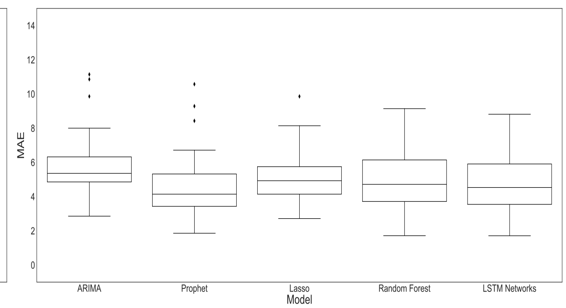

(d) 8 years rolling window, retraining every 7 days

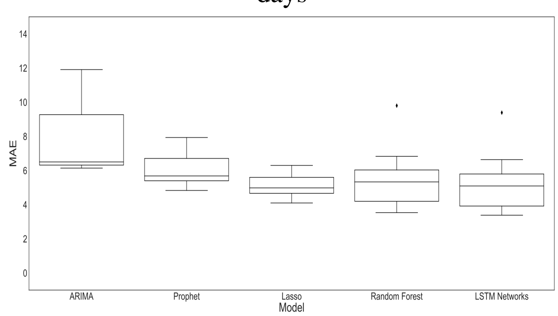

(e) 2 years rolling window, retraining every 30 days

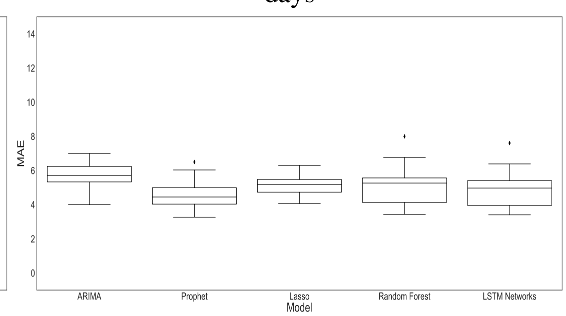

(f) 8 years rolling window, retraining every 30 days

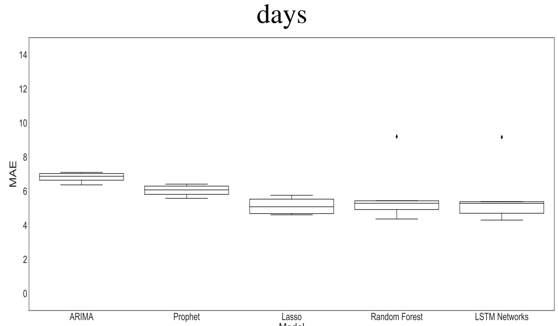

(g) 2 years rolling window, retraining every 90 days

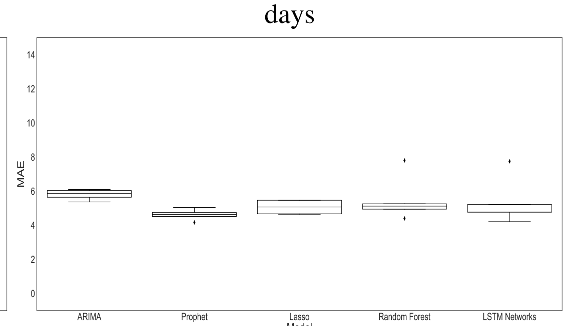

(h) 8 years rolling window, retraining every 90 days

Figure 4.7: MAE with different training window sizes and retraining periods

Figures 4.8 and 4.9 compare the MAPE and SMAPE of the models trained with different training window sizes and retraining periods. As we can see from these figures and from Table 4.2, increasing the training window size does not necessarily decrease the errors. There is similar behavior for the retraining periods, especially for the multivariate models, but we see that retraining less frequently results in less variable errors.

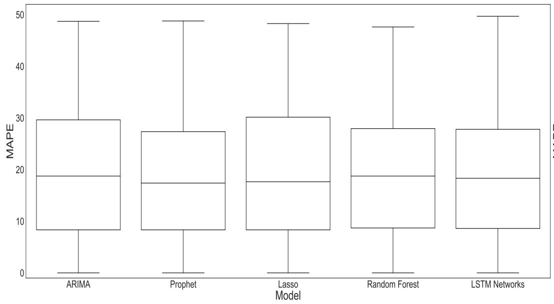

(a) 2 years rolling window, retraining every day

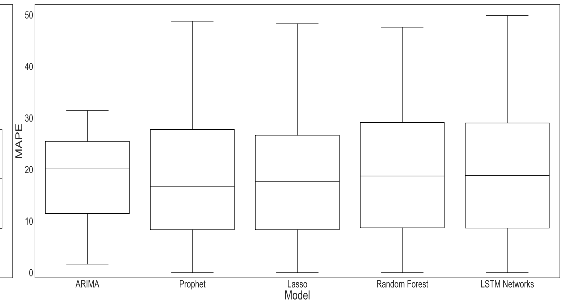

(b) 8 years rolling window, retraining every day

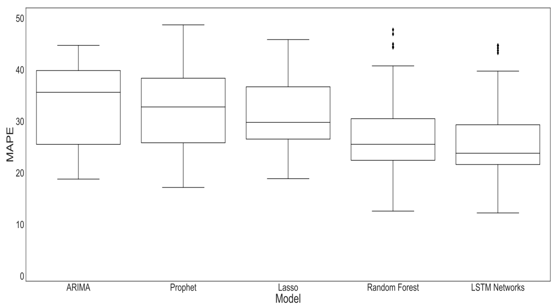

(c) 2 years rolling window, retraining every 7 days

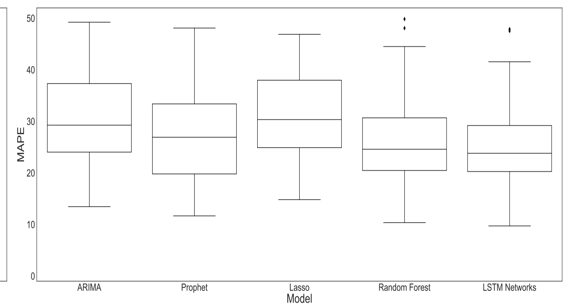

(d) 8 years rolling window, retraining every 7 days

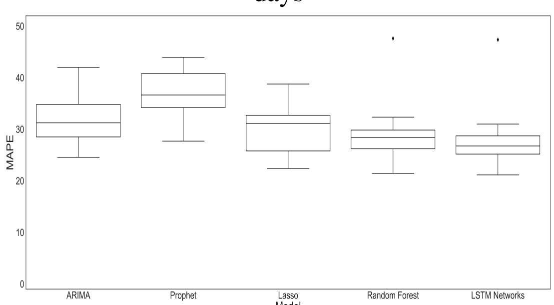

(e) 2 years rolling window, retraining every 30 days

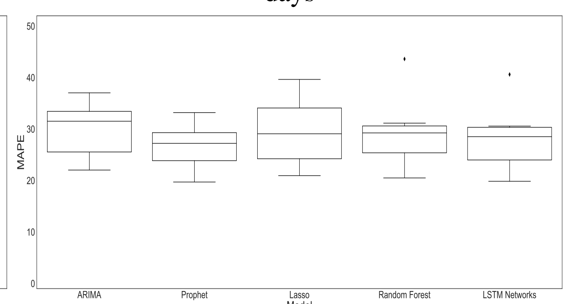

(f) 8 years rolling window, retraining every 30 days

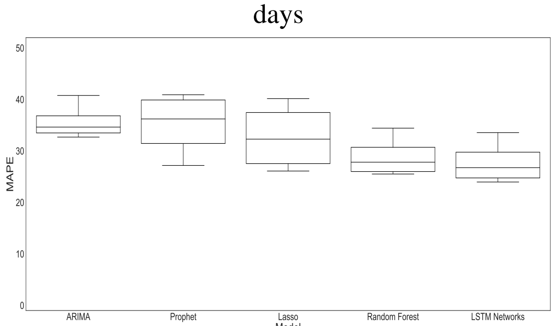

(g) 2 years rolling window, retraining every 90 days

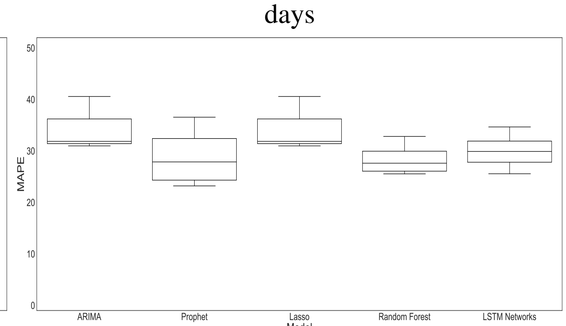

(h) 8 years rolling window, retraining every 90 days

Figure 4.8: MAPE with different training window sizes and retraining periods

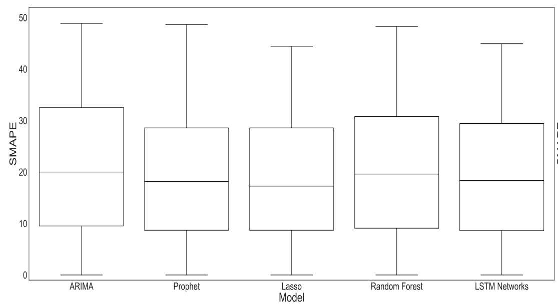

(a) 2 years rolling window, retraining every day

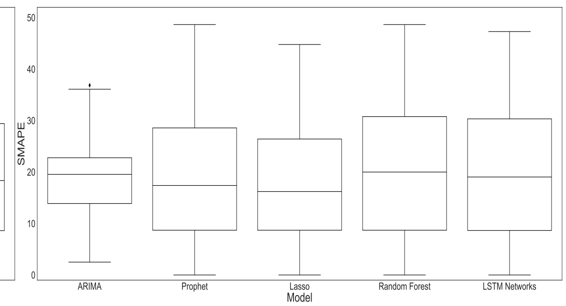

(b) 8 years rolling window, retraining every day

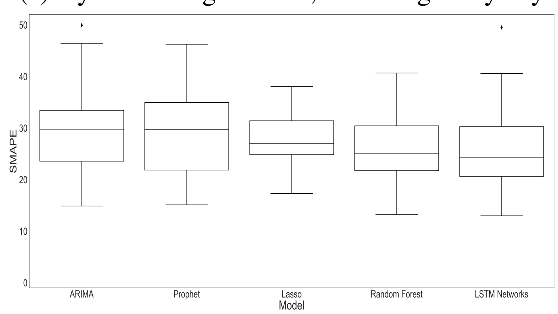

(c) 2 years rolling window, retraining every 7 days

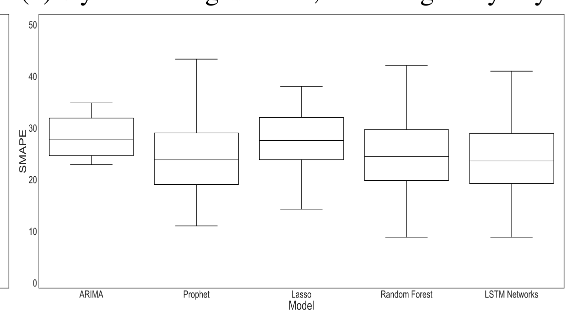

(d) 8 years rolling window, retraining every 7 days

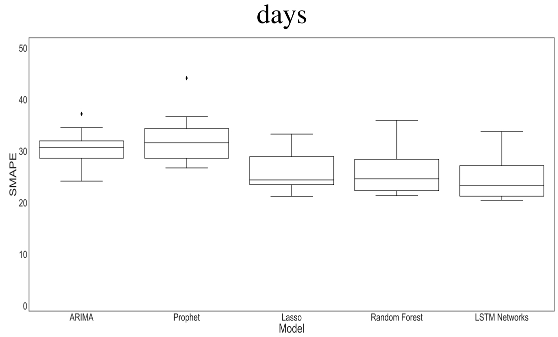

(e) 2 years rolling window, retraining every 30 days

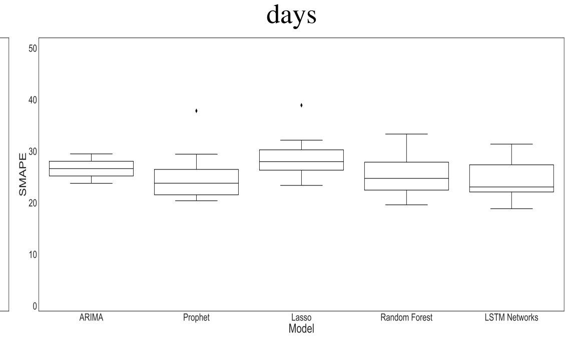

(f) 8 years rolling window, retraining every 30 days

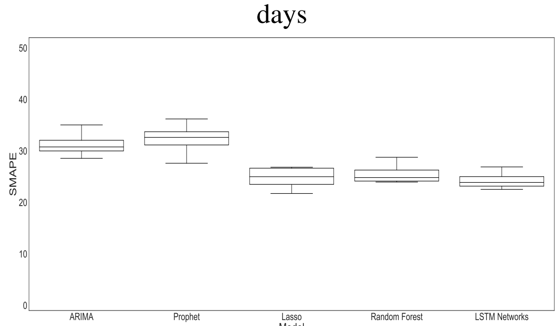

(g) 2 years rolling window, retraining every 90 days

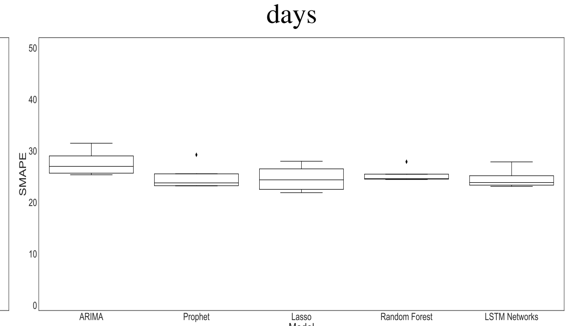

(h) 8 years rolling window, retraining every 90 days

Figure 4.9: SMAPE with different training window sizes and retraining periods

Overall, the results indicate that while univariate models can benefit from a larger training window size and frequent retraining, the performance of the multivariate models is not affected by a larger training window, meaning that these models have robust performance with different data volumes.

Table 4.2: Model performance with different training window sizes and retraining periods

		Two years				Eight years			
Retraining Period	**Model**	**RMSE**	**MAE**	**MAPE**	**SMAPE**	**RMSE**	**MAE**	**MAPE**	**SMAPE**
1 day	ARIMA	6.7±2.76	7.37±1.75	19.46±12.55	21.37±14.01	5.72±2.76	4.84±3.43	18.77±8.09	18.79±6.68
	Prophet	4.20±3.19	4.20±3.19	18.52±12.85	19.05±13.12	4.28±3.12	4.28±3.12	18.99±12.43	19.72±13.04
	Lasso Regression	4.88±3.40	4.88±3.40	19.66±12.86	18.92±12.25	4.80±3.39	4.80±3.39	18.96±12.43	18.08±11.63
	Random Forest	4.59±3.24	4.59±3.24	19.49±12.58	20.23±13.05	4.59±3.27	4.59±3.27	19.57±12.95	20.38±13.32
	LSTM Networks	4.31±2.99	4.54±3.28	19.4±12.69	19.23±12.36	3.09±2.49	4.49±3.24	19.84±13.37	19.77±12.94
7 days	ARIMA	6.81±2.09	6.43±0.79	33.19±8.33	30.0±7.74	5.85±0.84	5.56±1.75	31.14±8.75	28.26±3.86
	Prophet	7.34±2.01	5.95±1.74	32.32±8.36	28.65±8.25	5.6±1.94	4.61±1.69	27.22±9.09	24.32±6.76
	Lasso Regression	6.22±1.59	5.15±1.42	31.7±7.50	27.7±4.94	6.31±1.80	5.16±1.49	31.64±8.37	27.28±5.55
	Random Forest	6.3±2.00	5.17±1.68	27.48±8.70	25.64±6.62	6.12±1.94	5.01±1.63	26.8±8.78	25.43±7.33
	LSTM Networks	4.32±2.01	4.97±1.60	26.45±8.39	25.83±7.25	4.07±1.82	4.87±1.61	25.83±8.43	24.53±7.08
30 days	ARIMA	7.18±1.22	7.96±2.34	32.49±5.34	30.46±3.43	4.86±0.94	5.68 ± 0.88	30.06±4.90	26.63±2.86
	Prophet	6.62±1.23	6.01±0.87	36.92±5.06	32.38±4.80	4.85±1.20	4.62±0.95	26.42±4.05	25.01±4.78
	Lasso Regression	5.86±0.74	5.12±0.62	30.2±5.14	25.79±3.48	5.44±0.89	5.11±0.62	29.4±6.28	28.73±3.66
	Random Forest	6.13±1.61	5.42±1.60	29.09±6.41	25.93±4.33	5.41±1.36	5.19±1.22	28.66±5.70	25.52±3.87
	LSTM Networks	5.41±0.81	5.2±1.54	28.06±6.51	24.68±3.98	4.99±1.18	4.96±1.14	27.66±5.29	24.37±3.58
90 days	ARIMA	7.44±0.50	6.8±0.29	35.68±3.10	31.29±2.36	5.08±0.72	5.8±0.29	34.47±4.33	27.71±2.44
	Prophet	6.23±0.45	6.03±0.32	35.13±5.51	32.27±3.07	4.95±0.39	4.62±0.31	28.87±5.32	25.01±2.49
	Lasso Regression	6.38±0.50	5.13±0.49	32.71±5.83	24.82±1.81	4.94±0.58	5.06±0.40	34.47±4.33	24.66±2.46
	Random Forest	6.06±1.27	5.84±1.72	28.89±3.53	25.59±1.92	5.2±0.85	5.5±1.18	28.4±2.86	25.39±1.45
	LSTM Networks	5.26±0.63	5.77±1.74	27.76±3.74	24.29±1.64	5.48±0.55	5.34±1.24	29.95±3.03	24.68±1.89

4.4 Comparison and Discussion

In this section, we compare the models and provide recommendations for using these models in various scenarios. In Section 4.4.1 we compare the models based on a training window of two years, in Section 4.4.2 we discuss the impact of an increased amount of data on the forecasting models, and in Section 4.4.3 we discuss the effect of different retraining periods on the models. Finally, in Section 4.4.4 we provide the overall methodological implications of the study and in Section 4.4.5 discuss managerial implications.

4.4.1 Univariate versus Multivariate Models

We have presented five different models for platelet demand forecasting that can be divided into two groups: univariate and multivariate. Univariate models, ARIMA and Prophet, forecast future demand based only on the demand history. Although the ARIMA model only considers a limited number of previous values for forecasting the demand, retraining it every day, week or month leads to a slight performance improvement. The Prophet model incorporates the historical data, seasonality and holiday effects into the demand forecasting model which results in an improvement in the forecasting accuracy by approximately 10% compared to ARIMA. This highlights the impact of weekday/weekend and holiday effects in the platelet demand variation. As we discussed in Chapter 3, there is a weekday/weekend effect for platelet demand, which is not (directly) captured in the ARIMA model.

Multivariate models, on the other hand, incorporate clinical predictors as well as historical demand data for demand forecasting. We use lasso regression to select the dominant clinical predictors that affect the demand. Lasso regression examines the linear relationship among the clinical predictors and their influence on the demand. However, as presented in

Figure 3.5, there are correlations among the clinical predictors. There may also be nonlinear relationships among these clinical predictors that cannot be captured by a linear regression model. These issues motivated us to use two machine learning approaches, random forest and LSTM network. Random forest is capable of capturing nonlinearities among variables and its forecasting method of averaging past values provides some contrast to the LSTM network's modelling approach. An LSTM network can also account for nonlinearities among variables. Moreover, an LSTM network is capable of retaining past information while forgetting some parts of the historical data. As we can see from Table 4.2, in general, random forest, LSTM network and lasso regression have low forecast errors for different training window sizes, owing to the inclusion of the clinical predictors.

