Data Driven Personalized Management of

Hospital Inventory of Perishable and Substitutable Blood Units

by

Gina Dumkrieger

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved July 2020 by the
Graduate Supervisory Committee:

Pitu Mirchandani, Chair
John Fowler
Teresa Wu
Feng Ju

ARIZONA STATE UNIVERSITY

August 2020

ABSTRACT

The use of Red Blood Cells (RBCs) is a pillar of modern health care. Annually, the lives of hundreds of thousands of patients are saved through ready access to safe, fresh, blood-type compatible RBCs. Worldwide, hospitals have the common goal to better utilize available blood units by maximizing patients served and reducing blood wastage. Managing blood is challenging because blood is perishable, its supply is stochastic and its demand pattern is highly uncertain. Additionally, RBCs are typed and patient compatibility is required.

This research focuses on improving blood inventory management at the hospital level. It explores the importance of hospital characteristics, such as demand rate and blood-type distribution in supply and demand, for improving RBC inventory management. Available inventory models make simplifying assumptions; they tend to be general and do not utilize available data that could improve blood delivery. This dissertation develops useful and realistic models that incorporate data characterizing the hospital inventory position, distribution of blood types of donors and the population being served.

The dissertation contributions can be grouped into three areas. First, simulations are used to characterize the benefits of demand forecasting. In addition to forecast accuracy, it shows that characteristics such as forecast horizon, the age of replenishment units, and the percentage of demand that is forecastable influence the benefits resulting from demand variability reduction.

Second, it develops Markov decision models for improved allocation policies under emergency conditions, where only the units on the shelf are available for dispensing. In this situation the RBC perishability has no impact due to the short timeline for decision making. Improved location-specific policies are demonstrated via simulation models for two emergency event types: mass casualty events and pandemic influenza.

Third, improved allocation policies under normal conditions are found using Markov decision models that incorporate temporal dynamics. In this case, hospitals receive replenishment and units age and outdate. The models are solved using Approximate Dynamic Programming with model-free approximate policy iteration, using machine learning algorithms to approximate value or policy functions. These are the first stock- and age-dependent allocation policies that engage substitution between blood type groups to improve inventory performance.

DEDICATION

For my parents.

ACKNOWLEDGEMENTS

There are innumerable people that deserve recognition for their support over the years. I want to thank my committee for their patience and support during this arduous expedition.

Dr. Fowler's presentation on humanitarian logistics brought me into the program. Thank you for getting me started and connecting me with the Mayo Rochester group where this project originated. Thank you to Dr. Mirchandani for taking me on as a student and sticking with me.

Thank you to Dr. Wu for connecting me with the Mayo headache crew and for her sound advice and generally being a kind and encouraging presence.

Dr. Namir Saman originally made me think this was a good idea. Thank you?

I owe much gratitude to Edgar Hassler and Brian Stone without whom I wouldn't have made past the first year. I also owe Edgar for his service as my technical consultant, copyeditor, cloud computing provider and friend.

I would like to acknowledge Dr. Connie Borror who in addition to being technically brilliant was tremendously kind and encouraging to so many students including myself. I will attempt to emulate her spirit of generosity throughout my life.

Thank you to Dr. Rong Pan and Ms. Christina Sebring.

Thank you, Dr. Linda Chattin. I am forever grateful.

Thank you to everyone in the Neuroimaging of Headache Disorders lab at Mayo Clinic for their patience and support and for teaching me how to do research in an organized fashion.

Endless gratitude to Dr. Dimitri Bertsekas for his guidance and feedback especially on the problems of Chapter 5. Thank you for taking the time.

I am grateful to my esteemed lab mates and fellow students for the community and enlightening conversation. Special thanks to Kerem for the encouragement and to Gita for the working walks to Starbucks.

Any list of acknowledgements would be incomplete without some acknowledgement of the personal support I have received. It has been long process. Thank you to my friends who tolerated my endless complaints and bought me lunch more times than I can count. Thank you to Michael Moore for keeping me fed and coffee'd and provided with the appropriate amount of cats. Thank you for never giving me access to your Netflix account.

Thank you to my family. I love you very much.

TABLE OF CONTENTS

Page

LIST OF TABLES xi

LIST OF FIGURES xiv

CHAPTER

1. INTRODUCTION 1

1.1.Motivation for the Research 1

1.2.Blood Supply Network 1

1.2.1.Blood Supply Management 2

1.2.2.Blood Inventory Management 3

1.3.Preview of Dissertation Scope and Contributions 4

2. REVIEW OF RELEVANT MODELING APPROACHES 8

2.1Modeling Blood Inventory Management 8

2.1.1 Clinical Research Supporting Blood Availability Assumptions. 8

2.1.2 Interventions That Directly Impact Supply and Demand Variability. ..9

2.1.3 Inventory Controls that Respond to Uncertainty 11

2.2. Methods for Capturing System Complexity..12

2.2.1 Simulation to Model an Existing Process and See the Effect of New Policies..13

2.2.2.Markov Decision Processes (MDP) and Approximate Dynamic Programming (ADP) 15

2.3. Simulation to Show Effect of Proposed Shelf Life Changes................17

CHAPTER Page

3. USE OF SIMULATION AND FORECASTING MODELS TO EVALUATE THE BENEFIT OF FORECAST IMPROVEMENT ON BLOOD INVENTORY MANAGEMENT 21

3.1. Introduction 21

3.2. Literature on Forecasting Perishable Demand 22

3.2.1. Product Perishability. 22

3.2.2. Demand Uncertainty and Forecasting Error. 23

3.2.3. Prediction of Blood Demand. 23

3.3. Forecast and Simulation Models 25

3.3.1. Forecasting Error and Reorder Quantities. 25

3.3.2. Hospital Simulation Model. 27

3.4. Simulation Experiments 28

3.4.1.Preliminary Experiments: Non-forecast Related Factors Affecting Forecast Impact 28

3.4.2. Effect of Error Variability Under 100% Predictable Demand. 30

3.4.3. High Demand Case, Designed Experiment. 37

3.5. Discussion of Simulation Findings 39

4. MDP MODEL TO FIND OPTIMAL ALLOCATION UNDER EMERGENCY41

4.1 Introduction and Problem Definition: 41

4.2. Literature Review on Blood Allocation and Emergency Response 43

4.2.1. Emergency Context 43

4.2.2. Allocation Problems. 44

CHAPTER Page

4.2.3. Other Related Research....46

4.3. Exact MDP Model Formulation....46

4.4. Computational Results and Analysis....49

4.4.1. Results: Optimal Policy for U.S. Blood Type Distribution....50

4.4.2. Results: Effect of Population Distribution....55

4.4.3. Results: Effect of Initial Inventory.....59

4.4.4. Results: Effects of Timing of the Application of Optimal Policy.....61

4.4.5. Results: Rh-type Specific Policies.....63

4.5. Detailed Simulation Case Studies....66

4.5.1. Scenario 1: Mass Casualty Incident (MCI)....68

4.5.1.1. MCI Case Results....70

4.5.1.2. MCI Case Discussion....71

4.5.2. Scenario 2: Pandemic Influenza....71

4.5.2.1. Pandemic Case Results.....73

4.5.2.2. Pandemic Case Discussion.....76

4.5.3. Discussion of the Simulation Case Studies....77

4.6. Conclusions and Potential Future Research....78

4.6.1. Limitations and Potential Future Work....80

5. APPROXIMATE MDP MODEL TO FIND OPTIMAL ALLOCATION UNDER NORMAL OPERATIONS.... 82

5.1 Introduction and Problem Definition....82

5.2. Blood Inventory Management with Perishability and Substitutability Background....83

5.3. Approximate Dynamic Programming Background....88

5.3.1. Rollout Background....90

5.3.2. Approximate Policy Iteration with Q-factor Approximation Background....92

5.3.3. Policy Iteration Based on Policy Improvement by Rollout in Policy Space Approximation Background....95

5.4. Approximate Dynamic Programming Model Formulation & Solution 95

5.4.1. Changes to Model Formulation to Incorporate Temporal Dynamics.96

5.4.2. The Rollout model....100

5.4.3. Approximate Policy Iteration with Q-factor Approximation Model.102

5.4.4. Policy Iteration Based on Policy Improvement by Rollout and Policy Approximation in Policy Space Model....105

5.4.5. Procedures for Finding Good Candidate Policies....109

5.5. Computational Results and Analysis of Approximate MDPs....114

5.5.1. Comparison of Alternate Fixed Priority Policies....116

5.5.2. Rollout Results....121

5.5.1. Results of Policy Iteration based on Policy Improvement by Rollout and Policy Approximation in Policy Space....123