4.4.2 Two Years versus Eight Years of Data

As discussed in Section 4.1, we train our models for two training window sizes, with two years and eight years of data, respectively. Since there is no trend in the data from 2016 onwards (see Figure 3.2), in the first scenario the models are trained for two years (training window size of two years, starting from 2016). With this amount of data and by retraining every year, forecasts are not accurate for univariate time series approaches, and one needs to include the clinical predictors in the forecasting model. However, by considering a training window of eight years, the ARIMA model's performance improves by approximately 20%, compared to the case of a two year training window. The Prophet model's performance also improves when more data are available, specifically when it is trained less frequently (30 and 90 day retraining windows).

In general the multivariate models result in small forecasting errors for two years of data for training, and do not perform significantly better as the amount of data increases, which shows that there is not much sensitivity to the training window size. This highlights the importance of including the clinical predictors in the forecasting process.

4.4.3 Different Retraining Periods

We also compare different retraining periods and provide insight on how to choose the appropriate retraining period for this data (and in general). Our results show that considering different retraining periods does not affect the models in the same manner. While in general all the models benefit from retraining more frequently, univariate models benefit more. For the univariate models, the greatest performance increase is for the ARIMA model when retrained every day, resulting in a decrease of 50% in MAPE and SMAPE. For the multivariate models, lasso regression has an impressive performance increase when retrained every day, while random forest and LSTM networks show less sensitivity to the retraining period. So, by considering the overhead of retraining these models more frequently, one may decide to use a larger retraining window for random forest and LSTM networks.

Generally, if the retraining period is small, meaning that the models are retrained more frequently, the mean forecast accuracy representing the long-term overall performance is improved.

4.4.4 Methodological Implications of the Study

In general, when there is access only to previous demand values, using a univariate model and retraining it frequently is effective. In practice, applying forecasting models to real-world healthcare systems can be challenging. One common concern is data accessibility, especially in small or rural hospitals. In such healthcare facilities, laboratory test results may not be available or may be limited. In these systems, we recommend employing a simple model such as a univariate time series that forecast the demand based on the historical demand of platelets. Results show that using a simple univariate time series model can yield results comparable to those of more complicated models. Moreover, in resource-limited settings where access to data scientists and statisticians is limited, a straightforward univariate model like ARIMA or Prophet can be employed. These models do not demand extensive expertise for training and usage, making them more accessible and suitable for such environments. Notably, the World Health Organization (WHO) suggests that alongside the requirement for robust blood forecasting models, there is also a growing need for simple models that can be employed in all settings. Univariate models align well with these requirements (Organization et al., 2010).

In the case that several data variables are available, lasso regression, random forest models, and LSTM networks can forecast the demand with higher accuracy even when a small amount of data is available and without frequent retraining. Forecasting problems can have linear or nonlinear relationships among the model variables. Due to the fact that LSTM networks are appropriate for both linear and nonlinear time series, and are able to capture nonlinear dependencies, they can outperform linear regression models when long term correlations exist in the time series. Based on the LSTM results, we conclude that long term correlations and nonlinearity are not major issues for our data since the LSTM model does not significantly outperform lasso regression.

While LSTM networks perform well even with a limited amount of data and they can capture nonlinear relationships, they lack interpretability. Interpretability is an important feature of any prediction model used in a safety critical setting like blood product distribution. Considering the time and memory complexity, and interpretability of these models, lasso regression has lower time and memory complexity while it is also very interpretable.

Random forest models maintain interpretability while also having the ability to capture nonlinear relationships. Random forests do well when their training data has good coverage of the different feature combinations the model is forecasting. This is because random forest models make forecasts for a set of features by averaging together similar data points from the training data. This allows random forest models to extract nonlinear relationships but also means they cannot extract trends effectively and may need a large amount of data in order to work well. This can be seen in our model (see Table 4.2), a training window of eight years, with more training data points to reference, has a small improvement in the error measures over a training window of two years for different retraining periods.

Training random forest models and LSTM networks requires expertise in the machine learning area since poor training will cause low-precision results. It is also worth mentioning that the LSTM network is a robust learning model and is capable of learning linear and nonlinear relationships among the model variables even in very short time series data (Boulmaiz et al., 2020; Lipton et al., 2015). However, as the number of inputs increases, both the data variables that make data wide and the data rows that make data tall, LSTM performance tends to decrease because it is highly dependent on the input size. Moreover, wide data results in model overfitting (Lai et al., 2018). Having wide data, one can apply a feature selection method such as lasso regression to reduce the number of variables and regularize the input.

One limitation of forecasting models is that they cannot capture sharp peaks in demand. Figure 4.10 depicts the actual and predicted demands for the second half of 2018 with a training window of two years and a retraining period of 7 days (retraining weekly) using lasso regression. It appears that the model does some degree of smoothing and thus cannot detect the sharp peaks. One possible explanation is that regression models are regressed on the expectation of the outcome, and are not good at capturing the extreme deviations from this expectation. However, as shown in Figure 4.10, smoothing mostly occurs for the maxima rather than the minima. In other words, the model potentially has large errors when there is excess demand, for example in emergency situations. The results presented in Sections 4.3.1 and 4.3.2 support that all the models struggle with capturing the peaks in demand.

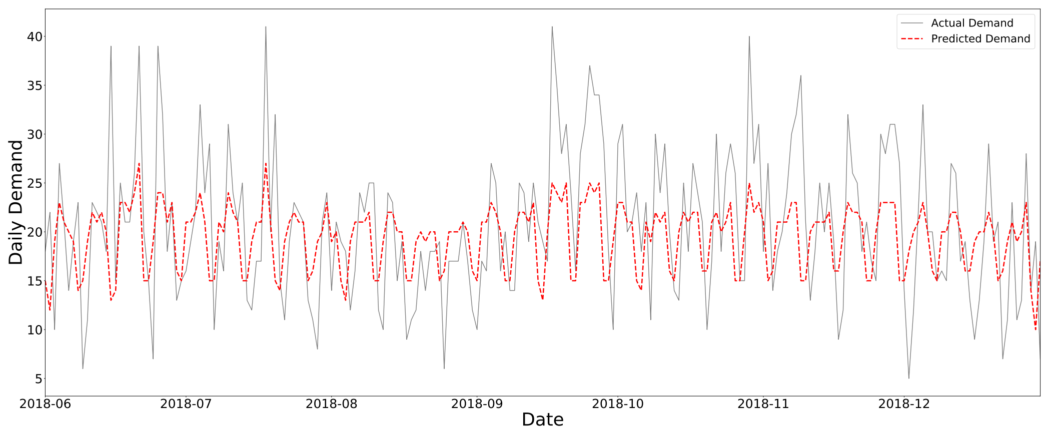

Figure 4.10: Demand forecasting with lasso regression with a training window of two years and a retraining period of seven days

To sum up, when a sufficient amount of data is available, using a univariate model results in a low forecast error, particularly in the case that it is retrained every day. Specifically, when there is only access to the previous demand (as is currently the case for CBS) and adequate historical data are available, one can benefit from a simple univariate model like ARIMA or Prophet, since univariate models are simpler than the multivariate models. Multivariate models are useful when there is access to a limited amount of data. Also, they do not necessarily require frequent retraining, which may be an important implementation concern.

4.4.5 Managerial Implications of the Study

The short shelf life of platelets results in wastages which not only incur large costs but also affect the environment since they cannot be reused, recycled, or recovered (Jemai et al., 2020). Moreover, since platelet demand is highly variable, urgent same-day deliveries are placed frequently. Apart from the high cost of urgent orders, platelet shortage can increase the risk of putting patients' lives in danger. Currently, blood suppliers are not aware of the demand at the hospitals since hospitals hold excess inventory to manage the highly variable platelet demand. Indeed, this has its roots in the bullwhip effect, that is hospitals tend to order more than their actual demand. We see an opportunity to better coordinate supply (number of units received) with demand (number of units transfused) through the development of a daily demand predictor. Forecasting the demand improves the transparency between blood suppliers and hospitals, and helps blood suppliers to make better-informed decisions.

From the clinical perspective, accurate demand forecasting is important for clinical and supply chain management purposes. Demand forecasting can be used for placing optimal platelet orders and for decision making in many parts of the supply chain such as donation planning, and resource and staff management. As we can see in Section 4.3, there appears to be a limit to how accurate the demand forecasts can be, so one important challenge would be how to use the demand forecasts to inform an ordering policy in an effective manner. Clearly, forecasts themselves do not reflect an optimal ordering decision but they can be used as additional information in building effective ordering/inventory management policies, which is the focus of the remainder of this dissertation.

Moreover, this research provides a holistic analysis of the predictors that affect the platelet demand, including the clinical predictors, hospital locations, day of the week and demand history. This can help blood suppliers with adapting clinically relevant factors into the decision making process, like decisions regarding the assignment of transfusion-related staff/resources (beds or equipment).

Overall, there is a significant caveat with all of these approaches in that there are still forecasting errors, in particular they all struggle with capturing peaks. These underestimations may cause significant concerns for using such forecasts directly as there is the danger of severe underestimation. Therefore, some adjustments may be required for using these forecasts according to specific objectives. In the next chapter, we propose an optimization model to incorporate demand forecasts in the inventory model.

4.5 Conclusion

In this chapter, we utilized two types of methods for platelet demand forecasting, univariate and multivariate methods. Univariate methods, ARIMA and Prophet, forecast platelet demand only by considering the historical demand information, while multivariate methods, lasso regression, random forest and LSTM networks, also consider clinical predictors. The error levels for the univariate models, particularly in the case that a small amount of data is available, motivate us to utilize clinical predictors to investigate their ability to improve the accuracy of forecasts. Results show that lasso regression, random forest and LSTM networks outperform the univariate methods when a limited amount of data are available. Moreover, since they include clinical predictors in the forecasting process, their results can aid in building a robust decision making and blood utilization system. However, their application is not limited to platelet products. We believe that they can be used in various areas when data features are available, including healthcare in general, finance and climate studies. On the other hand, when there is access to a sufficient amount of data, the marginal improvement for a simple univariate model such as ARIMA is higher than for multivariate models. In such scenarios, univariate models can be applied to historical data for demand forecasting, regardless of the product, which makes these models generalizable and widely applicable.

Chapter 5

Inventory Management of Perishable Products by Incorporating Demand Forecasts

In this chapter, we analyze the optimal ordering policy in a periodic-review single-item inventory system when demand forecasts are included in the inventory model. Demand forecasts are considered as a function of previous forecast value. We provide a mathematical characterization of the optimal ordering policy in terms of the structural properties under such assumptions. Also, we propose a heuristic as an alternate for the optimal policy. Our experiments suggest that by including forecasts in the inventory model, we can keep the cost at a very low level while keeping the on-hand inventory close to the actual demand.

We start this chapter with a formal description of the inventory problem. In Section 5.2, structural properties of the optimal ordering policy are explored. Next, we explain our proposed heuristic in Section 5.3. Finally, Section 5.4 concludes the chapter.

5.1 Problem Description and Model Formulation

We consider a periodic-review single-item inventory system over a finite horizon of T periods. We consider both lost-sales and backlogging cases. The product is perishable with a fixed shelf life of R periods. At the beginning of each period t, an order is placed and is received immediately (the lead time is zero). The per unit ordering cost is denoted by c. The inventory is described by a vector of size R, $\mathbf{x}_t = [x_t^1, x_t^2, ..., x_t^R]$, in which x_t^r is the number of items in inventory with a remaining shelf life of r periods. We assume that the new arrivals, q_t, are fresh products with a remaining shelf life of R periods. Next, demand d_t occurs and is satisfied based on the FIFO policy. This is a common assumption in the literature, specifically when the suppliers control the inventory issuing policy (Chen et al., 2014). Demand takes an additive form, which is also a common assumption in the literature (Chen and Simchi-Levi, 2004; Chen et al., 2014):

$$d_t = \hat{f}_t + e(\hat{f}_t) \tag{5.1.1}$$

where $\hat{f}_t$ is the predicted demand in period t and $e(\hat{f}_t)$ is the forecast error, whose distribution is dependent on $\hat{f}_t$. If there are any shortages, a penalty cost, p per unit of shortage, is incurred. At the end of the period, the expired units are disposed of and considered as wastage with a cost of w per each wasted unit. Finally, the inventory, $\mathbf{x}_t$, is updated based on the consumed units, and a holding cost, h per unit, is incurred based on the items remaining in the inventory. Figure 5.1 is a representation of the problem.

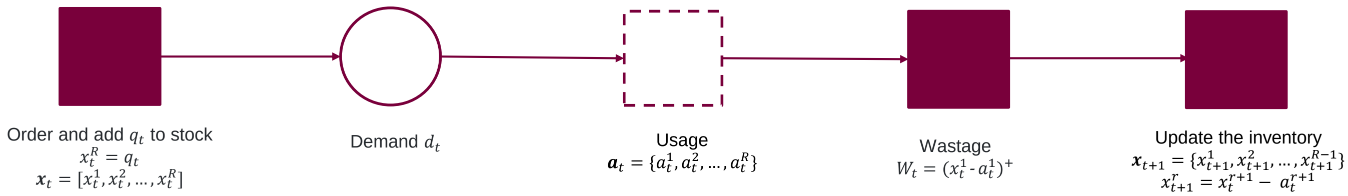

Figure 5.1: Inventory update process

The predicted demand, $\hat{f}_t$, is calculated as a function of $\hat{f}_{t-1}$, the previous period's predicted demand, plus an error ε_t as follows:

$$\hat{f}_t = \theta \hat{f}_{t-1} + \varepsilon_t \tag{5.1.2}$$

The errors ε_t are independently and identically distributed. We will also discuss a more general forecast function in which the forecast value is a more complex function of previous demand forecast values and demand values.

5.1.1 The Lost-Sales Case

In the lost-sales case, unmet demand is lost. To implement the FIFO policy, we follow a similar approach to Chen et al. (2014) and consider inventory as a vector $\mathbf{s}_t = [s_t^1, s_t^2, \dots, s_t^{R-1}]$, where

$$s_t^1 = x_t^1, \qquad s_t^2 = s_t^1 + x_t^2, \qquad \dots, \qquad s_t^{R-1} = s_t^{R-2} + x_t^{R-1}, \tag{5.1.3}$$

and $s_t^R = s_t^{R-1} + q_t$. Finally, the next state is:

$$\mathbf{s}_{t+1} = [(s_t^2 - d_t)^+, (s_t^3 - d_t)^+, \dots, (s_t^R - d_t)^+] \tag{5.1.4}$$

Now, using the relation $d_t = \hat{f}_t + e(\hat{f}_t)$, we can write the expected cost for period t as:

$$\mathbb{E}[c(s_t^R - s_t^{R-1}) + h(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^+ + w(s_t^1 - (\hat{f}_t + e(\hat{f}_t)))^+ + p(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^-] \tag{5.1.5}$$

The objective is to find an optimal ordering policy to minimize the total expected discounted cost over a finite planning horizon:

$$\min_{s_t^R \geq s_t^{R-1}} \left\{ \begin{aligned} &\sum_{t=1}^{T} \gamma^{t-1} \mathbb{E}[c(s_t^R - s_t^{R-1}) + h(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^+ \\ &+ w(s_t^1 - (\hat{f}_t + e(\hat{f}_t)))^+ + p(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^-] \end{aligned} \right\} \tag{5.1.6}$$

subject to:

$$s_{t+1}^r = (s_t^{r+1} - d_t)^+ \qquad \forall\, t \geq 1, 1 \leq r \leq R-1 \tag{5.1.7}$$

$$d_t = \hat{f}_t + e(\hat{f}_t) \tag{5.1.8}$$

where $\gamma \in [0,1]$ is the discount factor. The optimal cost from day t onward, $\hat{u}_t(\mathbf{s_t}, \hat{f}_t, \hat{f}_{t-1})$, can be written as follows:

$$\hat{u}_t(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1}) = \min_{s_t^R \geq s_t^{R-1}} \{\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1})\} \tag{5.1.9}$$

where

$$\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1}) = \mathbb{E}\begin{bmatrix} c(s_t^R - s_t^{R-1}) + h(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^+ + w(s_t^1 - (\hat{f}_t + e(\hat{f}_t)))^+ \\ + p(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^- + \gamma \mathbb{E}[\hat{u}_{t+1}(\mathbf{s}_{t+1}, \hat{f}_{t+1}, \hat{f}_t)] \end{bmatrix} \tag{5.1.10}$$

5.1.2 The Backlogging Case

In this section, we analyze the backlogging case where shortages are backordered. In the backlogging case, the on-hand inventory can be a negative value; $s_t^r \leq 0$ denotes that there is a backlog of $|s_t^r|$ items with shelf life of at most r periods. The next state is expressed as:

$$\mathbf{s}_{t+1} = [s_t^2 - d_t, s_t^3 - d_t, \ldots, s_t^R - d_t] \tag{5.1.11}$$

We can rewrite the expected cost for period t as:

$$\mathbb{E}[c(s_t^R - s_t^{R-1}) + h(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^+ + w(s_t^1 - (\hat{f}_t + e(\hat{f}_t)))^+ + b(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^-] \tag{5.1.12}$$

where b is the backorder cost per item. The assumption $c \leq b/(1-\gamma)$ implies that buying a unit in the current period has a lower cost than buying a unit in the next period, which prevents the model from intentionally holding backorders. Next, we rewrite the cost function for this case as:

$$\min_{s_t^R \geq s_t^{R-1}} \left\{ \begin{array}{l} \sum_{t=1}^{T} \gamma^{t-1} \mathbb{E}[c(s_t^R - s_t^{R-1}) + h(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^+ \\ + w(s_t^1 - (\hat{f}_t + e(\hat{f}_t)))^+ + b(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^-] \end{array} \right\} \tag{5.1.13}$$

subject to:

$$s_{t+1}^r = s_t^{r+1} - d_t \qquad \forall\, t \geq 1, 1 \leq r \leq R-1 \tag{5.1.14}$$

$$d_t = \hat{f}_t + e(\hat{f}_t) \tag{5.1.15}$$

Similar to the lost-sales case, the optimal cost from day t onward, $\hat{u}_t(\mathbf{s_t}, \hat{f}_t, \hat{f}_{t-1})$, can be written as follows:

$$\hat{u}_t(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1}) = \min_{s_t^R \geq s_t^{R-1}} \{\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1})\} \tag{5.1.16}$$

where

$$\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1}) = \mathbb{E}\begin{bmatrix} c(s_t^R - s_t^{R-1}) + h(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^+ + w(s_t^1 - (\hat{f}_t + e(\hat{f}_t)))^+ \\ + b(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^- + \gamma\mathbb{E}[\hat{u}_{t+1}(\mathbf{s}_{t+1}, \hat{f}_{t+1}, \hat{f}_t)] \end{bmatrix} \tag{5.1.17}$$

5.2 Structural Properties of the Optimal Policy

In this section, we explore structural properties of the optimal ordering policy for the lost-sales case. This structural analysis can be easily extended to the backlogging case. First, we characterize the form of the optimal policy. We show that the cost function satisfies the L$^\natural$-convexity property, which characterizes the optimal ordering policy. Next, we explore monotonicity properties of the optimal policy. The results presented in this section offer valuable perspectives on the optimal ordering policy, which motive us to propose easy to implement intuitive heuristics in the next section. L$^\natural$-convexity is a combination of convex analysis and lattice structure, so we start by introducing the definition of lattice structure. Let $\mathscr{F}$ be either the real space, $\mathbb{R}$, or the integer space, $\mathbb{Z}$. We define $\overline{\mathbb{R}} = \mathbb{R} \cup +\infty$ and $\mathscr{F}^n$ as n-tuples of real or integer numbers. Now, we introduce two operations in $\mathbb{R}^n$ named join ($\vee$) and meet ($\wedge$). For any $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ and $\mathbf{y} = (y_1, y_2, \ldots, y_n)$ in $\mathbb{R}^n$, their join is defined as:

$$\mathbf{x} \vee \mathbf{y} = (max\{x_1, y_1\}, max\{x_2, y_2\}, \ldots, max\{x_n, y_n\}) \tag{5.2.1}$$

and their meet is defined as:

$$\mathbf{x} \wedge \mathbf{y} = (min\{x_1, y_1\}, min\{x_2, y_2\}, \ldots, min\{x_n, y_n\}) \tag{5.2.2}$$

A set $X \subseteq \mathbb{R}^n$ is called lattice if for any $\mathbf{x}, \mathbf{y} \in X$, $\mathbf{x} \vee \mathbf{y}, \mathbf{x} \wedge \mathbf{y} \in X$. Also, X is called a sublattice of $\mathbb{R}^n$ since it takes the supremum and infimum from $\mathbb{R}^n$.

The concept of L$^\natural$-convexity was first introduced by Murota (1998) for functions defined on integer spaces.

Definition 1 (L$^\natural$-convexity) *A function $f : \mathscr{F}^n \to \mathbb{R}$ is L$^\natural$-convex, if for any $\mathbf{x}, \mathbf{y} \in \mathscr{F}^n$ and any $\alpha \in \mathscr{F}_+$:*

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f((\mathbf{x} + \alpha \mathbf{e}) \wedge \mathbf{y}) + f(\mathbf{x} \vee (\mathbf{y} - \alpha \mathbf{e})) \tag{5.2.3}$$

where $\mathbf{e}$ is the all-ones vector.

Definition 2 (Submodularity) *A function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is submodular, if for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$:*

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f(\mathbf{x} \wedge \mathbf{y}) + f(\mathbf{x} \vee \mathbf{y}) \tag{5.2.4}$$

We can present an equivalent definition of L$^\natural$-convexity based on submodularity as follows.

Proposition 1 *A function $f : \mathscr{F}^n \to \overline{\mathbb{R}}$ is L$^\natural$-convex if and only if for ξ in the intersection of $\mathscr{F}$ and any unbounded interval in $\overline{\mathbb{R}}$, $g(\mathbf{x}, \xi) = f(\mathbf{x} - \xi \mathbf{e})$ is submodular on $(\mathbf{x}, \xi)$.*

The proof of this proposition is provided in Simchi-Levi et al. (2005).

Now, in order to establish the L$^\natural$-convexity property, it is necessary to obtain several lemmas. We start by showing that L$^\natural$-convexity is preserved for specific changes in the state. Lemma 1 and Lemma 2 are developed using results from Zipkin (2008).

Lemma 1 *If $v : \mathscr{F}^n \to \mathbb{R}$ is an $L^\natural$-convex function, then the function*

$$\upsilon(\mathbf{s}_1, s_2) = v(\mathbf{s}_1 - s_2\mathbf{e})$$

is also $L^\natural$-convex.

Proof. The function $\upsilon(\mathbf{s}_1, s_2)$ is L$^\natural$-convex if $\omega(\mathbf{s}_1, s_2, \xi) = \upsilon[(\mathbf{s}_1, s_2) - \xi(\mathbf{e}, 1)]$ is submodular:

$$\upsilon[(\mathbf{s}_1, s_2) - \xi(\mathbf{e}, 1)] = \upsilon(\mathbf{s}_1 - \xi\mathbf{e}, s_2 - \xi) = v[(\mathbf{s}_1 - \xi\mathbf{e}) - (s_2 - \xi)\mathbf{e}] = v(\mathbf{s}_1 - s_2\mathbf{e}) = \upsilon(\mathbf{s}_1, s_2)$$

The function $\upsilon(\mathbf{s}_1, s_2)$ is submodular and so L$^\natural$-convex. $\square$

Lemma 2 *Let $v(\mathbf{s}_1, s_2) : \mathscr{F}^n \times \mathscr{F} \to \mathbb{R}$ be an $L^\natural$-convex function. The function $\upsilon(\mathbf{s}_1) =$ $\min_{s_2} v(\mathbf{s}_1, s_2)$ is also $L^\natural$-convex over $\mathscr{F}^n$.*

Proof. If $v(\mathbf{s}_1, s_2)$ is L$^\natural$-convex, then $\omega(\mathbf{s}_1, s_2, \xi) = v[(\mathbf{s}_1, s_2) - \xi(\mathbf{e}, 1)]$ is submodular. We can write:

$$\begin{aligned}\upsilon(\mathbf{s}_1 - s_2\mathbf{e}) &= \min_{\xi} v(\mathbf{s}_1 - s_2\mathbf{e}, \xi) = \min_{\xi} v[(\mathbf{s}_1, s_2 + \xi) - \xi(\mathbf{e}, 1)] = \min_{\xi} \omega[(\mathbf{s}_1, s_2 + \xi, \xi)] \\ &= \min_{\int \leq \xi} \omega(\mathbf{s}_1, \int, \xi)\end{aligned}$$

Since the constraint on the minimum is a sublattice of $\mathscr{F} \times \mathscr{F}$, using Theorem 2.7.6 of (Topkis, 1998), submodularity is preserved. $\square$

Lemma 3 *If v is $L^\natural$-convex and is nondecreasing in its variables, then the function*

$$\hat{v}(s_1, s_2, \ldots, s_n) = v((s_1 - a)^+, (s_2 - a)^+, \ldots, (s_n - a)^+)$$

is $L^\natural$*-convex for all* $(s_1, s_2, \ldots, s_n)$.

Proof. Let

$$\tilde{v}(s_1, s_2, \ldots, s_n) = v(s_1^+, s_2^+, \ldots, s_n^+)$$

Since v is nondecreasing in its variables, we can define $\int_i \geq s_i$ and write:

$$\tilde{v}(s_1, s_2, \ldots, s_n) = min_{\int_i \geq s_i, 0 \leq \int_1 \leq \cdots \leq \int_n} v(\int_1, \int_2, \ldots, \int_n)$$

Since the set associated with the constraints is lattice and the function v is $L^\natural$-convex, using Lemma 2, $\tilde{v}(s_1, s_2, \ldots, s_n)$ is also $L^\natural$-convex. Using Lemma 1, we can conclude that $\hat{v}(s_1, s_2, \ldots, s_n) = \tilde{v}(\mathbf{s} - a\mathbf{e})$ is $L^\natural$-convex. $\square$

Now, we show that the cost function is $L^\natural$-convex.

Theorem 1 *The cost function* $\hat{u}_t(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1})$ *is* $L^\natural$*-convex.*

Proof. We prove the $L^\natural$-convexity of $\hat{u}_t(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1})$ and $\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1})$ by induction. For a finite horizon, we consider $\hat{u}_T(\mathbf{s}_T, \hat{f}_T, \hat{f}_{T-1}) = 0$, so the result certainly holds for $\hat{u}_T(\mathbf{s}_T, \hat{f}_T, \hat{f}_{T-1})$. Assume that $\hat{u}_{t+1}(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1})$ is $L^\natural$-convex. Again, we can rewrite the cost-to-go function using the fulfilled demand:

$$\hat{k}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1} | e(\hat{f}_t), \varepsilon_{t+1}) = \min_{s_t^R \geq s_t^{R-1}} \left\{ \begin{array}{l} \psi_t = c(s_t^R - s_t^{R-1}) + h(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^+ + w(s_t^1 - (\hat{f}_t + e(\hat{f}_t)))^+ \\ \quad + p(s_t^R - (\hat{f}_t + e(\hat{f}_t)))^- + \gamma \hat{u}_{t+1}(\mathbf{s}_{t+1}, \hat{f}_{t+1}, \hat{f}_t)) : \hat{f}_{t+1} = \theta \hat{f}_t + \varepsilon_{t+1} \end{array} \right\} \tag{5.2.5}$$

The first four terms in the objective function ψ_t are $L^\natural$-convex. Since ε_{t+1} is given, the constraint is lattice. By Lemma 3, $L^\natural$-convexity of $\hat{u}_{t+1}(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1})$ implies that $\hat{u}_{t+1}(\mathbf{s}_{t+1}, \hat{f}_{t+1}, \hat{f}_t)$ is $L^\natural$-convex. Since $\hat{u}_{t+1}(\mathbf{s}_{t+1}, \hat{f}_{t+1}, \hat{f}_t)$ is $L^\natural$-convex, by using Lemma 2, $\hat{k}_t$ is $L^\natural$-convex because of the $L^\natural$-convexity of ψ_t. Now, we can calculate the expected value of

$\hat{k}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1} | e(\hat{f}_t), \varepsilon_{t+1})$ over the forecast error distribution, $F_{e(\hat{f})}$, and the predicted demand residuals, F_ε:

$$\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1}) = \mathbb{E}_{F_{e(\hat{f})}, F_\varepsilon}[\hat{k}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1} | e(\hat{f}_t), \varepsilon_{t+1})] \tag{5.2.6}$$

$$\hat{u}_t(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1}) = \min_{s_t^R \geq s_t^{R-1}} \{\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1})\} \tag{5.2.7}$$

Using Corollary 2.6.2 of (Topkis, 1998), L$^\natural$-convexity is preserved by expectation, so $\hat{g}_t$ is L$^\natural$-convex. Also, by using Lemma 2, $\hat{u}_t$ is L$^\natural$-convex. □

Theorem 1 verifies the L$^\natural$-convexity of the cost function. We conjecture that the theorem holds in more generality for the forecast function (5.1.2), in particular when $\hat{f}_t$ is a linear function of i previous forecast values and j previous demand values. The L$^\natural$-convexity property of the cost function allows us to prove that the optimal ordering policy is a base-stock policy in which the order-up-to level is a function of the state. We define $\hat{f}_-$ as the predicted demand value in the previous period.

We next explore monotonicity properties of the optimal order-up-to level s^R. Due to the dependence of the forecast error on the forecast, the order quantity is in general not monotone in the demand forecasts. For example, a demand forecast may lead to a lower order-up-to quantity than a lower demand forecast that exhibits higher variance in its error distribution. To address this, we assume that the forecast error is a multiplicative function of the forecast as follows:

$$e(\hat{f}_t) = Z_t \hat{f}_t \tag{5.2.8}$$

The random multipliers Z_t are assumed to be independently and identically distributed with bounded support on [-1, ∞). Next, we rewrite the demand as a function of the predicted demand and the forecast error as:

$$d_t = \hat{f}_t + Z_t \hat{f}_t = \hat{f}_t(1 + Z_t) = \zeta_t \hat{f}_t \tag{5.2.9}$$

where the random variable $\zeta_t = 1 + Z_t$. Hence, demand takes the form of a multiplicative function of the expected demand, which is a common assumption in the literature (Bernstein et al., 2016; Simchi-Levi et al., 2005).

Theorem 2 *(Monotonicity properties of the model) The cost functions $\hat{g}_t(\mathbf{s}_t, s_t^R, \hat{f}_t, \hat{f}_{t-1})$ and $\hat{u}_t(\mathbf{s}_t, \hat{f}_t, \hat{f}_{t-1})$ are L$^\natural$-convex. Thus, the optimal order-up-to level s^R is nondecreasing in the current inventory level* $\mathbf{s}$ *and current and previous demand forecasts, $\hat{f}, \hat{f}_-$.*

Proof. Due to the L$^\natural$-convexity of the cost function $\hat{u}_t$, and based on Lemma 2.3.5 in (Simchi-Levi et al., 2005), the optimal solution s^R is nondecreasing in the inventory values $\mathbf{s}$, the current forecast value $\hat{f}$, and previous predicted demand value $\hat{f}_-$. $\square$

Theorem 2 establishes a monotonicity result for the optimal order-up-to level in the demand forecasts. It states that the optimal order-up-to level is a nondecreasing function of each element of the state, including demand forecasts. Since order quantity is the difference between order-up-to and inventory levels, the order quantity is nondecreasing in the demand forecasts and nonincreasing in the inventory levels.

5.3 A Proposed Heuristic

In the previous section, we have shown L$^\natural$-convexity of the cost function. L$^\natural$-convexity of the cost function guarantees that a local optimum is the global optimum. Thus, we have shown that the optimal policy is a state-dependent order-up-to level, s^R, that depends on the inventory values, and current and previous demand forecast values. We have also proved that the order-up-to levels s^R are nondecreasing in the forecast values. Motivated by this property and due to the dependence of the optimal policy on the state, we propose a heuristic as a substitute for the optimal policy. The ordering quantity is chosen as follows: in each period t, the inventory is filled up to an order-up-to level, S_t. The order-up-to levels are defined as an affine function of the current period's forecasts, $\hat{f}_t$, as follows:

$$S_t(\hat{f}_t) = \alpha \hat{f}_t + \beta \tag{5.3.1}$$

The coefficient of the predicted demand, α, captures the dependence of the order-up-to level on the demand forecast. One can intuitively think of this heuristic as one in which there is a fixed base-stock level (β) that is adjusted depending on the demand forecast (through the coefficient α). At one extreme, $\alpha = 0$ would correspond to a fixed base-stock policy, while $\alpha = 1$ and $\beta = 0$ would correspond to the policy where orders are equal to the forecasts. This heuristic is consistent with the monotone nature of the optimal policy. Given that demand forecasts are generated by models such as ARIMA or LSTM networks, which determine the forecasts based on previous demand and demand forecasts, we can infer that $\hat{f}$ contains a combination of the past history. Thus, only $\hat{f}$ is included in (5.3.1) for computing the order quantity. In the next chapter, we will evaluate the heuristic and compare it with other ordering policies.

5.4 Conclusion

In this chapter, we studied the structural properties of the optimal ordering policy for a perishable product with a fixed shelf life over a finite horizon for both lost-sales and backlogging cases when demand forecasts are included in the inventory model. Using a property called $L^\natural$-convexity, we showed that the optimal ordering policy is state-dependent, depending on the current inventory level and the forecast history. Moreover, we showed the monotonicity properties of the optimal policy in the state which motivated us to propose a heuristic in which the order-up-to levels are an affine function of demand forecasts. These results show that incorporating demand forecasts in the inventory model results in the state-dependent optimal policy being dependent on demand forecasts. This is in contrast with recent work that includes factors that influence demand, but does not make explicit demand forecasts (Guan et al. (2017); Abouee-Mehrizi et al. (2022)). This dynamic approach to inventory management offers several advantages, such as allowing the inventory level to adjust based on demand fluctuations and generating order quantities that closely align with actual demand. This not only reduces the wastage rate but enhances the transparency of suppliers and retailers, thereby aiding suppliers in their planning. Moreover, it gives freedom to managers to choose suitable demand estimators which can be plugged into our proposed ordering policy (these demand estimators may serve other purposes beyond inventory management). In the next chapter, we evaluate the performance of the heuristic by implementing it for a real dataset. We also consider the cases when more history is included in the heuristic (as the constructed estimators include more history than (5.1.2)) to further investigate our conjecture.

Chapter 6

Experiments and Results

In this chapter, we evaluate the performance of the heuristic proposed in Chapter 5. As was the case for our forecasting models, we consider the daily demand aggregated over four hospitals due to internal inventory management procedures. Firstly, we compare the performance of the heuristic with the case that more history is included in the heuristic. As discussed in the previous chapter, this would be appropriate if the constructed estimators use more history, which will be the case for the estimators constructed as per the approaches in Chapter 4. This will help us to further explore our conjecture numerically. Secondly, we perform sensitivity analysis to explore the generality of the heuristic. Thirdly, we investigate a special case of the general inventory problem by considering just shortage and wastage costs. For platelet products, the key cost factors revolve around shortages (given their impact on human health) and wastage costs (due to their price). This concept can be extended to other perishable items, leading us to contemplate the use of a simplified cost function. We evaluate this case for three different forecasts and compare it with a classic base-stock policy and the policy that is currently used at the hospitals. Different sensitivity analyses are done for this case as well. Finally, we explore the possibility of reducing the ordering frequency.

6.1 Results for the Proposed Heuristic

This section includes the experimental results for the heuristic proposed in Chapter 5. For all evaluations, we use the platelet transfusion data described in Chapter 3. Forecasts are generated based on models described in Chapter 4. The current inventory management that Hamilton hospitals use for their blood supply chain is similar to a base-stock policy. In a base-stock policy (Glasserman and Tayur, 1994), a replenishment order is placed to increase the inventory level up to the base-stock level, S, when the inventory level at the end of a period is below S. Base-stock policies are known to perform well, especially when the shortage cost is much higher than the holding cost. However, they can result in high inventory levels, in particular when the base-stock level is high. This can be problematic in systems with limited inventory capacities. The base-stock level should be sufficiently large to prevent shortages, but at the same time should not result in significant wastages (van Sambeeck et al., 2022). In this section, we evaluate the performance of the heuristic and compare it with that of a classic base-stock policy and the current practice at the hospitals.

6.1.1 Training the Heuristic

The values of α and β are trained using historical data. With the current time denoted by t, we choose α and β using historical data from a time window of length T (the training window), i.e., data from time $t_0 - T + 1$ to t_0, where $T < t_0 < t$. In our case, α and β are updated weekly, with a training window of the most recent two years of data. The historical data consists of the realized demands d_u and the corresponding estimates $\hat{f}_u$ over the training window $u = t_0 - T + 1, ..., t_0$, where t_0 is the end of a training window. The chosen values of α and β are then used to determine the base-stock function (5.3.1).

The values of α and β are determined by solving the following optimization problem:

$$\min_{\alpha,\beta} \qquad \sum_{u=t_0-T+1}^{t_0} cq_u + h(\sum_{r=1}^{R} x_u^r - d_u)^+ + p(d_u - \sum_{r=1}^{R} x_u^r)^+ + w\left(x_u^1 - d_u\right)^+ \tag{6.1.1}$$

subject to:

$$x_u^R = q_u, \qquad u = t_0 - T + 1, \ldots, t_0 \tag{6.1.2}$$

$$q_u = \alpha \hat{f}_u + \beta - (\sum_{r=1}^{R} x_{u-1}^r - d_{u-1})^+, \qquad u = t_0 - T + 1, \ldots, t_0 \tag{6.1.3}$$

$$x_u^{r-1} = x_{u-1}^r - a_{u-1}^r, \qquad u = t_0 - T + 1, \ldots, t_0; \quad r = 1, 2, \ldots, R \tag{6.1.4}$$

$$a_u^r = \begin{cases} d_u - \sum_{j=1}^{r-1} x_u^j & \text{if } d_u - \sum_{j=1}^{r-1} x_u^j < x_u^r \\ x_u^r & otherwise \end{cases} \tag{6.1.5}$$

$$q_u \geq 0, \qquad \alpha \geq 0 \tag{6.1.6}$$

The cost function, (6.1.1), consists of ordering, holding, shortage and wastage costs. The ordering cost is the product a constant ordering cost (c) and the number of ordered units on day u, q_u. The holding cost is the product of a constant holding cost (h) and the number of units in the inventory on day u. The shortage cost is the product of a constant shortage cost for each shortage (p) and the number of shortages, defined as the difference between the actual demand, d_u, and the inventory level on day u, if the inventory level is lower than the actual demand. Finally, the wastage cost is the product of a constant wastage cost (w) and the number of expired units on day u, $\left(x_u^1 - d_u\right)^+$. In (6.1.2), the inventory update is defined based on new arrivals, q_u (note that in the solution to this optimization problem, the quantities q_u will in general be different than the historical order quantities). New arrivals (the ordering quantity) are the difference between the target inventory level, $S(\hat{f}_u)$, and the inventory level at the end of the previous day, as shown in (6.1.3). The inventory values, x_u^r, are updated based on new arrivals and the number of consumed units on day u, a_u^r, as stated in (6.1.4). The consumed units of r remaining days after withdrawing all products having remaining days from 1 to $r-1$ using the FIFO withdrawal policy is denoted by a_t^r. The usage of each unit with r remaining days based on the FIFO policy is calculated in (6.1.5).