5.5.1.1. Cobre API Results....123

5.5.1.2. Mayo API Results....130

CHAPTER Page

5.5.1.3. Sensitivity Analysis..........134

5.5.1.4. Application of Mayo Policy to Cobre Scenario..........142

5.5.1.5. Feature Importance for Best API Model..........143

5.6. Conclusions..........144

6. SUMMARY AND CONCLUSIONS..........147

REFERENCES..........153

APPENDIX

A MIXTURE DISTRIBUTIONS FOR AGE AT RECEIPT..........167

B COBRE ALTERNATIVE FIXED PRIORITY POLICY COMPARISON 174

C DIRECT API SOLUTION SETUP CONFIGURATIONS: COBRE..........176

D DIRECT API SOLUTION SETUP CONFIGURATIONS: MAYO..........179

E API SIMULATION SETTINGS..........181

F PIPI_directMayo0.221 FINAL XGBCLASSIFIER PARAMETERS AND VARIABLE IMPORTANCE..........183

LIST OF TABLES

Table Page

1. Results of Previous Research on Effect off Change in Maximum Shelf Life......19

2. Exploratory Simulation Results Showing Effect of Reorder Schedule and Unit Age at Delivery on Annual Average Unmet Demand and Outdating..........................30

3. Results of Simulations for Varying error-SD and Safety Factor (%). 32

4. Coefficient Estimates and Variable Significance from Analysis of Variance on Results of Experiment on Effect of Percent of Predictable Demand, Forecast Error Mean and SD, and Safety Percentages on High Demand Case...........................38

5. US Population ABO Type Distribution..50

6. First Unserved patient: Standard, U.S. Optimal, Hungary Optimal, India Optimal and Chicago Optimal Policies Applied to U.S., Hungary, India and Chicago Average Populations 58

7. Additional patients served through earlier application of optimal policy 62

8. US and Chicago type distribution, all and negative 64

9. Results of Mass Casualty Incident DES simulations 70

10. Results of Pandemic Flu DES simulations. 74

11. Direct PI models selected for Head to Head Comparison in Further Simulations: Cobre 112

12. Direct PI Models Selected for Head To Head Comparison in Further Simulation Studies: Mayo...114

13. Simulation Results: Cobre Alternate Fixed Priority Simulation Comparison...119

Table Page

14. Simulation Results: Mayo Alternate Fixed Priority Policy Comparison 120

15. Simulation Results: Cobre Rollout Comparison 122

16. API Simulation Results: Cobre Standard Policy Comparison..........125

17. API Simulation Results: Additional Basic Comparisons Runs for Cobre 126

18. Units Allocated by type, by policy: All standard Cobre policies..........128

19. To:From Matrix: Cobre Default..........129

20. To:From Matrix: Cobre PIPI_direct_2extension50 129

21. API Simulation Results: Mayo Standard Policy Comparison 132

22. Units Allocated by Type, by Policy: Mayo 133

23. To:From Matrix: Mayo Default..........133

24. To:From Matrix: Mayo PIPI_direct_mayo0.221 134

25. API Simulation Results: Cobre Increased Demand 136

26. API Simulation Results: Cobre Decreased Demand..........136

27. API Simulation Results: Cobre Older Age..........137

28. API Simulation Results: Cobre Young Age..........137

29. API Simulation Results: Mayo Increased Demand..........140

30. API Simulation Results: Mayo Decreased Demand 140

31. API Simulation Results: Mayo Older Age 141

32. API Simulation Results: Cobre Alternate Fixed Priority Policy Simulation Comparisons..........175

33. Direct API Solution Setup Configurations: Cobre..........177

34. Direct API Solution Setup Parameter Configurations: Mayo..........180

Table		Page
35.	Settings for API Policy Comparison Simulations.	182
36.	PIPI_directMayo0.221 Final XGBClassifier Parameters	184