We also compare the heuristic with policies that have the order-up-to level as a linear function of current and previous demand forecasts. Considering that forecasts are generated by models that have in themselves a history of demand and demand forecasts, like ARIMA, including the current period's forecast has in itself demand history. To evaluate the effect of directly including forecast history in the order-up-to level function, we redefine it as:

$$S_t = \Sigma_{k=1}^{M} \alpha_k \hat{f}_{t-k+1} + \beta \qquad M = 1, 2, \ldots, T \tag{6.1.7}$$

where $M = 1$ corresponds to the proposed heuristic. As M grows, the correlation between the order-up-to level and previous forecasts increases. This would allow us to evaluate our conjecture that the optimal ordering policy depends on a history of forecasts and demand values.[1] Demand forecasts are generated from an LSTM network (forecast MAPE (Mean Absolute Percentage Error) = Mean [sd]: 26.45 [8.39] and forecast RMSE (Root Mean Squared Error) = Mean [sd]: 4.32 [2.01]) as presented in Chapter 4.

[1] Since forecasts are generated by models such as ARIMA or LSTM networks, they are derived based on previous demand and demand forecasts. So, one can infer that they contain a combination of the past history.

6.1.2 Results

To compare the performance of the proposed heuristic and the cases when previous forecasts are considered in the order-up-to function, we use a data-driven approach; values of α and β are obtained by minimizing the total cost over a planning horizon of one year (2017) in the data. These values are used to calculate the cost over a six month period, from January 2018 to June 2018. The initial inventory is 20. Given that the actual costs are not a determining factor (only the relative cost values matter), the ordering cost is set at 1 per unit, the holding cost at 0.25 per unit, the wastage cost at 1 per wasted unit, and the shortage cost at 5 per unit. Shortages are handled by placing backorders so no demand is lost. The inventory is updated based on the FIFO policy.

Table 6.1 gives the total cost, number of wastages and backorders for the proposed heuristic, order-up-to functions with history, a data-driven base-stock policy and the policy that is currently used at the hospitals. To train the data-driven base-stock policy, we force α to be zero in the optimization problem in Section 6.1.1. We train the base-stock model for one year (2017), similar to how the heuristic is trained, and update the inventory for January 2018 to June 2018. The policy currently used at the hospitals is a base-stock policy in which the base-stock value (S) is chosen based on the assessment of clinical staff.

The first five rows correspond to (6.1.7), where $M = 1$ represents the proposed heuristic. As we see in Table 6.1, including more history in the order-up-to-level function results in a slight decrease in total cost and the number of backorders.

Table 6.1: Costs of the proposed heuristic, an order-up-to function with history, a base-stock policy and the current practice

Model	**Total cost**	**Number of wastages**	**Number of backorders**
$M = 1$	4250	0	37
$M = 2$	4187	0	35
$M = 3$	4166	0	33
$M = 4$	4162	0	32
$M = 5$	4156	0	30
The base-stock policy	4397	0	42
The actual policy at hospitals	*unknown*	231	*unknown*

Figure 6.1 illustrates cost improvements of the order-up-to functions with history and base-stock policy compared to the proposed heuristic. As we see in Figure 6.1, as more forecast history is added to the order-up-to-level function, in general there is a cost improvement which becomes smaller as more history is added to the model. These results support our conjecture in Chapter 5 that the forecasts can depend on more than one previous forecast value.

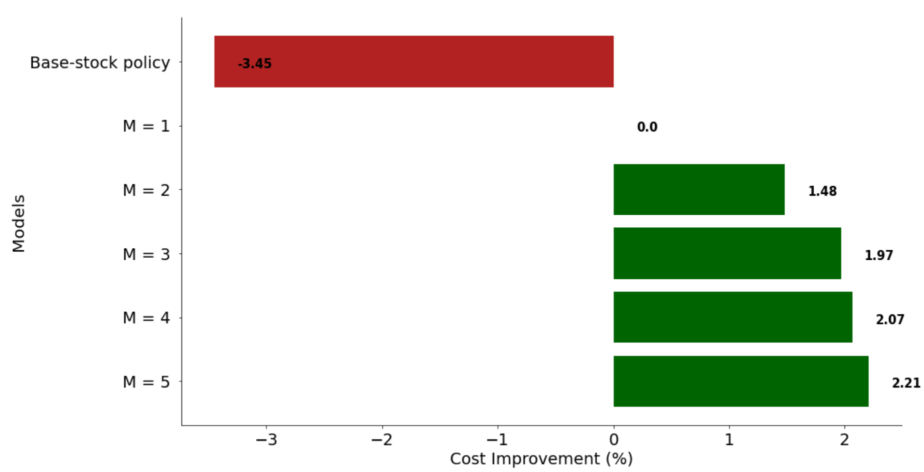

Figure 6.1: Cost improvement for different policies in comparison to the proposed heuristic

Table 6.2 shows the mean and standard deviation of the inventory, ordering quantities, and actual demands for the proposed heuristic, order-up-to functions with history, a base-stock policy, and the actual policy at the hospitals. The proposed heuristic and order-up-to functions with history result in significantly lower inventory levels compared to a base-stock policy. That reduces not only the ordering, wastage, and holding costs but also keeps the shortages low. The proposed policy with/without previous history and data-driven base-stock policy have a significantly lower inventory level than the actual policy at the hospitals. Furthermore, the mean daily order quantity is very close to the actual demand for all the models, including the base-stock policy. The mean order quantity of the actual policy is slightly larger than the actual demand and has high variability.

Table 6.2: Inventory levels of the proposed heuristic, an order-up-to function with history, a base-stock policy and the current practice

Model	**Inventory (mean $\pm$ std)**	**Order Quantity (mean $\pm$ std)**	**Demand (mean $\pm$ std)**
$M = 1$	7.52 ± 5.52	19.82 ± 8.61	20.55 ± 7.34
$M = 2$	7.26 ± 5.38	19.71 ± 9.10	20.55 ± 7.34
$M = 3$	7.26 ± 5.38	19.71 ± 9.10	20.55 ± 7.34
$M = 4$	7.07 ± 5.13	19.73 ± 8.64	20.55 ± 7.34
$M = 5$	6.20 ± 5.06	19.43 ± 8.66	20.55 ± 7.34
The base-stock policy	11.23 ± 6.99	20.21 ± 7.02	20.55 ± 7.34
The actual policy at the hospitals	43.35 ± 7.49	22.08 ± 9.95	20.55 ± 7.34

6.1.2.1 Sensitivity to the Shelf Life

We investigate the sensitivity of our proposed heuristic to different shelf lives of 2, 3, 4, and 5 days and compare it with a base-stock policy. Figure 6.2 illustrates the total cost, wastages, shortages, and inventory levels for different shelf lives. We see in Figure 6.2 (a) that shorter shelf life leads to an increase in total cost, primarily due to a lower inventory level (see Figure 6.2 (c)) that increases shortages. As shown in Figure 6.2 (b), shorter shelf lives increase shortages as a result of decreased on-hand inventory. Additionally, shorter shelf lives lead to significantly higher wastage as their shelf life is comparatively smaller. Products with a shelf life of 3, 4, and 5 days have zero wastage, while a product with a shelf life of 2 days has a wastage of around 45 and 110 for the proposed heuristic and the base-stock policy, respectively. This can be attributed to the fact that longer shelf lives allow for items to be stored for longer periods, allowing for higher on-hand inventory, which leads to lower wastages and shortages.

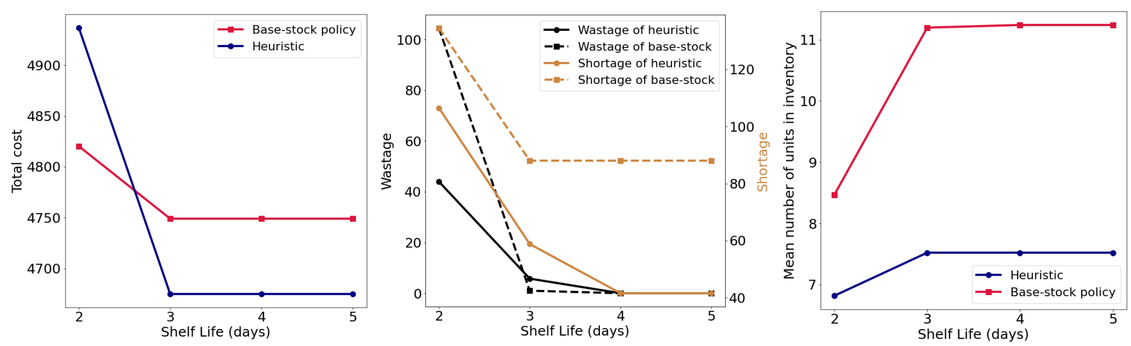

(a) (b) (c)

Figure 6.2: Sensitivity to the shelf life for the proposed heuristic

6.1.2.2 Sensitivity to the Costs

We also explore the sensitivity of the heuristic and the base-stock policy to ordering, holding, wastage, and shortage costs. As we change the wastage cost, the wastages remain zero for both the heuristic and the base-stock policy. This is mainly due to the ordering quantities being close to the actual demand and keeping the on-hand inventory low, which may result in shortages but keeps the wastage zero. Figure 6.3 represents the sensitivity of the total cost, mean daily number of units in the inventory, and mean daily order quantities to the ordering, holding, and shortage costs. We see in Figure 6.3 (a) that as the ordering cost increases, the total cost increases for both the heuristic and the base-stock policy, with a small gap between the costs for the heuristic and the base-stock policy. As we expect, by increasing the ordering cost, the mean daily ordering quantity reduces for both models (Figure 6.3 (c)), but the reduction has a gradual slope. Similarly, the mean daily number of units in inventory decreases gradually as the ordering cost increases.

Figures 6.3 (d) - 6.3 (f) illustrate sensitivity to the holding cost. As the holding cost rises, the total cost for both models rises with a constant gap. The mean daily number of units in the inventory has a significant decrease for a slightly larger holding cost, an increase of 20%. Figure 6.3 (f) shows that by increasing the holding cost, the gap between the base-stock and the heuristic increases while both models see a decrease in the number of ordered units.

Finally, Figures 6.3 (g) - 6.3 (i) depict the sensitivity to the shortage cost. As mentioned earlier, we consider the backlogging case, meaning that the shortages are backordered. We see in Figure 6.3 (g) that increasing the shortage (backorder) cost leads to a larger gap between the total cost of the heuristic and the base-stock policy. For a shortage cost of 2.5, the total costs of both models are very close, but for a shortage cost of almost 10 times larger, there is a significant gap between the costs of the two models. Moreover, as shown in Figure 6.3 (h), as the shortage cost increases, while the on-hand inventory for the heuristic and the base-stock policy rises, the gap between the heuristic and the base-stock policy widens. The increase in the inventory level shows an initial steep slope, followed by a gradual slope from a shortage cost of 10 per unit. The on-hand inventory for the base-stock policy is almost 50% higher than the heuristic when the shortage cost is 20 per unit. Similarly, as the shortage cost increases, there is a sudden increase in the mean daily order quantity followed by a gentle slope.

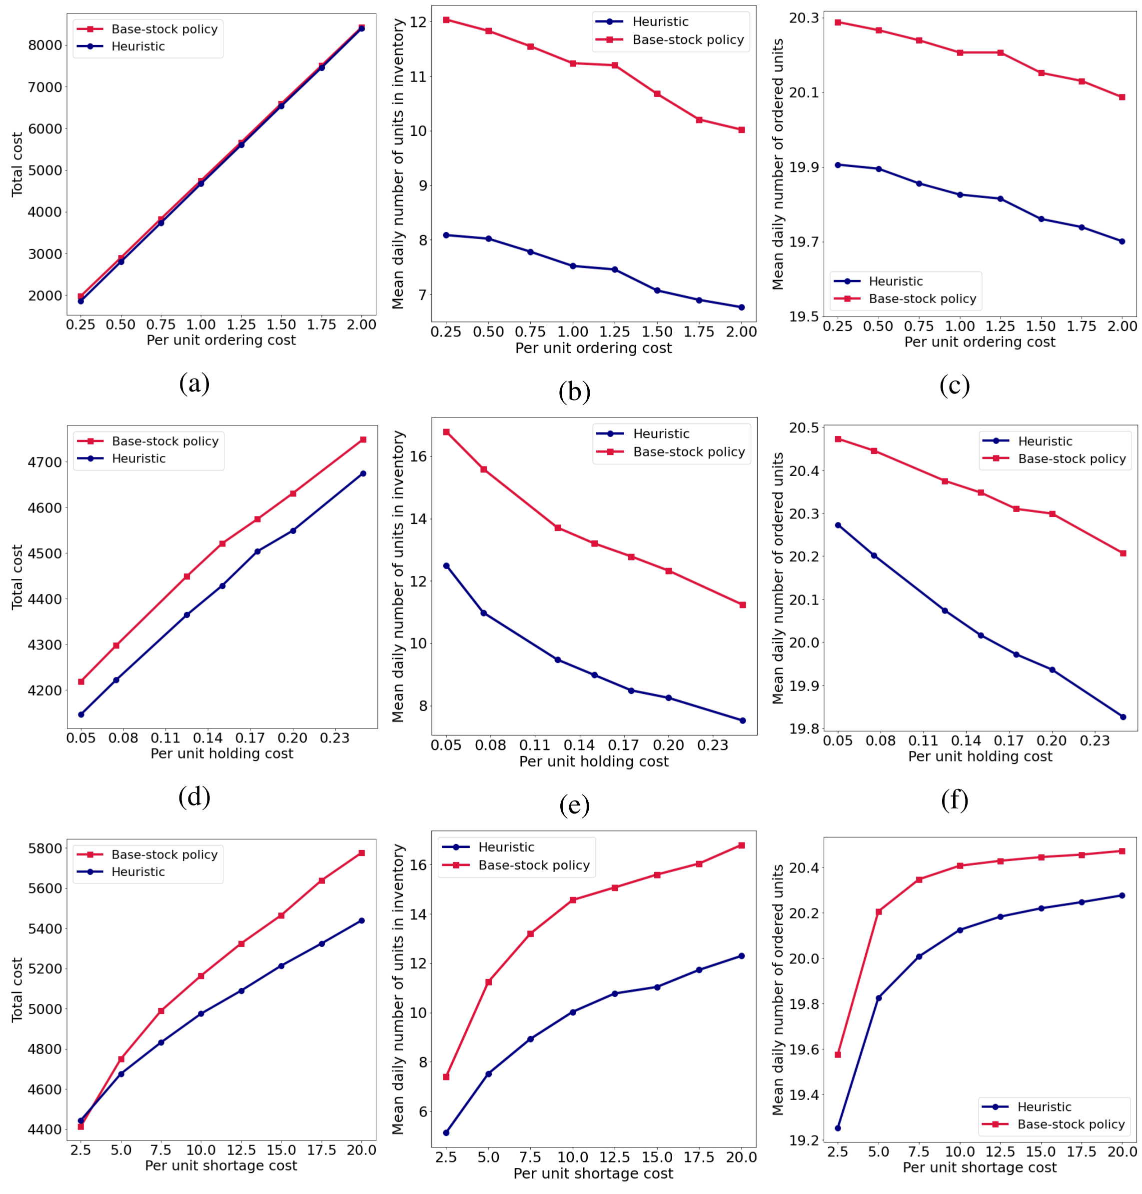

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Figure 6.3: Sensitivity to the costs for the proposed heuristic

6.1.2.3 Impact of Forecast Error on the Heuristic

In this section, we investigate the impact of forecast error on the performance of the proposed heuristic and order-up-to functions with history. Specifically, we examine the effect of two different forecast errors on shortages, wastages, order-up-to levels, and inventory levels. Table 6.3 presents the results for two forecasting models, an LSTM network (forecast MAPE = Mean [sd]: 26.45 [8.39] and forecast RMSE = Mean [sd]: 4.32 [2.01]) and an ARIMA model (forecast MAPE = Mean [sd]: 33.19 [8.33] and forecast RMSE = Mean [sd]: 6.81 [2.09]), and the actual policy at the hospitals.

First, we explore the Coefficient of Variation (CV) of the order-up-to levels, inventory mean and standard deviation, and subsequently the CV of order quantities. As we see in Table 6.3, when the forecast error is lower, the order-up-to levels exhibit lower variation. Inventory mean and standard deviation are comparable for both forecast errors and significantly lower than the current inventory level. Also, the order quantity has larger variation when the forecast error is higher, which is a direct cause of the higher inventory level. Additionally, as we can see in Table 6.3, when the forecast error is large, including more history results in an increase in the CV of the order-up-to levels and consequently an increase in the inventory mean and standard deviation.

Regarding shortages and wastages, we observe that regardless of the forecast error, the rates of shortages and wastages consistently remain low (almost the same for this dataset). When the forecast error is lower, the inclusion of additional history leads to a slight reduction in the average daily shortage. However, when the forecast error is higher, incorporating more history results in the opposite effect, causing a small increase in the average daily shortage. Overall, the proposed heuristic demonstrates consistent performance across different forecast errors.

Table 6.3: Impact of forecast error on the order-up-to level, inventory level, shortages, and wastages

Model	CV of Order-up-to levels	Inventory (mean + std)	CV of order quantity	Average wastage per period	Average shortage per period
Forecast Error: RMSE = 4.32 M = 1	7.71	7.52 ± 5.52	43.44	0	0.20
Forecast Error: RMSE = 4.32 M = 2	7.86	7.26 ± 5.38	46.16	0	0.19
Forecast Error: RMSE = 4.32 M = 3	7.86	7.26 ± 5.38	46.16	0	0.18
Forecast Error: RMSE = 6.81 M = 1	11.24	8.51 ± 6.07	39.37	0	0.20
Forecast Error: RMSE = 6.81 M = 2	14.06	9.99 ± 7.80	38.37	0	0.21
Forecast Error: RMSE = 6.81 M = 3	14.06	9.99 ± 7.80	38.37	0	0.21
The actual policy at the hospitals	*unknown*	43.35 ± 7.49	45.06	1.27	*unknown*

6.2 A Special Case for Platelet Data - Shortage and Wastage Costs

In Section 6.1.1, we trained a data-driven optimization model to obtain α and β values. In this section, we consider a scenario when only shortage and wastage costs are considered, which closely matches the practical use case for platelet inventory management. The experimental results are presented as follows. Firstly, performance of the proposed heuristic in terms of cost, inventory levels, and ordering quantities is explored. Secondly, in Section 6.2.2 we compare the proposed heuristic's performance with a base-stock policy and the current ordering policy. Next, sensitivity analyses are done for the proposed model's behaviour in different scenarios. Finally, different ordering frequencies are explored.

6.2.1 Results for the Proposed Policy

We present and investigate the results acquired from the experiments where the model is trained with the assumption that the initial inventory is 20. The cost ratio (shortage to wastage cost) is five, with the wastage cost set to one (Haijema, 2013). The model is trained considering demand forecasts from an ARIMA model (forecast MAPE = Mean [sd]: 33.19 [8.33] and forecast RMSE = Mean [sd]: 6.81 [2.09]), a lasso regression (forecast MAPE = Mean [sd]: 31.70 [7.50] and forecast RMSE = Mean [sd]: 6.22 [1.59]), and an LSTM network (forecast MAPE = Mean [sd]: 26.45 [8.39] and forecast RMSE = Mean [sd]: 4.32 [2.01]), as presented in Chapter 4. The optimization model is retrained every week using a rolling window method with a training window size of two years of the most recent data, including demand forecasts, starting from 2016. The inventory is updated for 2018 by retraining every week.

Training the model results in a range of α and β values with a zero cost over the training window. Selecting from multiple α and β values is challenging. One way to force unique optimal α and β values is adding a holding cost, but that may reduce inventory levels and induce shortages. We choose α and β values that would result in less variant inventory levels, since in real situations highly variable inventory levels may be seen as problematic. Highly variable inventory levels are more likely to lead to shortages and wastages, so we propose a heuristic to find (sub)optimal α and β values with the minimum inventory variations. Since α is the coefficient of demand forecasts in the base-stock function (see (5.3.1)), the base-stock level standard deviation and the inventory standard deviation are controlled by α, and so the optimal α value can be obtained by minimizing the inventory standard deviation over the training window. Since β is the constant in the base-stock level function (5.3.1), it impacts the mean of the base-stock levels and consequently the mean of the inventory. We introduce a measure, combined variation (ψ), which considers both inventory standard deviation and coefficient of variation, and select the α and β values that result in the minimum value of ψ over the training window. First, we define $\boldsymbol{I}$ as the total inventory over the training window.

$$\boldsymbol{I} = \sum_{u=t_0-T+1}^{t_0} \left(\sum_{r=1}^{R} I_u^r \right) \tag{6.2.1}$$

Using (6.2.1), we define $\boldsymbol{I}_\mu$ as the inventory mean, $\boldsymbol{I}_\sigma$ as the inventory standard deviation, and $\boldsymbol{I}_{CV} = \frac{\boldsymbol{I}_\sigma}{\boldsymbol{I}_\mu}$ as the inventory coefficient of variation over the training window. Since the inventory standard deviation and coefficient of variation have different ranges, we use z-score normalization to have both variables on the same scale. The combined variation, ψ, is then defined as the combination of the square of the scaled inventory standard deviation and the square of the scaled inventory coefficient of variation:

$$\psi = (scale(\boldsymbol{I}_\sigma))^2 + (scale(\boldsymbol{I}_{CV}))^2 \tag{6.2.2}$$

When the inventory standard deviation is large, minimizing ψ is driven by minimizing the variation. So, to choose the (sub)optimal α value, the minimum ψ value can be used. On the other hand, when the inventory standard deviation is small, the coefficient of variation is also small, minimizing ψ is driven by maximizing inventory mean, which can be used to choose the (sub)optimal β value. Thus, the two terms in (6.2.2) can be used to find the (sub)optimal α and β pair. We use the StandardScaler() function in the sklearn package in Python to scale inventory standard deviation and coefficient of variation to have unit variance.

To calculate the order-up-to levels, the order-up-to level function (5.3.1) is used, and the order-up-to levels are computed based on the selected α and β values. The ordering quantity for the next period is the difference of the base-stock level and the inventory level at the end of the current period.

For the data used in this study, a wide range of α and β values result in zero shortage and wastage. We then restrict α to be between 0 and 1, to take a subset that can be interpreted as a forecast-adjusted base-stock policy. The optimal $[\alpha, \beta]$ values for the proposed model when it is retrained every week and by using ARIMA forecasts for 2018 are [0.6, 37] for weeks 1 to 14, and [0.4, 39] for weeks 15 to 52. Using lasso regression forecasts, the optimal $[\alpha, \beta]$ values are [0.4, 36] for weeks 1 to 33, and [0.8, 34] for weeks 34 to 52. For LSTM network forecasts, the optimal $[\alpha, \beta]$ values are [0.6, 35.5] for weeks 1 to 13, and [0.6, 34] for weeks 15 to 52. Figure 6.4 is a representation of choosing optimal α and β values for LSTM network forecasts based on ψ. In Figure 6.4 (a), for each α value, each point corresponds to a specific β value. The optimal α value is the one that gives the minimum ψ value, $\alpha = 0.6$, which can directly be obtained from the first term in (6.2.2). Figure 6.4 (b) shows ψ as a function of β for $\alpha = 0.6$. The optimal β value is obtained as the minimum point in the graph, 35.5.

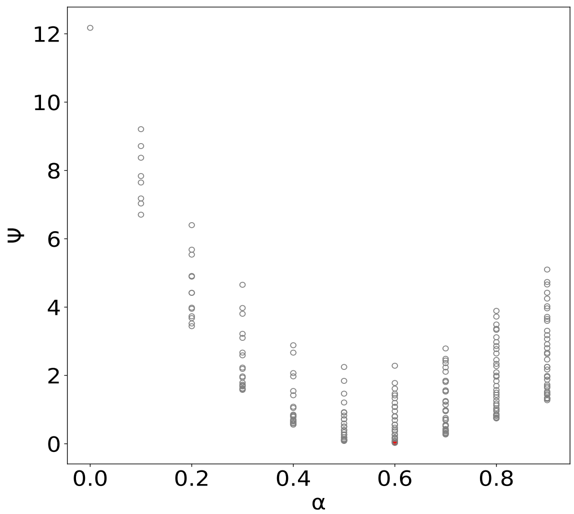

(a) α vs. ψ

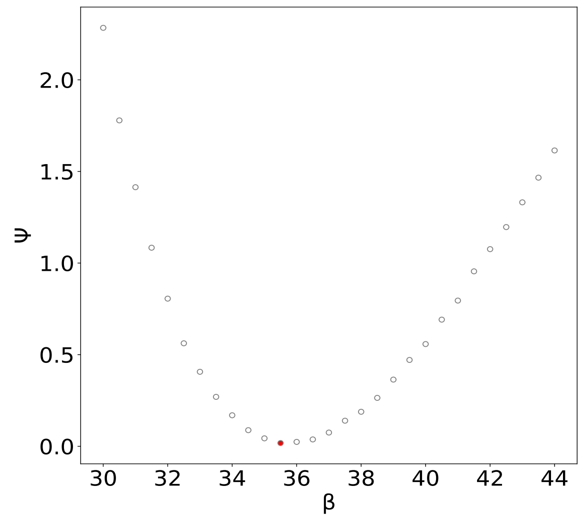

(b) β vs. ψ for $\alpha = 0.6$

Figure 6.4: Calculating optimal α and β values

Using forecasts from any of the considered methods results in a zero cost, meaning zero shortage and wastage, which shows that the performance does not highly depend on forecast quality. The mean ordering quantity for all methods is 19.39, which is very close to the mean actual demand, 19.49, and much smaller than the mean actual ordering quantity, 21.54. The mean actual ordering quantity is calculated from the hospital data. Figure 6.5 compares the order quantities for the proposed heuristic trained with data from each forecasting model for 2018. As we can see from this figure, the order quantities of these models are quite similar, which means that using the proposed heuristic results in having order quantities close to the actual demand, regardless of the forecast quality.

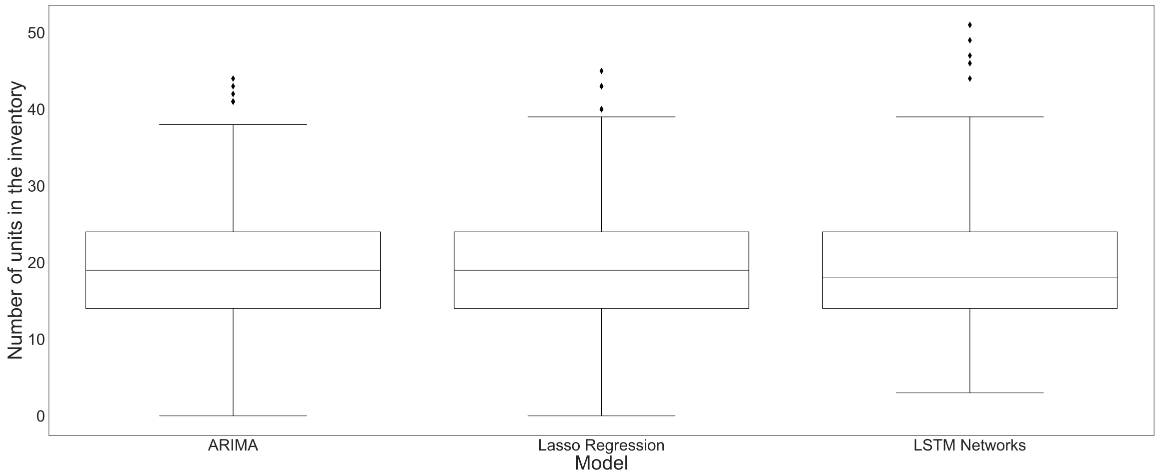

Figure 6.5: Order quantities for the proposed heuristic trained with data from ARIMA, lasso regression and LSTM network for 2018

Next, we explore the inventory level of the proposed heuristic using the ARIMA, lasso regression and LSTM network demand forecasts for 2018. As we can see in Figure 6.6, the mean inventory level using the ARIMA model forecasts is 29.67 and the standard deviation is 7.52 for 2018. By using the lasso regression forecasts, the mean inventory level is 28.85 and the standard deviation is 7.23, and by using the LSTM network forecasts, the inventory mean and standard deviation are 26.96 and 6.82, respectively. These values imply that having a smaller forecast error results in a lower mean inventory level and variation (as a reminder using forecasts from either of the forecasting approaches resulted in no shortages or wastages).

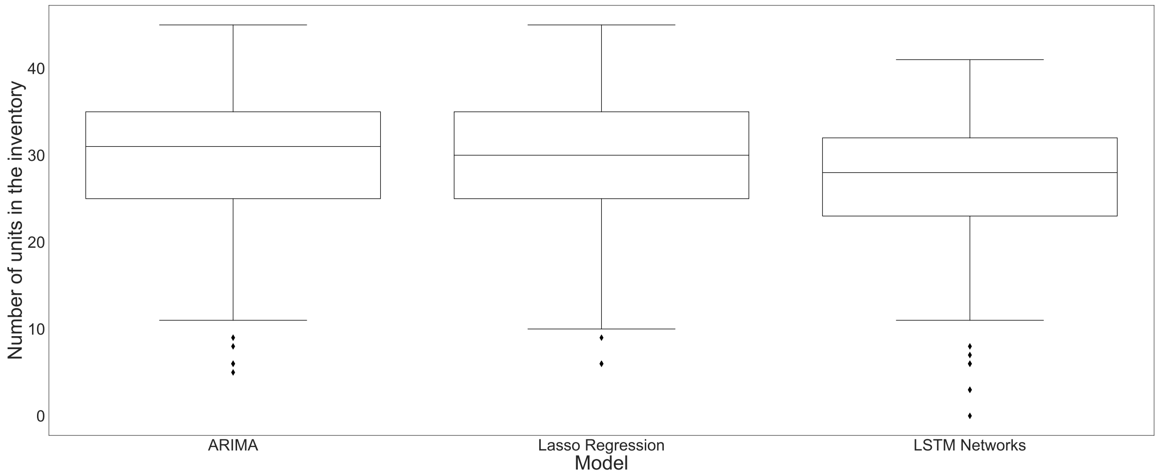

Figure 6.6: Inventory levels of the proposed model trained with data from ARIMA, lasso regression and LSTM network for 2018

6.2.2 Comparative Analysis

In this section, we compare the results for the proposed heuristic with the policy currently used in hospitals (which we call the actual case), and a data-driven base-stock policy. We train a data-driven base-stock model with two years of data (2016-2017) with a cost ratio of five (shortage to wastage cost), similar to how the proposed heuristic is trained. The base-stock policy results in a cost of four with four wasted units for 2018.

Figure 6.7 gives a comparison of the inventories of five methods for 2018: i) the actual inventory with the current policy in hospitals, ii) order under actual demand and consume using the FIFO withdrawal policy, iii) the proposed heuristic using LSTM network forecasts, iv) the proposed heuristic using ARIMA forecasts, and v) the base-stock policy. The minimum inventory mean and standard deviation are when the order quantity is equal to the actual demand, and as we see in Figure 6.7, the mean inventory is equal to the initial inventory, 20, and the inventory standard deviation is 0. For the actual case, the inventory mean and standard deviation are mean[sd]: 42.60[7.30]. The high inventory level, almost twice as large as the actual demand, can be considered as a reason for the current high wastage rate. The inventory levels in the proposed model for the LSTM network and ARIMA forecasts are mean[sd]: 26.90[6.74] and mean[sd]: 28.82[8.11], respectively. Finally, the mean and standard deviation for the base-stock policy are mean[sd]: 31.58[7.66]. These results show that using a forecast-dependent base-stock policy not only leads to a zero cost for our data, but keeps the on-hand inventory much lower. Less on-hand inventory for perishable products is important for the following reasons: i) it results in fresh inventory; and ii) it can be beneficial when there exists a high holding cost or when the inventory capacity is limited.

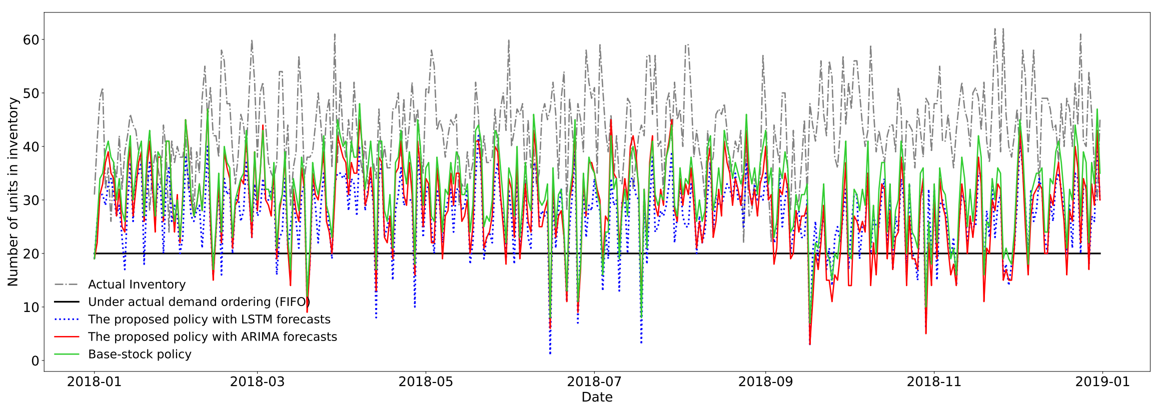

Figure 6.7: Comparison of the inventory levels for the proposed heuristic using LSTM network and ARIMA forecasts, base-stock policy, actual inventory, and ordering by actual demand (FIFO withdrawal policy)

We also compare the ordering quantities for the proposed heuristic with the actual case and the actual demand. Figure 6.8 is a representation of the ordering quantities for the proposed heuristic using LSTM forecasts, the actual ordering quantities, and ordering by actual demand. As we can see in Figure 6.8, ordering quantities for the proposed heuristic are very close to the actual demand, with a mean[sd]: 19.39[7.95], and mean[sd]: 19.36[7.65] for ordering by actual demand. For the actual case, the order quantities mean and standard deviation are mean[sd]: 21.54[9.90].

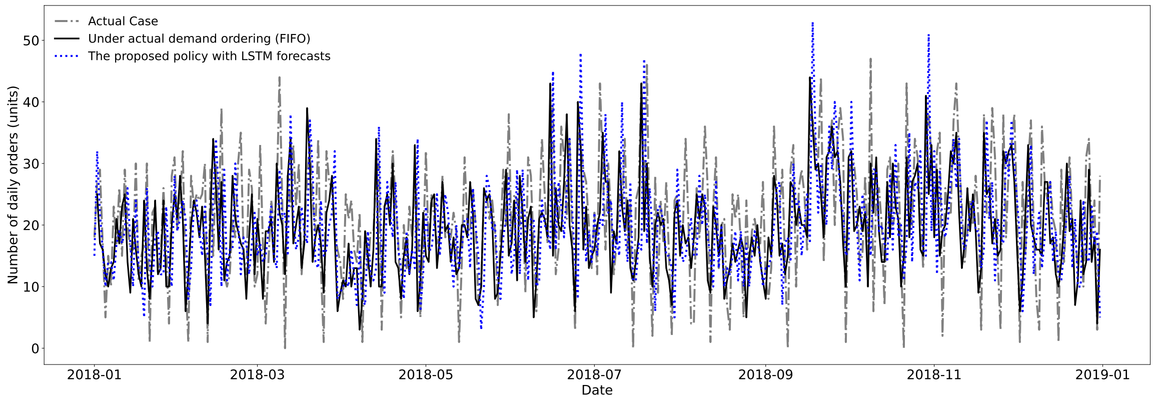

Figure 6.8: Comparison of the ordering quantities for the proposed heuristic using LSTM network, the actual case, and ordering by actual demand (FIFO withdrawal policy)

6.2.3 Sensitivity Analysis

In this section, we explore sensitivity of the proposed model to three different factors, the minimum inventory level, the cost ratios in the cost function, and the remaining shelf life of the product. Sensitivity analysis to the minimum inventory level is done due to the fact that in practice, low inventory levels may trigger manual orders. This is understandable because having a low inventory in case of an emergency can put patients' lives in danger. As a result, it may be desirable to ensure that the inventory remains above a certain value. Sensitivity analysis to the cost ratios shows the generalizability of the proposed model in different systems with different shortage and wastage costs. Last but not least, by doing a sensitivity analysis to the remaining shelf life, we explore how the model behaves when the shelf life of received products differs from our assumption. It is important because the received platelet units may vary with respect to their remaining shelf lives. For all the results presented in this section, we use demand forecasts from the LSTM network.

6.2.3.1 Sensitivity to the Minimum Inventory Level

We first explore sensitivity of the proposed heuristic to the minimum inventory level. No minimum inventory was considered for the results presented in Section 6.1. Here, we consider minimum inventory levels of five, ten and fifteen units when training the model with LSTM forecasts. Table 6.4 gives the results for different minimum inventory values. As we can see from Table 6.4, as the minimum inventory level increases, the number of wastages rises while the number of shortages remains zero. Furthermore, as we expect, by incrementing the minimum inventory level, the mean inventory level increases, which is mainly due to the fact that the order-up-to level is set to a higher value. The mean ordering quantity is, however, quite similar for all three minimum inventory values, being very close to the mean actual demand, 19.49. This implies that when the shortage cost is much higher than the wastage cost, one can ensure a certain minimum inventory level without a significant increase in ordering quantities and associated wastages. This may be helpful in practice to avoid manual interventions as a result of low inventory levels.

Table 6.4: Sensitivity to the minimum inventory level for 2018 using LSTM network forecasts

Minimum Inventory Level	Cost	No. of Shortages	No. of Wastages	Inventory Mean	Inventory Std	Mean Ordering Quantity
5	0	0	0	30.32	6.75	19.40
10	6	0	6	33.30	6.73	19.42
15	9	0	9	34.70	6.73	19.44

6.2.3.2 Sensitivity to the Cost Ratio

In order to explore the generalizability of the proposed heuristic under different shortage and wastage costs, we perform a sensitivity analysis for the cost ratio. The results presented so far are based on a cost ratio of five (shortage to wastage cost). However, in different systems, this ratio may be different. In fact, the proposed model's performance depends on the cost ratio rather than the absolute values of the costs. We consider different cost ratios (shortage to wastage cost) equal to 5, $\frac{1}{5}$, 1, $\frac{1}{10}$ and 10. The results show that the number of shortages and wastages are zero for all the cost ratios. The optimal α and β values are the same for the cost ratios, and consequently the inventory mean and standard deviation are the same (mean[sd]: 27.32[6.74]). The results show that the proposed model is not highly sensitive to the cost ratio, and so changing the cost ratio does not have significant impact on the performance of the policy.

6.2.3.3 Sensitivity to the Remaining Shelf Life

We begin by considering different remaining shelf lives upon arrival. Since platelet products have five to seven days shelf life and spend two days in test and screening processes, the remaining shelf life of the platelets used for transfusion is three to five days. As a result, we investigate the sensitivity of our model to the product's shelf life by considering five, four, three, and two days of remaining shelf life upon arrival.

First, we consider a scenario in which the model is trained for a product with five days of shelf life, but the shelf life of received units can be different from what the model is trained for. Table 6.5 gives the results for different shelf life values by considering a cost ratio of five (shortage to wastage cost). Each row of the table considers a constant shelf life for received units, five, four, three and two days. As the shelf life of platelets decreases, the total cost increases. This is mainly due to the large increase in the number of wastages. As a result of these wastages, the ordering quantity compensates by increasing as the shelf life decreases.

Table 6.5: Sensitivity to the remaining shelf life when the model is trained for five days shelf life and is tested for different shelf lives (for 2018 using LSTM network forecasts)

Remaining Shelf Life	Cost	No. of Shortage	No. of Wastage	Inventory Mean	Inventory Std	Mean Ordering Quantity
Five	0	0	0	27.32	6.74	19.39
Four	16	0	16	27.27	6.71	19.45
Three	253	0	253	26.60	6.45	20.11
Two	1913	0	1913	22.04	5.50	24.65

Second, we consider a scenario for training the model for products with four days remaining shelf life and explore the sensitivity to remaining shelf lives of five, four, three, and two days when updating the inventory. As we can see in Table 6.6, having a shelf life of five days results in no wastage and 20 units of shortage. However, smaller shelf lives result in a large number of wasted units because the order-up-to level is set for a larger shelf life than actual. The number of shortages is the same in all of the scenarios.