LIST OF FIGURES

Figure Page

1. Blood Supply Network Illustration 1

2. Dissertation Research Scope and Problems Addressed 5

3. Plot of outcomes vs error SD for high demand case. Outdated units plot shows quadratic fit line. 33

4. Plot of Outcomes Vs Error SD for Low Demand Case. Outdated units plot shows quadratic fit line. 33

5. Daily and Weekly Demand Distributions for High Demand Case 36

6. Daily and Weekly Demand Distributions for Low Demand Case. 36

7. ABO Allowable Blood Type Substitution 49

8. Recommendations for AB Patients When There is No AB Inventory (Projection). 51

9. Serving AB Patient, Give A/Give B Boundary 52

10. First Unserved Patient: US policy and population. Patient number of the first unserved patient, optimal vs default AB policy, based on U.S. population average for initial inventory [20A, 20B,.0AB,20O] 53

11. Third Unserved Patient. Patient number of the third unserved patient: optimal vs standard AB policy, based on U.S. population average for initial inventory [20 A, 20 B, 0 AB, 20 O] 54

12. Count of Unserved Patients in First 60 Patients; Comparisons with Theoretical Minimums 55

Figure Page

13. Blood Type Distributions: U.S., Hungary, India, Chicago 56

14. Pairwise Comparison of Location-specific Polices with Simulated Patient Populations... 58

15. Effect of Inventory Level on Policy Improvement. 60

16. Effect of Inventory Level on Policy Improvement 60

17. Timing of Optimal Policy. Additional patients served through earlier application of optimal policy, compared to allocating *using* standard policy until the inventory of any type reaches 10 or fewer*, then* allocating according to optimal policy 62

18. Approximate Separating Planes Between Inventory Positions where A is Recommended (below the planes) and Areas Where B is Allocated (above the planes), for AB Patients When There is No AB Inventory On Hand. From top to bottom are the planes derived from Chicago Rh positive optimal policy, Chicago Rh negative optimal policy, US Rh positive optimal policy, US Rh negative optimal policy.. 64

19. Policy Recommendation Boundary Planes Intersection with Zero O Inventory Plane.. 65

20. Extended Projections of planes separating “Give A” and “Give B” areas for U.S., India and Hungary averages. A is given on the side of the plane where quantity of A >0... 65

21. Unmet Demand Over The Course of a Pandemic. Shows the average number of patients with unserved demand over the duration of the pandemic and in the days following. Day 0 is the last day of the pandemic (last day of reduced supply). ... 75

Figure Page

22. Unmet Demand and Outdated Units: (a) Unmet demand and outdated units for all scenarios; (b) closeup of unmet demand and outdated units, showing mass casualty scenarios only. 76

23. Diagram of Rollout Process 90

24. Illustration of the MDP Model for the Improved Allocation of Units Under Normal Conditions Problem 97

25. Allowable Blood Type Substitutions in 8-type Model 100

26. Pseudocode for Single Period Rollout Implementation 101

27. Pseudocode for Testing Rollout Implementation 102

28. API with Q-factor Approximation Implementation. 104

29. Collection of Samples in Two API models; Q-factor Approximation and Approximation in Policy Space 107

30. Pseudocode for API with Policy Estimation Implementation 108

31. Value by Iteration: DirectPI_Cobre2 110

32. Value by Iteration: DirectPI_Cobre2 Extended Iterations 111

33. Value by Iteration API Mayo 113

34. Fixed Priority Policies 117

35. Cobre Standard Comparison Reward Boxplots 127

36. Cobre PIPI_direct_2extension50 Standard Extra Runs Boxplot 127

37. Cobre Improvement of Rollout Over API 128

Figure Page

38. Mayo API Standard Comparison Boxplots. Figures Show the Improvement of Each New Policy Over the Default Policy in Terms of Undiscounted Reward and Discounted Reward. 131