Table 6.6: Sensitivity to the remaining shelf life when the model is trained for four days shelf life and is tested for different shelf lives (for 2018 using LSTM network forecasts)

Remaining Shelf Life	Cost	No. of Shortage	No. of Wastage	Inventory Mean	Inventory Std	Mean Ordering Quantity
Five	100	20	0	19.82	6.52	19.32
Four	102	20	2	19.81	6.52	19.33
Three	155	20	55	19.65	6.45	19.49
Two	1076	20	976	17.11	5.65	22.02

Considering the results in the above tables, the main issue of training for constant shelf life is that if the shelf life of received units is smaller than what the model is trained for, the wastage rate increases.

6.2.4 Ordering Frequency

The inventory policy that has been used throughout this paper places an order at the end of every day. In this section the performance of the proposed heuristic when only placing orders at the end of certain days is explored.

Up to this point, different forecasting models have been used, all of them achieving similar performance. For this section, we focus on a lasso model. This model is trained with data where the target for each order is the cumulative demand of all days from the day following the order up to and including the day of the next order. For example, if there is no order on Monday the forecast on Sunday will cover demand for Monday and Tuesday combined. In addition, the model is fed all features from the previous few days before the order, with the exact number of days being the longest period the schedule goes without ordering. For example, if the schedule skips two days in a row, creating a gap of three days, the model will be fed three days worth of features, from the day when an order is placed and the two days before.

The ordering policy is the same as defined in the previous chapter (Chapter 5) but (5.3.1) is modified to be:

$$S_t(\hat{f}_t) = \begin{cases} \alpha \hat{f}_t + \beta & \text{if } weekday(t) \in schedule \\ 0 & otherwise, \end{cases} \tag{6.2.3}$$

where *weekday* gives the day of the week of a period t, and *schedule* is the set of weekdays that the current schedule places orders on:

$$schedule \subset (\{Mo, Tu, We, Th, Fr, Sa, Su\}) \tag{6.2.4}$$

The notation used to denote an ordering schedule is as follows: MoTuWeThFrSaSu is the notation used to represent daily ordering, MoTuThFrSu is the notation for ordering on all days except Saturday and Wednesday, etc. In this section eight schedules are explored, one six-day schedule: MoTuWeThFrSu; four five-day schedules: MoTuWeThFr, TuWeThFrSu, MoTuThFrSu, and MoTuWeThSa; and three four-day schedules: MoTuThSa, MoWeThSa, and TuThFrSu.

MoTuWeThFr is current practice, staff levels are low over the weekend and placing routine orders is not feasible. The other schedules are included to compare the proposed heuristic to a base-stock policy. Table 6.7 shows statistics for schedules both with the proposed heuristic and with a base-stock policy. All schedules see an improvement in performance with the proposed heuristic, with the current schedule half the cost of a base-stock policy. In a number of schedules this improvement is achieved by greatly reducing wastage and slightly increasing shortage. This trade-off can be controlled by changing the shortage to wastage cost ratio, but this is not explored here, only a ratio of 5 is tested. Based on these schedules the base-stock policy and the proposed heuristic struggle when they go two or more days without ordering. This can be seen with the four day schedules, which do better than MoTuWeThFr despite ordering on fewer days, as they keep gaps in the schedule to one day only. The base-stock policies keep around 5-15 more units in inventory than the proposed heuristic with the current schedule keeping 15 fewer units under the proposed heuristic.

Table 6.7: Performance of scheduled order policies

	Schedule	Cost	No. of Shortage	No. of Wastage	Inventory Mean	Inventory Std	Mean Ordering Quantity
Performance of proposed heuristic	MoTuWeThFr	129	20	29	31.18	17.6	28.94
	MoTuWeThFrSu	0	0	0	29.85	9.09	23.56
	TuWeThFrSu	44	7	9	35.84	13.06	28.44
	MoTuThFrSu	20	1	15	31.37	11.17	28.29
	MoTuWeThSa	39	5	14	32.88	12.17	28.31
	MoTuThSa	51	5	26	36.02	13.01	35.31
	MoWeThSa	72	12	12	35.92	14.43	35.91
	TuThFrSu	124	22	14	35.06	12.9	35.23
Performance of base-stock policy	MoTuWeThFr	269	37	84	46.66	13.14	28.31
	MoTuWeThFrSu	29	1	24	35.93	10.68	23.52
	TuWeThFrSu	168	22	58	46.27	12.19	28.35
	MoTuThFrSu	46	0	46	43.85	12.4	28.37
	MoTuWeThSa	41	1	36	42.39	11.11	28.31
	MoTuThSa	53	0	53	41.66	13.98	35.48
	MoWeThSa	78	4	58	42.58	14.09	35.49
	TuThFrSu	173	20	73	44.31	14.6	35.52

Figure 6.9 displays how the number of units ordered each day in a schedule are distributed. Of note is that the proposed heuristic tends to place large orders before a multi-day gap while the base-stock policy places larger orders after a gap. Including forecasts in the proposed heuristic helps smooth over gaps by predicting them and preemptively ordering extra inventory while the base-stock policy keeps more inventory and recovers from the gap after it has passed.

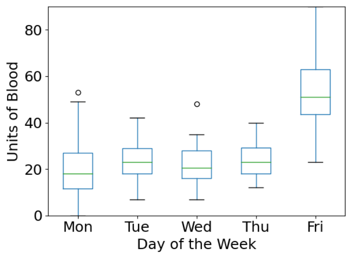

(a) Order quantities for ordering MoTuWeThFr with the proposed heuristic

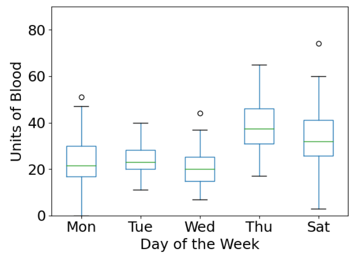

(c) Order quantities for ordering MoTuWeThSa with the proposed heuristic

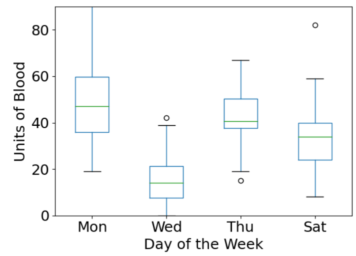

(e) Order quantities for ordering MoWeThSa with the proposed heuristic

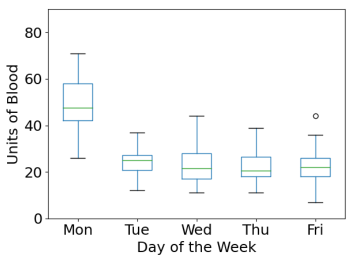

(b) Order quantities for ordering MoTuWeThFr with a base-stock policy

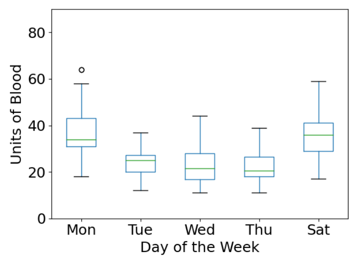

(d) Order quantities for ordering MoTuWeThSa with a base-stock policy

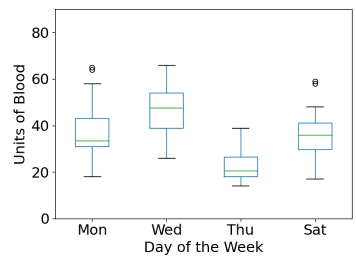

(f) Order quantities for ordering MoWeThSa with a base-stock policy

Figure 6.9: Distribution of order quantities by day for select schedules

6.3 Discussion

In this section, we provide some discussion of our results for the heuristic and when it is trained with a cost function consisting of only shortage and wastage costs.

6.3.1 The Heuristic

As discussed in Chapter 5, calculating the optimal policy may become difficult as the optimal policy is state-dependant, and so the proposed heuristic provides a simpler and more practical alternative. While we cannot compare the proposed heuristic to the optimal policy, the fact that the wastages and shortages are not high suggest that the heuristic is near optimal. We evaluate the heuristic with forecasts generated from two different forecast models to explore the impact of forecasting accuracy on the heuristic. Results show that regardless of the quality of the forecasts, both wastage and shortage rates remain consistently low. Thus, one can still make effective decisions with lower quality forecasts. These findings highlight that the proposed heuristic can be an effective alternative to the optimal policy. The proposed heuristic adopts an affine function of the next time period's demand forecast to set the order-up-to level, which makes it intuitively explainable and easy to implement. We also explored variations of the heuristic that include more forecast history. Numerical results of such cases suggest that the forecast value can be considered as a more complex function of previous history, rather than simply a function of the previous value.

Sensitivity analysis conducted on the product's shelf life shows that both the proposed heuristic and the base-stock policy perform well with products that have a shelf life of 3, 4, and 5 days, but their performance degrades as the shelf life becomes smaller. However, the proposed heuristic shows relatively lower sensitivity to the cost function than the base-stock policy. Additionally, the heuristic has lower on-hand inventory levels than the base-stock policy in all the considered cases. This suggests that by including additional demand information in the inventory model, order quantities close to the actual demand can be achieved while keeping the on-hand inventory low. This approach has several benefits such as low on-hand inventory levels that are beneficial when there is limited inventory capacity or a high holding cost. It also results in fresher inventory and reduces wastage.

6.3.2 Heuristic Trained for Shortages and Wastages

In the second part of this chapter, we evaluated a special case of the heuristic for the platelet transfusion data, when it is trained by only considering shortage and wastage costs. It is trained by using forecasts from three different forecasting models with different forecast accuracy. Results show that due to non-uniqueness of the optimal policy from the training data, it is important to carefully choose the most suitable policy to use for the test data. Besides that, the proposed heuristic is an efficient alternative for the optimal policy even by considering just shortages and wastages in cost function, which not only leads to a zero cost for the platelet data, but keeps the on-hand inventory very close to the actual demand. Moreover, there is no need for the forecasts to be of great quality, even a simple forecasting model with large error can be used. The reason is that the proposed ordering policy has forecast-dependent order-up-to levels that hedge against forecast inaccuracies.

We also consider various ordering frequencies using the proposed ordering policy and the base-stock policy, to explore the possibility of reducing the ordering frequency. Using the proposed heuristic and ordering six days a week, skipping Saturday, does not affect the performance, and can be considered as an alternative choice for daily ordering. Ordering five days a week results in a moderate performance decrease, but still appears to be a feasible option. Results show that when orders are placed less frequently, considering forecast-dependent base-stock levels, large orders are placed before a multi-day gap while with a classic base-stock policy these orders are placed after the gap. That being the case, demand forecasts compensate for demand over gaps by ordering extra inventory before the gap while the policies without forecasts generally keep more inventory and place large orders after the gap is passed. This suggests that if a blood centre were to move to less frequent ordering, it is important to include demand forecasts.

The proposed heuristic offers benefits throughout the supply chain, beyond benefits for the hospitals. Considering forecast-dependent order-up-to levels benefits suppliers of the blood products in several ways. First, incorporating additional information about the actual demand in the inventory management process results in close-to-actual-demand ordering quantities. This not only reduces the wastage rate in hospitals, but increases the transparency of blood utilization between blood suppliers and hospitals which can help blood suppliers in their planning. Second, the on-hand inventory tends to be low, which is not only helpful when there is limited inventory capacity or a high holding cost but results in fresher inventory. Third, the wastage is very low since the order-up-to level varies as the predicted demand (and consequently the actual demand) changes. By comparing the actual mean inventory (42.60) and the proposed heuristic's mean inventory using LSTM network forecasts (26.96), we see that the on-hand inventory is reduced significantly. One of the consequences of this is the wastage rate being reduced from 9% to zero.

The current inventory management practice is very similar to a base-stock policy, which is simple and easy to implement, but can result in large on-hand inventory, high wastage rate, and frequent shortages. The current base-stock levels being used are very high, causing the current high wastage rate. According to the results presented here, a base-stock policy can be effective, but the base-stock level should be chosen in a data-driven manner. Base-stock levels in a base-stock policy can be reduced using data, and one can further improve the inventory model by including demand forecasts in the model. Since platelet products are very expensive ($504 per unit), and hospitals have limited inventory capacities to store them, using a smart ordering policy (forecast, then order) can considerably reduce costs. Furthermore, forecasting the demand can help blood centres to make decisions in other parts of a blood supply chain, such as the clinical management for staff and resource allocation. Thus, by incorporating demand forecasts in the inventory management problem, specifically for products with short shelf lives, one can benefit from both the advantages of demand forecasting and including additional demand information in the inventory management process. As the results show, there is no need for high quality, highly accurate forecasts for generating the forecast-dependent base-stock levels, even a simple forecasting model is sufficient for incorporating additional information in the inventory model.

6.4 Conclusion

In this chapter, we evaluated a data-driven inventory management policy for platelet products, which can be generalized to perishable products. The data we used for evaluation is for platelet transfusion in Hamilton hospitals. The numerical results showed that even with the forecasts not being highly accurate, the proposed heuristic resulted in a cost close to the policies that consider history. It achieved significant improvements over the actual wastages in the data. Moreover, in comparison to the base-stock policy, the heuristic resulted in lower shortage rates while keeping the on-hand inventory much lower. Overall, it achieved a 3.45% reduction in the total cost compared to the base-stock model. Also, by adding more history of demand forecasts in the order-up-to-level function, the wastage cost remained zero, and the number of shortages decreased slightly. The total cost improvement for the case when the current and four previous forecasts were included in the inventory model was 2.21%. Considering the negligible reduction in the total cost when one, two, three, and four previous forecasts were included in the order-up-to-level function, we concluded that it is not necessary to include a large amount of forecast history to calculate the order-up-to levels and only considering the current forecast is sufficient.

We also considered a cost function consisting of only shortage and wastage costs. Experimental results show that while our model does not require highly accurate forecasts, it results in a lower inventory, resulting in a reduction of the wastage rate by 9%. It also reduces same-day urgent orders by 14%. Furthermore, the proposed ordering policy improves the transparency between blood suppliers and hospitals which has the potential to improve the overall efficiency of the platelet blood supply chain. We also explored different ordering frequencies for the proposed model and the base-stock policy. For this particular application, reducing the ordering frequency to five days a week is feasible and does not impact the performance. Moreover, based on the presented results, incorporating forecasts in the inventory model increases the viability of reduced ordering frequencies.

Chapter 7

Conclusions and Future Work

In this final chapter, we summarize the work presented in previous chapters, re-state the major contributions of this dissertation and state directions for future work that are based on combinations of ideas from Chapters 4 and 5.

7.1 Summary and Contributions

In this dissertation, we utilized five forecasting models to forecast demand for perishable blood products at hospitals. The forecasting models included two univariate models, ARIMA and Prophet, and two multivariate models, lasso regression, random forest, and LSTM networks. We compared these models in terms of forecast accuracy, different re-training periods, and the amount of data used for training. Next, we used these forecasts as additional demand information in the inventory model with fixed costs to determine the optimal ordering policy. Lastly, based on structural results for the optimal policy, we proposed a heuristic that is a variation of the base-stock policy. The heuristic considered the order-up-to levels as an affine function of the current period's forecasts.

7.2 Future Work

In order to further progress in the development of models of real perishable inventory problems, in future work it is necessary to study:

1. more systematic ways for selecting predictors for demand forecasting models
2. inventory management models under uncertainty in the remaining shelf life of products or lead time
3. optimality of the ordering policy for a more general forecast function, under different issuing policies, and for worst-case scenarios
4. generalizability of the demand forecasting and inventory management models

We will provide a more comprehensive explanation of these topics in the subsequent sections.

7.2.1 Selection of Predictors for Demand Forecasting

In Chapter 4, we utilized multivariate forecasting models that use predictors for forecasting the demand. For the data used in this study, predictors are selected based on lasso regression. We also add additional predictors that affect the platelet demand based on the exploratory analysis done for detecting trends, seasonality or holiday patterns. We see this step as necessary for selecting the most appropriate predictors for the forecasting model. However, this approach might not always be practical, necessitating a more automated solution for predictor selection. Another issue is that selected predictors should be interpretable to allow users to determine the effect of different predictors on demand values.

This is not only helpful for more robust demand prediction, but can help the decision makers/suppliers to plan for the future. Thus, we see an important next step as further exploring the lasso regression approach to enhance variable selection, with a particular focus on interpretability. This will not only affect the lasso regression itself, but also may improve other multivariate forecasting models such as LSTM forecasting since we selected the LSTM inputs using lasso regression. Moreover, it would be interesting to apply other variable-selection techniques such as recursive feature selection, stepwise forward and backward regression, genetic algorithms, or Support Vector Machines (SVM). While lasso regression works well for this data, it may be of benefit to explore nonlinearity or multicollinearity among features for other datasets.

7.2.2 Modelling Uncertainty

For the inventory management part of this study, we consider perishable products with a fixed shelf life. Also, we assume that received products are fresh products with a remaining shelf life equal to the product's shelf life. While this is a reasonable assumption, in real problems, parameters like the remaining shelf life of received units or the lead time could be subject to uncertainty. Thus, we see as a next step the examination of worst-case scenarios when variability in the remaining shelf life of the products is ignored. This can be helpful to assess the impact of that ignoring uncertainty in the inventory model. Secondly, it would be interesting to explore ordering policies that are robust to the uncertainty in the remaining shelf life or lead time. For example, see the work of Abouee-Mehrizi et al. (2022) for uncertainty in the remaining shelf life and the work of Hansen et al. (2023) for uncertainty in the lead time. A challenge for future work, then, is to model the effect of such uncertainties on ordering decisions and the optimal policy.

7.2.3 Further Investigation of the Optimal Ordering Policy

In Chapter 5, we explored the structural properties of the optimal ordering policy when demand forecasts are included in the inventory model. The forecasts are considered as an affine function of previous forecast and (indirectly) demand values, like in an ARIMA model. However, in real systems, forecasts from models such as an ANN may be used that do not follow this linearity assumption. In future work, we would therefore like to investigate more complex forecasts such as nonlinear, nonconvex functions and explore how they affect the ordering decisions.

In a perishable inventory system, assuming a FIFO ordering policy is a common assumption due to the limited shelf life of products. For products with a short shelf life, this assumption becomes even more reasonable. However, certain systems, like the platelet transfusion system, may require products with varying levels of freshness. In the red blood cell transfusion system, due to a requirement for matching blood types, demand is satisfied from different product groups. Additionally, the presence of priority demand classes within the system highlights the significance of incorporating other issuing policies, such as LIFO (Last In, First Out) or priority queues. This would contribute significantly to achieving the objective of effectively managing realistic perishable inventory systems.

Another future extension of this work would involve studying the asymptotic optimality of the inventory model when demand forecasts are incorporated into the model. Given the complexity of large scale systems, analyzing properties of the system as it grows in size in one or more aspects is of interest (Goldberg et al., 2016; Xin and Goldberg, 2016; Zhang et al., 2020; Bu et al., 2023). This would render the inventory model applicable to larger systems or scenarios where one parameter becomes large. Specifically asymptotic optimality results by considering the following asymptotic regimes of the system parameters are interesting for this problem: (i) large demand size, (ii) large per unit shortage cost, (iii) large per unit wastage cost, and (iv) large per unit holding cost.

In Chapter 5, we proposed a heuristic motivated by the structure of the optimal policy. Our results show that the heuristic is an efficient alternate for the optimal policy. It would be interesting to calculate the gap between the optimal policy and the proposed heuristic to formally demonstrate robustness of the results.

7.2.4 Exploring Generalizability of Forecasting and Inventory Models

For the demand forecasting part of this research, we use data for platelet transfusion in Hamilton hospitals. We observe that multivariate models that include additional predictors in the model perform more efficiently than univariate models that only use previous demand values. In future work, it would be of interest to utilize these models on different platelet transfusion datasets (outside of Hamilton) to explore generality of the results. Moreover, it is of interest to perform time series analysis and train the models for general perishable products. This can be helpful not only for the demand forecasting part, but also can be useful for investigating effect of different forecast accuracies on the proposed heuristic. Another direction for future work is to apply and evaluate the proposed heuristic on different datasets to explore the generalizability of the heuristic.

As discussed in Chapter 2, there are two general data-driven inventory management approaches, (1) prediction and optimization as a single step, (2) predict then optimize (the work in this study). In Chapter 5, we provide a brief comparison between these two models and the reason that the second approach is used in this research. However, a more extensive comparison of the two approaches would be of interest.

7.3 Conclusions

The central thesis of this dissertation is that by including additional demand information in the inventory model, we can better model, understand and solve real-world perishable inventory problems. In the presence of abundant data, one can benefit from using historical data to estimate future demand values which can in turn be used in the inventory model for determining order quantities. The work of this dissertation followed a similar manner as the data-driven predict then optimize approach. Firstly, a holistic data analysis was done and the demand was predicted. This is helpful not only for decision makers to make better-informed decisions, but can be used as added information for determining future ordering quantities. Next, the structure of the optimal policy when these forecasts are incorporated in the inventory model was explored. Lastly, motivated by the structure of the optimal ordering policy, a forecast-dependent base-stock policy was proposed. We presented an extensive evaluation of the heuristic. Experimental results showed that while the heuristic is simple and intuitive, it worked well for different forecasting accuracies. Thus, it can be considered as an effective alternative for the optimal policy. To our knowledge, this dissertation is the first work that studies structure of the optimal policy when demand forecasts are incorporated in the inventory model.

Appendix A

Supplement to Chapter 3

The data cleaning process is an essential step aimed at extracting relevant information from the TRUST database. The TRUST database consists of various tables for different blood transfusion products. We extract data from the inventory table within the database that holds information regarding transfusions. Each row in this table corresponds to a single blood product unit. We extract more than 200 variables that can be grouped into four main categories:

- Product data: product name, product ABO group, Rh type, collection date and time, issue date and time, issue volume, expiry date
- Patient data: gender, age, ABO group, Rh type
- Admission data: admission date and time, discharge date and time, transfusion location
- Laboratory test data: Complete Blood Count (CBC) such as RBC count, hemoglobin count, platelet count, hematocrit count

Our data collection and cleaning procedure can be divided into two main phases. Initially, we focus on retrieving data specific to platelet transfusions. Next, we collaborate with medical professionals to identify and select the relevant variables from the platelet transfusion data. We consider more than 100 potential clinical variables during the selection process for model variables. Variables with more than 70% missing values are excluded from variable selection. The techniques used for handling missing values in individual variables rely on both clinical definitions and the practical application of the respective variable. For instance, in cases a specific test such as daily platelet count lacks its value, we make the assumption that the patient had normal platelet counts on that specific day. Furthermore, we employ the min-max method to normalize the variables. Finally, the data is aggregated to a daily dataset using a structured data processing framework (Li et al., 2022).

Furthermore, our discussions with physicians have revealed that most of the outliers in the dataset are not random occurrences; they have specific underlying reasons. For instance, we note some extremely large values for transfusion volumes. These outliers arise from situations where small bags are used for transfusions but are recorded as standard transfusions in the inventory database. To rectify this, we standardize these values accordingly.

Another data-related concern involves the lack of consistent data types for certain variables across different entries. For instance, when examining specific laboratory test results, we encounter both integer values and comparative values (e.g., UREA ≤ 0.5). To address this inconsistency, we standardize these values by referencing either the maximum or minimum allowable value for the test or by utilizing the mean value for the test, depending on the specific test itself. Other cases involve data types that are initially integers but include non-integer values (e.g., 300ML for transfusion volume); we opt to retain only the integer part while discarding the non-integer components.

Appendix B

Supplement to Chapter 4

Table B.1 gives the selected predictors using lasso regression. Considering the coefficients for the predictors and their corresponding confidence intervals in Table B.1, and based on (Ranstam, 2012), variables that have a coefficient of zero and confidence intervals that are symmetric around zero are candidates to be eliminated. As we can see from Table B.1, abnormal_plt has the highest coefficient. The predictors abnormal_hb and abnormal_redcellwidth can be considered as two other important lab tests for forecasting the demand. Day of the week, last week's platelet usage and yesterday's platelet usage also have notable impact on the platelet demand. As we can see in Table B.1, unexpectedly, some of the predictors have a negative coefficient in the demand forecasting model. The reason is that, as we can see from Figure 3.5, there are high correlations among the predictors that result in interactions among the model predictors, which may cause multicollinearity issues. Specifically, the predictors abnormal_hb, abnormal_INR, abnormal_hematocrit, and abnormal_MPV are correlated with abnormal_plt. The predictors abnormal_hematocrit and abnormal_hb also have high correlations with most of the other abnormal laboratory test results. The coefficients for lab tests are high. This is consistent with the observation that

Table B.1: Predictors and their corresponding coefficients for lasso regression

predictors	**Coefficients**	95% **Confidence Interval**
abnormal_ALP	-0.02	(-0.08 , 0.04)
abnormal_MPV	0.01	(-0.06 , 0.11)
abnormal_hematocrit	0.00	(-0.11 , 0.14)
abnormal_PO2	-0.11	(-0.19 , 0.00)
abnormal_creatinine	0.03	(-0.03 , 0.11)
abnormal_INR	0.06	(-0.02 , 0.22)
abnormal_MCHb	-0.03	(-0.10 , 0.04)
abnormal_MCHb_conc	-0.03	(-0.10 , 0.04)
abnormal_hb	0.05	(-0.04 , 0.19)
abnormal_mcv	-0.03	(-0.11 , 0.04)
abnormal_plt	0.23	(0.02 , 0.36)
abnormal_redcellwidth	0.07	(0.00 , 0.15)
abnormal_wbc	-0.02	(-0.09 , 0.03)
abnormal_ALC	0.01	(-0.05 , 0.08)
location_GeneralMedicine	-0.11	(-0.21 , 0.00)
location_Hematology	0.04	(-0.02 , 0.16)
location_IntensiveCare	0.05	(-0.01 , 0.15)
location_CardiovascularSurgery	0.04	(-0.03 , 0.11)
location_Pediatric	0.04	(-0.02 , 0.10)
Monday	0.07	(0.00 , 0.16)
Tuesday	0.07	(0.00 , 0.14)
Wednesday	0.00	(-0.04 , 0.07)
Thursday	0.01	(-0.03 , 0.09)
Friday	-0.39	(-0.46 , -0.31)
Saturday	-0.31	(-0.39 , -0.23)
Sunday	0.10	(0.03 , 0.18)
lastWeek_Usage	0.12	(0.05 , 0.19)
yesterday_Usage	0.10	(0.02 , 0.17)
yesterday_ReceivedUnits	0.06	(0.00 , 0.14)

the lab test results are significant indicators for platelet transfusion. The predictors abnormal_plt, abnormal_hb, abnormal_ALC and abnormal_wbc have higher coefficients and consequently higher impact on platelet demand. For day of the week, Friday and Saturday have negative coefficients due to the fact that they cover the weekend (Friday: -0.39 and Saturday: -0.31). For hospital census data, except for location_GeneralMedicine, all the coefficients are in a similar range to the lab tests.

Bibliography

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., et al. (2016). Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.

Abdulwahab, U. and Wahab, M. (2014). Approximate dynamic programming modeling for a typical blood platelet bank. *Computers & Industrial Engineering*, 78:259–270.

Abouee-Mehrizi, H., Mirjalili, M., and Sarhangian, V. (2022). Data-driven platelet inventory management under uncertainty in the remaining shelf life of units. *Production and Operations Management*, 31(10):3914–3932.

Abouee-Mehrizi, H., Mirjalili, M., and Sarhangian, V. (2023). Platelet inventory management with approximate dynamic programming. *arXiv preprint arXiv:2307.09395*.

Alcaina, P. S. (2020). Platelet transfusion: and update on challenges and outcomes. *Journal of Blood Medicine*, 11:19.

Amiri, S. A. H. S., Zahedi, A., Kazemi, M., Soroor, J., and Hajiaghaei-Keshteli, M. (2020). Determination of the optimal sales level of perishable goods in a two-echelon supply chain network. *Computers & Industrial Engineering*, 139:106156.

Attari, M. Y. N., Pasandideh, S. H. R., Aghaie, A., and Niaki, S. T. A. (2018). A bi-objective robust optimization model for a blood collection and testing problem: an accelerated stochastic benders decomposition. *Annals of Operations Research*, pages 1–39.

Ban, G.-Y. and Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1):90–108.

Bar-Lev, S. K., Boxma, O., Mathijsen, B., and Perry, D. (2017). A blood bank model with perishable blood and demand impatience. *Stochastic Systems*, 7(2):237–263.

Baron, O., Berman, O., and Perry, D. (2010). Continuous review inventory models for perishable items ordered in batches. *Mathematical Methods of Operations Research*, 72:217–247.

Baron, O., Berman, O., and Perry, D. (2020). Continuous review inventory models for perishable items with leadtimes. *Probability in the Engineering and Informational Sciences*, 34(3):317–342.

Barron, Y. (2019). A state-dependent perishability (s, S) inventory model with random batch demands. *Annals of Operations Research*, 280:65–98.

Barron, Y. and Baron, O. (2020). QMCD approach for perishability models: The (S, s) control policy with lead time. *IISE Transactions*, 52(2):133–150.

Beamon, B. M. (1998). Supply chain design and analysis:: Models and methods. *International journal of production economics*, 55(3):281–294.

Beliën, J. and Forcé, H. (2012). Supply chain management of blood products: A literature review. *European Journal of Operational Research*, 217(1):1–16.

Benjaafar, S., Cooper, W. L., and Mardan, S. (2011). Production-inventory systems with imperfect advance demand information and updating. *Naval Research Logistics (NRL)*, 58(2)(2):88–106.

Berk, E. and Gürler, Ü. (2008). Analysis of the (Q, r) inventory model for perishables with positive lead times and lost sales. *Operations Research*, 56(5):1238–1246.

Bernstein, F., Li, Y., and Shang, K. (2016). A simple heuristic for joint inventory and pricing models with lead time and backorders. *Management Science*, 62(8):2358–2373.

Blake, J. T. (2009). On the use of operational research for managing platelet inventory and ordering. *Transfusion*, 49(3):396–401.

Blake, J. T. (2010). An introduction to platelet inventory and ordering problems. *Wiley Encyclopedia of Operations Research and Management Science*.

Boulmaiz, T., Guermoui, M., and Boutaghane, H. (2020). Impact of training data size on the LSTM performances for rainfall–runoff modeling. *Modeling Earth Systems and Environment*, 6:2153–2164.

Brodheim, E., Derman, C., and Prastacos, G. (1975). On the evaluation of a class of inventory policies for perishable products such as blood. *Management Science*, 21(11):1320–1325.

Bruno, G., Diglio, A., Piccolo, C., and Cannavacciuolo, L. (2019). Territorial reorganization of regional blood management systems: Evidences from an Italian case study. *Omega*, 89:54–70.

Bu, J., Gong, X., and Chao, X. (2023). Asymptotic optimality of base-stock policies for perishable inventory systems. *Management Science*, 69(2):846–864.

Chao, X., Gong, X., Shi, C., Yang, C., Zhang, H., and Zhou, S. X. (2018). Approximation algorithms for capacitated perishable inventory systems with positive lead times. *Management Science*, 64(11):5038–5061.

Chao, X., Gong, X., Shi, C., and Zhang, H. (2015). Approximation algorithms for perishable inventory systems. *Operations Research*, 63(3):585–601.

Chen, S., Li, Y., Yang, Y., and Zhou, W. (2021). Managing perishable inventory systems with age-differentiated demand. *Production and Operations Management*, 30(10):3784–3799.

Chen, S., Li, Y., and Zhou, W. (2019). Joint decisions for blood collection and platelet inventory control. *Production and Operations Management*, 28(7):1674–1691.

Chen, X. (2017). $L^\natural$-convexity and its applications in operations. *Frontiers of Engineering Management*, 4(3):283–294.

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

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

Chen, X. and Simchi-Levi, D. (2004). Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Operations Research*, 52(6):887–896.

Cheung, Y.-W. and Lai, K. S. (1995). Lag order and critical values of the augmented Dickey-Fuller test. *Journal of Business & Economic Statistics*, 13(3):277–280.

Civelek, I., Karaesmen, I., and Scheller-Wolf, A. (2015). Blood platelet inventory management with protection levels. *European Journal of Operational Research*, 243(3):826–838.

Clarkson, J., Voelkel, M. A., Sachs, A.-L., and Thonemann, U. W. (2023). The periodic review model with independent age-dependent lifetimes. *Production and Operations Management*, 32(3):813–828.

Cohen, M. A. (1976). Analysis of single critical number ordering policies for perishable inventories. *Operations Research*, 24(4):726–741.

Cohen, M. A. and Pierskalla, W. P. (1979). Target inventory levels for a hospital blood bank or a decentralized regional blood banking system. *Transfusion*, 19(4):444–454.

Cooper, W. L. (2001). Pathwise properties and performance bounds for a perishable inventory system. *Operations Research*, 49(3):455–466.

Critchfield, G. C., Connelly, D. P., Ziehwein, M. S., Olesen, L. S., Nelson, C. E., and Scott, E. P. (1985). Automatic prediction of platelet utilization by time series analysis in a large tertiary care hospital. *American Journal of Clinical Pathology*, 84(5):627–631.

Croson, R. and Donohue, K. (2006). Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science*, 52(3):323–336.

Da Veiga, C. P., Da Veiga, C. R. P., Catapan, A., Tortato, U., and Da Silva, W. V. (2014). Demand forecasting in food retail: A comparison between the holt-winters and arima models. *WSEAS Transactions on Business and Economics*, 11(1):608–614.

Dejonckheere, J., Disney, S. M., Lambrecht, M. R., and Towill, D. R. (2004). The impact of information enrichment on the bullwhip effect in supply chains: A control engineering perspective. *European Journal of Operational Research*, 153(3):727–750.

Deniz, B., Karaesmen, I., and Scheller-Wolf, A. (2010). Managing perishables with substitution: Inventory issuance and replenishment heuristics. *Manufacturing & Service Operations Management*, 12(2):319–329.

Dillon, M., Oliveira, F., and Abbasi, B. (2017). A two-stage stochastic programming model for inventory management in the blood supply chain. *International Journal of Production Economics*, 187:27–41.

Drackley, A., Newbold, K. B., Paez, A., and Heddle, N. (2012). Forecasting Ontario's blood supply and demand. *Transfusion*, 52(2):366–374.

Du, X. F., Leung, S. C., Zhang, J. L., and Lai, K. K. (2013). Demand forecasting of perishable farm products using support vector machine. *International Journal of Systems Science*, 44(3):556–567.

Duong, L. N. K., Wood, L. C., and Wang, W. Y. (2018). A review and reflection on inventory management of perishable products in a single-echelon model. *International Journal of Operational Research*, 31(3):313–329.

Efron, B. and Tibshirani, R. J. (1994). *An introduction to the bootstrap*. CRC press.

Elmachtoub, A. N. and Grigas, P. (2022). Smart "predict, then optimize". *Management Science*, 68(1):9–26.

Ensafian, H. and Yaghoubi, S. (2017). Robust optimization model for integrated procurement, production and distribution in platelet supply chain. *Transportation Research Part E: Logistics and Transportation Review*, 103:32–55.

Ensafian, H., Yaghoubi, S., and Yazdi, M. M. (2017). Raising quality and safety of platelet transfusion services in a patient-based integrated supply chain under uncertainty. *Computers & Chemical Engineering*, 106:355–372.

Eskandari, M. and Sharifi, O. (2018). Designing efficient spoof detection scheme for face biometric. In *Image and Signal Processing: 8th International Conference, ICISP 2018, Cherbourg, France, July 2-4, 2018, Proceedings 8*, pages 427–434. Springer.

Everitt, H., van der Werf, P., Seabrook, J. A., Wray, A., and Gilliland, J. A. (2022). The quantity and composition of household food waste during the covid-19 pandemic: A direct measurement study in canada. *Socio-economic Planning Sciences*, 82:101110.

Fang, F., Nguyen, T.-D., and Currie, C. S. (2021). Joint pricing and inventory decisions for substitutable and perishable products under demand uncertainty. *European Journal of Operational Research*, 293(2):594–602.

Fanoodi, B., Malmir, B., and Jahantigh, F. F. (2019). Reducing demand uncertainty in the platelet supply chain through artificial neural networks and ARIMA models. *Computers in Biology and Medicine*, 113:103415.

Fontaine, M. J., Chung, Y. T., Rogers, W. M., Sussmann, H. D., Quach, P., Galel, S. A., Goodnough, L. T., and Erhun, F. (2009). Improving platelet supply chains through collaborations between blood centers and transfusion services. *Transfusion*, 49(10):2040–2047.

Fortsch, S. M. and Khapalova, E. A. (2016). Reducing uncertainty in demand for blood. *Operations Research for Health Care*, 9:16–28.

Fortsch, S. M. and Perera, S. (2018). A resilient donor arrival policy for blood. *Operations Research for Health Care*, 19:165–174.

Frankfurter, G. M., Kendall, K. E., and Pegels, C. C. (1974). Management control of blood through a short-term supply-demand forecast system. *Management Science*, 21(4):444–452.

Fries, B. E. (1975). Optimal ordering policy for a perishable commodity with fixed lifetime. *Operations Research*, 23(1):46–61.

Fu, K., Gong, X., and Liang, G. (2019). Managing perishable inventory systems with product returns and remanufacturing. *Production and Operations Management*, 28(6):1366–1386.

Glasserman, P. and Tayur, S. (1994). The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy. *Operations Research*, 42(5):913–925.

Goldberg, D. A., Katz-Rogozhnikov, D. A., Lu, Y., Sharma, M., and Squillante, M. S. (2016). Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *Mathematics of Operations Research*, 41(3):898–913.

Goldberg, D. A., Reiman, M. I., and Wang, Q. (2021). A survey of recent progress in the asymptotic analysis of inventory systems. *Production and Operations Management*, 30(6):1718–1750.

Gong, X. and Chao, X. (2013). Optimal control policy for capacitated inventory systems with remanufacturing. *Operations Research*, 61(3):603–611.

Guan, L., Tian, X., Gombar, S., Zemek, A. J., Krishnan, G., Scott, R., Narasimhan, B., Tibshirani, R. J., and Pham, T. D. (2017). Big data modeling to predict platelet usage and minimize wastage in a tertiary care system. *Proceedings of the National Academy of Sciences*, 114(43):11368–11373.

Gunpinar, S. and Centeno, G. (2015). Stochastic integer programming models for reducing wastages and shortages of blood products at hospitals. *Computers & Operations Research*, 54:129–141.

Gürler, Ü. and Özkaya, B. Y. (2008). Analysis of the (s, S) policy for perishables with a random shelf life. *IIE Transactions*, 40(8):759–781.

Haijema, R. (2011). Optimal issuing of perishables with a short fixed shelf life. In *Computational Logistics: Second International Conference, ICCL 2011, Hamburg, Germany, September 19-22, 2011. Proceedings 2*, pages 160–169. Springer.

Haijema, R. (2013). A new class of stock-level dependent ordering policies for perishables with a short maximum shelf life. *International Journal of Production Economics*, 143(2):434–439.

Haijema, R. (2014). Optimal ordering, issuance and disposal policies for inventory management of perishable products. *International Journal of Production Economics*, 157:158–169.

Hajek, B. (1985). Extremal splittings of point processes. *Mathematics of Operations Research*, 10(4):543–556.

Hamdan, B. and Diabat, A. (2019). A two-stage multi-echelon stochastic blood supply chain problem. *Computers & Operations Research*, 101:130–143.