39. Cobre Policy Improvement over Mayo Policies in Cobre Scenario. 143

40. Original Data vs Model Fits for A Positive Age at Receipt Mixture Distribution 171

41. Fit Characteristic Curves for Each Blood Type 173

42. Variable Importance of PIPI_directMayo0.221 Final XGBClassifier Model. 185

CHAPTER 1

INTRODUCTION

1.1.Motivation for the Research

Blood is a critical weapon the arsenal of modern medicine. Each day thousands of patients around the world are saved by the transfusion of blood products. It is a vital product, but its availability is never assured and when the supply chain fails patients can die. It is estimated that 150,000 deaths from maternal hemorrhage alone could be prevented by access to safe blood (Schantz-Dunn & Nour, 2011). While in some areas blood is in constant shortage, in other areas units expire on the shelves. Five million patients in the US receive blood transfusions every year but its collection and distribution pose significant challenges. The supply of blood products is fraught with challenges both logistical and personal, both global and local. It is this intersection of human and technical problems that motivates this dissertation.

1.2. Blood Supply Network

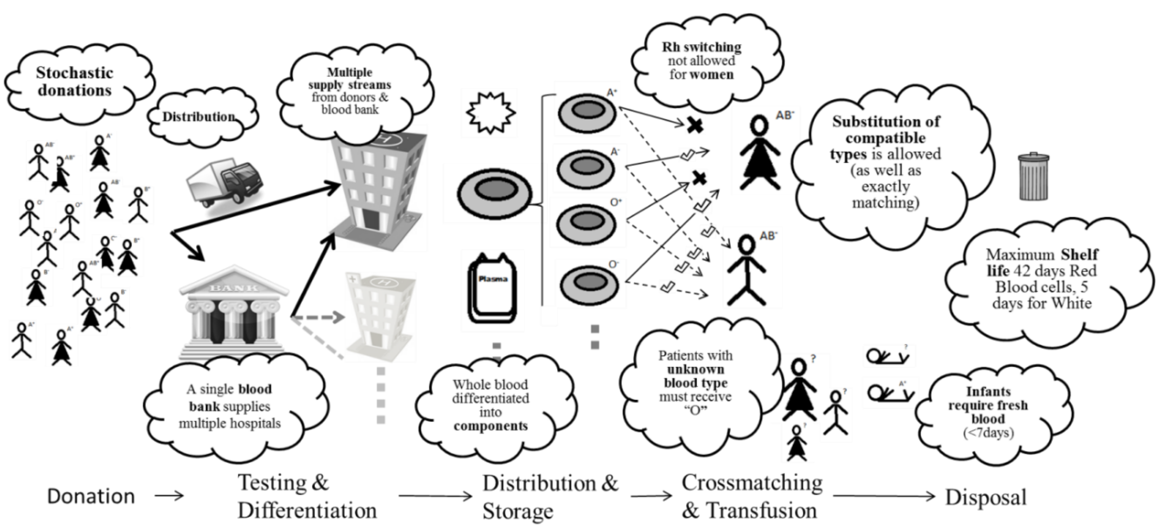

Figure 1. Blood Supply Network Illustration

Figure 1describes the blood supply chain and highlights some of the challenges found in it. The supply chain described here reflects the most common processes in the United States. Countries with well-established donation systems will have similar processes but outside of these countries it can look very different.

1.2.1. Blood Supply Management

Every unit of blood product transfused must be collected from a donor, typed and screened for disease. Though some hospitals possess internal blood donor centers, typically these collection and processing functions are performed by centralized blood providing organizations like the Red Cross or America's Blood Services. For safety reasons, the World Health Organization (WHO) recommends collecting blood from 100% voluntary, unpaid donors. In the United States virtually all blood collected meets this description. The reliance on volunteer donors turning out to blood drives, bloodmobiles and local blood centers is a major supply-side source of uncertainty. Blood providers attempt to motivate blood donors of specific types through telephone calls or non-monetary renumeration (e.g. t-shirts) and must make decisions about where and when to schedule blood drives and bloodmobile appearances.

From a donation site, blood must be taken to a processing facility to be typed and screened for disease. For units of whole blood decisions must be made about whether to differentiate it into components (red cells, plasma, platelets etc.) or leave it as whole blood. Each of these blood components has different uses and different shelf lives.

The regional blood center, which may or may not be the same as the processing center, is responsible for distributing the blood to hospitals based on their orders. The regional blood center responds to orders from the hospitals but chooses which units (e.g. of which age) to use to fulfill the orders. A regional blood center will work with hospitals across its region to provide blood for transfusion in that region. They may also rotate blood from smaller hospitals to larger, in order to increase the chances of it being used before it expires. The national demand for blood is met through such a network of regional systems.

1.2.2. Blood Inventory Management

At the hospital, exactly matched or compatible blood is dispensed to patients that need it. Blood demand is largely stochastic. Some units will be used in the treatment of chronic diseases, some will be used in elective procedures and some will be used in emergent procedures. Decisions on what blood type to give a patient are made heuristically or ad hoc in the hospital setting based on availability of different blood types and other factors. Often hospitals performing elective procedures will perform crossmatch and reserve units for elective procedures. The units reserved become temporarily unavailable; however, if not used for that purpose they will be returned to inventory. Blood that reaches the maximum shelf life before use must be discarded. Currently the maximum shelf life for red cells is 42 days, with the oldest units being assigned first among compatible blood types. An exception to this rule is that neonates are given blood 7 days or younger whenever possible and some facilities will attempt to provide specific patient classes (e.g. cardiac patients) with fresh blood.

The importance of blood in healthcare is indisputable. The many functional tasks required in the blood supply chain; collection, testing, distribution, rotation, replenishment and allocation, combined with the features of substitutability, perishability and multiple sources of uncertainty result in an incredibly complex supply chain. Together they form fertile ground for use-inspired and technically challenging research.

1.3. Preview of dissertation scope and contributions

There are too many tasks and challenges inherent in the blood supply chain to tackle them all simultaneously. This dissertation addresses challenges surrounding the provision of red blood cells at the local (hospital) level that have not been adequately address in previous research. Throughout the focus is on more accurately approximating the complex reality of the blood supply chain, local perspectives, and the importance of context in inventory management. Having discussed the reality of the blood supply chain in Chapter 1, Chapter 2 gives an overview of existing modeling approaches to blood related logistics problems. Blood inventory management is by no means a new problem in operations research and Chapter 2 discusses the main approaches that have historically been used for the most studied blood related problems. It also highlights the need for multiple modeling approaches in the complex problems addressed in this dissertation.

Figure 2 is a thematic map of the research addressed by this dissertation and how the chapters of the dissertation relate to each other. The use of data-backed simulation and the focus on hospital level issues are common elements found in each chapter. Not captured in the map is the lack of prior work on the selected research topics.

Chapter 3 discusses the potential benefits of demand forecasting improvement by hospitals and how facility characteristics and inventory management decisions interact with forecasting to increase or decrease the benefits of accurate forecasting. It is often assumed that improved forecasting is always beneficial. However, through the use of detailed discrete event simulation models, this chapter shows that, in the complicated milieu of the real world, the benefit is highly dependent on the characteristics of the hospital environment.