Hamdan, B. and Diabat, A. (2020). Robust design of blood supply chains under risk of disruptions using Lagrangian relaxation. *Transportation Research Part E: Logistics and Transportation Review*, 134:101764.

Hansen, O., Transchel, S., and Friedrich, H. (2023). Replenishment strategies for lost sales inventory systems of perishables under demand and lead time uncertainty. *European Journal of Operational Research*, 308(2):661–675.

Hastie, T. and Tibshirani, R. (1987). Generalized additive models: some applications. *Journal of the American Statistical Association*, 82(398):371–386.

Ho, T. K. (1995). Random decision forests. In *Proceedings of the Third International Conference on Document Analysis and Recognition (Volume 1) - Volume 1*, ICDAR '95, page 278, USA. IEEE Computer Society.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.

Hospital Liaison Specialists (2020). Canadian Blood Services. `https://www.blood.ca/en/hospital-services/customer-service/hospital-liaison-specialists`. Accessed Date: 2020-11-01.

Huber, J., Gossmann, A., and Stuckenschmidt, H. (2017). Cluster-based hierarchical demand forecasting for perishable goods. *Expert Systems with Applications*, 76:140–151.

Huber, J., Müller, S., Fleischmann, M., and Stuckenschmidt, H. (2019). A data-driven newsvendor problem: From data to decision. *European Journal of Operational Research*, 278(3):904–915.

Huber, J. and Stuckenschmidt, H. (2020). Daily retail demand forecasting using machine learning with emphasis on calendric special days. *International Journal of Forecasting*, 36(4):1420–1438.

Jemai, J., Do Chung, B., and Sarkar, B. (2020). Environmental effect for a complex green supply-chain management to control waste: A sustainable approach. *Journal of cleaner production*, 277:122919.

Jennings, J. B. (1973). Blood bank inventory control. *Management Science*, 19(6):637–645.

Kalpakam, S. and Sapna, K. (1994). Continuous review (s, S) inventory system with random lifetimes and positive leadtimes. *Operations Research Letters*, 16(2):115–119.

Kara, A. and Dogan, I. (2018). Reinforcement learning approaches for specifying ordering policies of perishable inventory systems. *Expert Systems with Applications*, 91:150–158.

Karaesmen, I. Z., Scheller-Wolf, A., and Deniz, B. (2011). Managing perishable and aging inventories: review and future research directions. *Planning Production and Inventories in the Extended Enterprise: A State of the Art Handbook, Volume 1*, pages 393–436.

Kaur, M. N., Arnold, D. M., Heddle, N. M., Cook, R. J., Hsia, C., Blostein, M., Jamula, E., Sholzberg, M., Lin, Y., Kassis, J., et al. (2022). Cost-effectiveness of eltrombopag vs intravenous immunoglobulin for the perioperative management of immune thrombocytopenia. *Blood Advances*, 6(3):785–792.

Keskin, N. B., Min, X., and Song, J.-S. J. (2021). The nonstationary newsvendor: Data-driven nonparametric learning. *Available at SSRN 3866171*.

Ketzenberg, M., Gaukler, G., and Salin, V. (2018). Expiration dates and order quantities for perishables. *European Journal of Operational Research*, 266(2):569–584.

Khaldi, R., El Afia, A., Chiheb, R., and Faizi, R. (2017). Artificial neural network based approach for blood demand forecasting: Fez transfusion blood center case study. In *Proceedings of the 2nd International Conference on Big Data, Cloud and Applications*, pages 1–6.

Khalilpourazari, S. and Hashemi Doulabi, H. (2022). A flexible robust model for blood supply chain network design problem. *Annals of Operations Research*, pages 1–26.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kouki, C., Babai, M. Z., Jemai, Z., and Minner, S. (2016). A coordinated multi-item inventory system for perishables with random lifetime. *International Journal of Production Economics*, 181:226–237.

Kouki, C., Jemaï, Z., and Minner, S. (2015). A lost sales (r, Q) inventory control model for perishables with fixed lifetime and lead time. *International Journal of Production Economics*, 168:143–157.

Kouki, C., Jemai, Z., Sahin, E., and Dallery, Y. (2014). Analysis of a periodic review inventory control system with perishables having random lifetime. *International Journal of Production Research*, 52(1):283–298.

Kouki, C., Legros, B., Babai, M. Z., and Jouini, O. (2020). Analysis of base-stock perishable inventory systems with general lifetime and lead-time. *European Journal of Operational Research*, 287(3):901–915.

Kumar, A., Mhaskar, R., Grossman, B. J., Kaufman, R. M., Tobian, A. A., Kleinman, S., Gernsheimer, T., Tinmouth, A. T., Djulbegovic, B., and Panel, A. P. T. G. (2015). Platelet transfusion: a systematic review of the clinical evidence. *Transfusion*, 55(5):1116–1127.

Kumari, D. and Wijayanayake, A. (2016). An efficient inventory model to reduce the wastage of blood in the national blood transfusion service. In *2016 Manufacturing & Industrial Engineering Symposium (MIES)*, pages 1–4. IEEE.

Lai, G., Chang, W.-C., Yang, Y., and Liu, H. (2018). Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 95–104.

Larimi, N. G., Yaghoubi, S., and Hosseini-Motlagh, S.-M. (2019). Itemized platelet supply chain with lateral transshipment under uncertainty evaluating inappropriate output in laboratories. *Socio-Economic Planning Sciences*, 68:100697.

Lestari, F., Anwar, U., Nugraha, N., and Azwar, B. (2017). Forecasting demand in blood supply chain (case study on blood transfusion unit). In *Proceedings of the World Congress on Engineering*, volume 2.

Levi, R., Perakis, G., and Uichanco, J. (2015). The data-driven newsvendor problem: New bounds and insights. *Operations Research*, 63(6):1294–1306.

Li, N., Arnold, D. M., Down, D. G., Barty, R., Blake, J., Chiang, F., Courtney, T., Waito, M., Trifunov, R., and Heddle, N. M. (2022). From demand forecasting to inventory ordering decisions for red blood cells through integrating machine learning, statistical modeling, and inventory optimization. *Transfusion*, 62(1):87–99.

Li, N., Chiang, F., Down, D. G., and 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*, 29:100290.

Li, Q. and Yu, P. (2014). Multimodularity and its applications in three stochastic dynamic inventory problems. *Manufacturing & Service Operations Management*, 16(3):455–463.

Li, Q., Yu, P., and Wu, X. (2016). Managing perishable inventories in retailing: Replenishment, clearance sales, and segregation. *Operations Research*, 64(6):1270–1284.

Lipton, Z. C., Kale, D. C., Elkan, C., and Wetzel, R. (2015). Learning to diagnose with LSTM recurrent neural networks. *arXiv preprint arXiv:1511.03677*.

Liu, L. and Lian, Z. (1999a). Continuous review models for products with fixed lifetimes. *Operations Research*, 47(1):150–158.

Liu, L. and Lian, Z. (1999b). (s, S) continuous review models for products with fixed lifetimes. *Operations Research*, 47(1):150–158.

Liu, L. and Shi, D.-H. (1999). An (s, S) model for inventory with exponential lifetimes and renewal demands. *Naval Research Logistics (NRL)*, 46(1):39–56.

Liu, Y., Feng, Y., and Lai, G. (2019). Analysis of optimal inventory management for dual blood products with dual sources. *Available at SSRN 3323514*.

Lu, Y. and Song, J.-S. (2005). Order-based cost optimization in assemble-to-order systems. *Operations Research*, 53(1):151–169.

Mahmoodi, A., Haji, A., and Haji, R. (2016). A two-echelon inventory model with perishable items and lost sales. *Scientia Iranica*, 23(5):2277–2286.

Minner, S. and Transchel, S. (2010). Periodic review inventory-control for perishable products under service-level constraints. *OR Spectrum*, 32:979–996.

Mor, R. S., Jaiswal, S. K., Singh, S., and Bhardwaj, A. (2019). Demand forecasting of the short-lifecycle dairy products. *Understanding the Role of Business Analytics: Some Applications*, pages 87–117.

Moshtagh, M. S., Zhou, Y., and Verma, M. (2022). Coordinating a decentralized blood supply chain with interactions between supply-side and demand-side operational decisions. *Available at SSRN 4287750*.

Müller, S., Huber, J., Fleischmann, M., and Stuckenschmidt, H. (2020). Data-driven inventory management under customer substitution. *Available at SSRN 3624026*.

Muriana, C. (2016). An EOQ model for perishable products with fixed shelf life under stochastic demand conditions. *European Journal of Operational Research*, 255(2):388–396.

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

Nahmias, S. (1975). Optimal ordering policies for perishable inventory—ii. *Operations Research*, 23(4):735–749.

Nahmias, S. (1976). Myopic approximations for the perishable inventory problem. *Management Science*, 22(9):1002–1008.

Nahmias, S. (1977). Higher-order approximations for the perishable-inventory problem. *Operations Research*, 25(4):630–640.

Nahmias, S. (2011). *Perishable inventory systems*, volume 160. Springer Science & Business Media.

Nandakumar, P. and Morton, T. E. (1993). Near myopic heuristics for the fixed-life perishability problem. *Management Science*, 39(12):1490–1498.

Office of the Auditor General of Ontario (2020). Value-for-money audit blood management and safety. `https://www.auditor.on.ca/en/content/annualreports/arreports/en20/20VFM_02bloodmgmt.pdf`. Accessed Date: 2020-10-21.

Olsson, F. (2010). An inventory model with unidirectional lateral transshipments. *European Journal of Operational Research*, 200(3):725–732.

Or, I. and Pierskalla, W. P. (1979). A transportation location-allocation model for regional blood banking. *AIIE Transactions*, 11(2):86–95.

Organization, W. H. et al. (2010). Who experts' consultation on estimation of blood requirements: 03-05 february 2010, who-hq, geneva: meeting report. Technical report, World Health Organization.

Osorio, A. F., Brailsford, S. C., and Smith, H. K. (2018). Whole blood or apheresis donations? a multi-objective stochastic optimization approach. *European Journal of Operational Research*, 266(1):193–204.

Parlar, M., Perry, D., and Stadje, W. (2011). FIFO versus LIFO issuing policies for stochastic perishable inventory systems. *Methodology and Computing in Applied Probability*, 13:405–417.

Pegels, C. C. and Jelmert, A. E. (1970). An evaluation of blood-inventory policies: A markov chain application. *Operations Research*, 18(6):1087–1098.

Perera, S. C. and Sethi, S. P. (2023a). A survey of stochastic inventory models with fixed costs: Optimality of (s, S) and (s, S)-type policies—continuous-time case. *Production and Operations Management*, 32(1):154–169.

Perera, S. C. and Sethi, S. P. (2023b). A survey of stochastic inventory models with fixed costs: Optimality of (s, S) and (s, S)-type policies—discrete-time case. *Production and Operations Management*, 32(1):131–153.

Perry, D. and Posner, M. J. (1998). An (S- 1, S) inventory system with fixed shelf life and constant lead times. *Operations Research*, 46(3-supplement-3):S65–S71.

Pierskalla, W. P. (2004). Supply chain management of blood banks. *Operations research and health care: A handbook of methods and applications*, pages 103–145.

Pirabán, A., Guerrero, W. J., and Labadie, N. (2019). Survey on blood supply chain management: Models and methods. *Computers & Operations Research*, 112:104756.

Prak, D., Teunter, R., and Syntetos, A. (2017). On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research*, 256(2):454–461.

Prastacos, G. P. (1984). Blood inventory management: an overview of theory and practice. *Management Science*, 30(7):777–800.

Priyadarshi, R., Panigrahi, A., Routroy, S., and Garg, G. K. (2019). Demand forecasting at retail stage for selected vegetables: a performance analysis. *Journal of Modelling in Management*, 14(4):1042–1063.

Puchalsky, W., Ribeiro, G. T., da Veiga, C. P., Freire, R. Z., and dos Santos Coelho, L.

(2018). Agribusiness time series forecasting using wavelet neural networks and metaheuristic optimization: An analysis of the soybean sack price and perishable products demand. *International Journal of Production Economics*, 203:174–189.

Rajendran, S. and Ravindran, A. R. (2017). Platelet ordering policies at hospitals using stochastic integer programming model and heuristic approaches to reduce wastage. *Computers & Industrial Engineering*, 110:151–164.

Rajendran, S. and Ravindran, A. R. (2019). Inventory management of platelets along blood supply chain to minimize wastage and shortage. *Computers & Industrial Engineering*, 130:714–730.

Rajendran, S. and Srinivas, S. (2020). Hybrid ordering policies for platelet inventory management under demand uncertainty. *IISE Transactions on Healthcare Systems Engineering*, 10(2):113–126.

Ramezanian, R. and Behboodi, Z. (2017). Blood supply chain network design under uncertainties in supply and demand considering social aspects. *Transportation Research Part E: Logistics and Transportation Review*, 104:69–82.

Ranstam, J. (2012). Why the P-value culture is bad and confidence intervals a better alternative. *Osteoarthritis and Cartilage*, 20(8):805–808.

Ravichandran, N. (1995). Stochastic analysis of a continuous review perishable inventory system with positive lead time and Poisson demand. *European Journal of Operational Research*, 84(2):444–457.

Rutherford, C., Cheng, S. Y., and Bailie, K. (2016). Evidence of bullwhip in the blood supply chain. In *23rd EurOMA Conference 2016*.

Samani, M. R. G., Hosseini-Motlagh, S.-M., and Ghannadpour, S. F. (2019). A multilateral perspective towards blood network design in an uncertain environment: Methodology and implementation. *Computers & Industrial Engineering*, 130:450–471.

Samuel, C., Gonapa, K., Chaudhary, P., and Mishra, A. (2010). Supply chain dynamics in healthcare services. *International Journal of Health Care Quality Assurance*, 23(7):631–642.

Saoud, P., Kourentzes, N., and Boylan, J. E. (2022). Approximations for the lead time variance: A forecasting and inventory evaluation. *Omega*, 110:102614.

Sarhangian, V., Abouee-Mehrizi, H., Baron, O., and Berman, O. (2018). Threshold-based allocation policies for inventory management of red blood cells. *Manufacturing & Service Operations Management*, 20(2):347–362.

Shen, L., Li, F., Li, C., Wang, Y., Qian, X., Feng, T., and Wang, C. (2020). Inventory optimization of fresh agricultural products supply chain based on agricultural superdocking. *Journal of Advanced Transportation*, 2020:1–13.

Silva Filho, O. S., Carvalho, M. A., Cezarino, W., Silva, R., and Salviano, G. (2013). Demand forecasting for blood components distribution of a blood supply chain. *IFAC Proceedings Volumes*, 46(24):565–571.

Silva Filho, O. S., Cezarino, W., and Salviano, G. R. (2012). A decision-making tool for demand forecasting of blood components. *IFAC Proceedings Volumes*, 45(6):1499–1504.

Silver, E. A., Pyke, D. F., and Thomas, D. J. (2016). *Inventory and production management in supply chains*. CRC Press.

Simchi-Levi, D., Chen, X., Bramel, J., et al. (2005). The logic of logistics. *Theory, algorithms, and applications for logistics and supply chain management*.

Sirelson, V. and Brodheim, E. (1991). A computer planning model for blood platelet production and distribution. *Computer Methods and Programs in Biomedicine*, 35(4):279–291.

Stanworth, S. J., New, H. V., Apelseth, T. O., Brunskill, S., Cardigan, R., Doree, C., Germain, M., Goldman, M., Massey, E., Prati, D., et al. (2020). Effects of the COVID-19 pandemic on supply and use of blood for transfusion. *The Lancet Haematology*, 7(10):E756–E764.

Stokes, E. A., Wordsworth, S., Staves, J., Mundy, N., Skelly, J., Radford, K., and Stanworth, S. J. (2018). Accurate costs of blood transfusion: a microcosting of administering blood products in the united kingdom national health service. *Transfusion*, 58(4):846–853.

Tashman, L. J. (2000). Out-of-sample tests of forecasting accuracy: an analysis and review. *International Journal of Forecasting*, 16(4):437–450.

Taylor, S. J. and Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1):37–45.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.

Tiwari, S., Jaggi, C. K., Gupta, M., and Cárdenas-Barrón, L. E. (2018). Optimal pricing and lot-sizing policy for supply chain system with deteriorating items under limited storage capacity. *International Journal of Production Economics*, 200:278–290.

Toomey, J. W. (2000). *Inventory management: principles, concepts and techniques*, volume 12. Springer Science & Business Media.

Topkis, D. M. (1998). *Supermodularity and complementarity*. Princeton university press.

Trapero, J. R., Cardos, M., and Kourentzes, N. (2019). Empirical safety stock estimation based on kernel and GARCH models. *Omega*, 84:199–211.

Twumasi, C. and Twumasi, J. (2022). Machine learning algorithms for forecasting and backcasting blood demand data with missing values and outliers: A study of Tema General Hospital of Ghana. *International Journal of Forecasting*, 38(3):1258–1277.

Van Dijk, N., Haijema, R., Van Der Wal, J., and Sibinga, C. S. (2009). Blood platelet production: a novel approach for practical optimization. *Transfusion*, 49(3):411–420.

Van Donselaar, K. H., Peters, J., de Jong, A., and Broekmeulen, R. A. (2016). Analysis and forecasting of demand during promotions for perishable items. *International Journal of Production Economics*, 172:65–75.

van Sambeeck, J. H., van Brummelen, S., van Dijk, N., and Janssen, M. (2022). Optimal blood issuing by comprehensive matching. *European Journal of Operational Research*, 296(1):240–253.

Volken, T., Buser, A., Castelli, D., Fontana, S., Frey, B. M., Rüsges-Wolter, I., Sarraj, A., Sigle, J., Thierbach, J., Weingand, T., et al. (2018). Red blood cell use in Switzerland: trends and demographic challenges. *Blood Transfusion*, 16(1):73–82.

Weiss, H. J. (1980). Optimal ordering policies for continuous review perishable inventory models. *Operations Research*, 28(2):365–374.

Weraikat, D., Zanjani, M. K., and Lehoux, N. (2016). Two-echelon pharmaceutical reverse supply chain coordination with customers incentives. *International Journal of Production Economics*, 176:41–52.

Williams, C. L. and Patuwo, B. E. (1999). A perishable inventory model with positive order lead times. *European Journal of Operational Research*, 116(2):352–373.

Xin, L. and Goldberg, D. A. (2016). Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. *Operations Research*, 64(6):1556–1565.

Xu, G., Feng, J., Chen, F., Wang, H., and Wang, Z. (2019). Simulation-based optimization of control policy on multi-echelon inventory system for fresh agricultural products. *International Journal of Agricultural and Biological Engineering*, 12(2):184–194.

Yang, C.-L. and Sutrisno, H. (2018). Short-term sales forecast of perishable goods for franchise business. In *2018 10th International Conference on Knowledge and Smart Technology (KST)*, pages 101–105. IEEE.

Zahiri, B. and Pishvaee, M. S. (2017). Blood supply chain network design considering blood group compatibility under uncertainty. *International Journal of Production Research*, 55(7):2013–2033.

Zahiri, B., Torabi, S., Mousazadeh, M., and Mansouri, S. (2015). Blood collection management: Methodology and application. *Applied Mathematical Modelling*, 39(23-24):7680–7696.

Zhang, C., Ayer, T., White, C. C., Bodeker, J. N., and Roback, J. D. (2022). Inventory sharing for perishable products: Application to platelet inventory management in hospital blood banks. *Operations Research*.

Zhang, H., Zhang, J., and Zhang, R. Q. (2020). Simple policies with provable bounds for managing perishable inventory. *Production and Operations Management*, 29(11):2637–2650.

Zheng, M., Wu, K., and Shu, Y. (2016). Newsvendor problems with demand forecast updating and supply constraints. *Computers & Operations Research*, 67:193–206.

Zhou, D., Leung, L. C., and Pierskalla, W. P. (2011). Inventory management of platelets in hospitals: Optimal inventory policy for perishable products with regular and optional expedited replenishments. *Manufacturing & Service Operations Management*, 13(4):420–438.

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