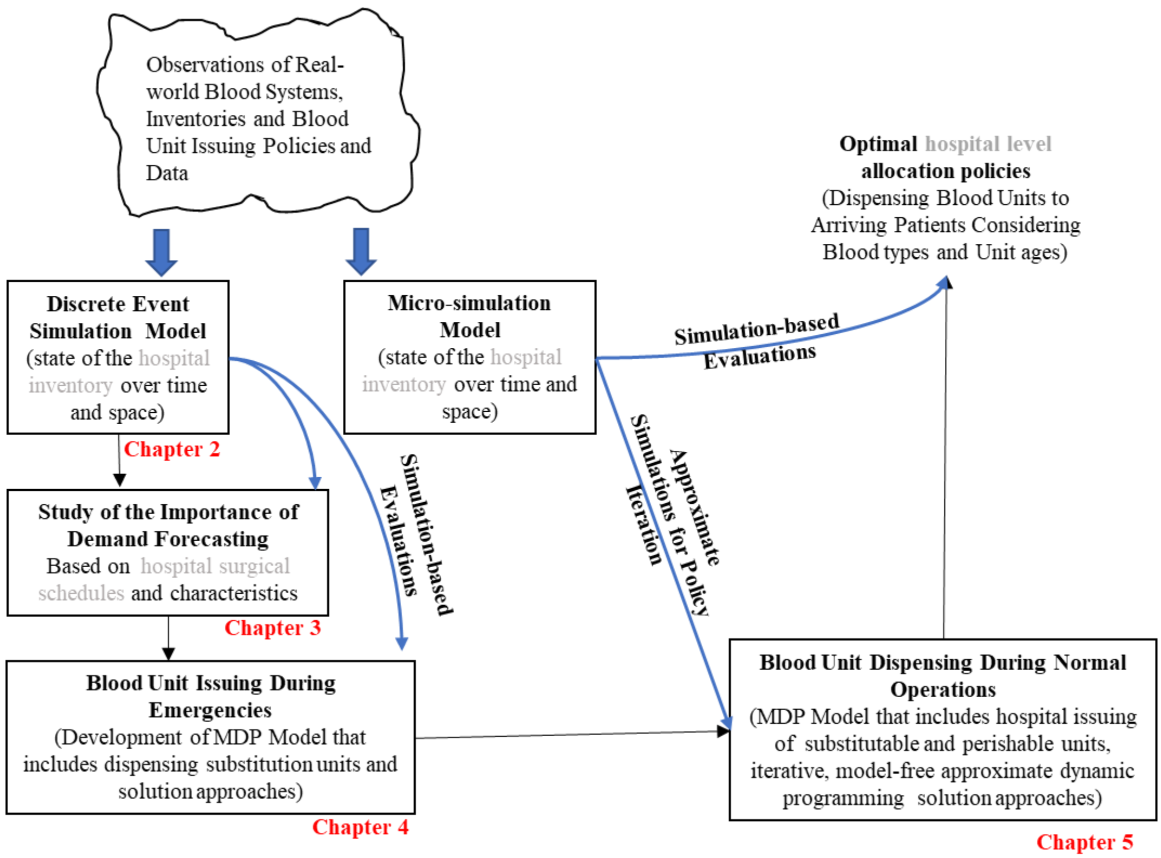

Figure 2. Dissertation Research Scope and Problems Addressed

Chapter 4 is a major contribution of this dissertation. It addresses the problem of unit allocation under emergency scenarios. New stock level-dependent Markov Decision Process (MDP)-derived allocation policies that consider the substitutability of blood units are developed and tested extensively. The policies are tailored to the blood type distribution of the patient population and the necessity of such local policies is shown. Monte Carlo simulations show the MDP-derived policies improve the utilization of existing inventories in the absence of perishability. They confirm the importance of location-specific polices. Further examination via discrete event simulations (DES) evaluates the benefits of the developed allocation policies over the default policy in a more complicated environment, during simulated influenza pandemics or mass casualty events, and over a longer time frame. These results confirm the Monte Carlo simulation results and indicate also the benefits of the MDP-based policies to daily (non-emergency) allocation.

Chapter 5 develops another major contribution in this dissertation. It addresses the development of an age-sensitive, stock level-dependent, proactively substituting, allocation policy for daily use. As opposed to Chapter 4 where the MDP model featured a short time frame needed for emergencies and which excluded perishability and replenishment, the MDP-based policies developed in Chapter 5 include both perishability and replenishment. Large problem sizes necessitate an approximation; the development and solutions of the approximate MDP-based policies are discussed.

This dissertation focuses on bringing an extended degree of realism to the problem of allocating red blood units. It focuses on a hospital perspective, and considers the traits of the hospital and the impact of those traits on its blood inventory management functions. Briefly, the dissertation contributes to the study of hospital level blood inventory management by increasing the understanding of how the benefits of demand forecasting based on surgical schedules are influenced by hospital characteristics and policies. Additionally, it makes two significant contributions on policies of blood unit allocation, a highly understudied area. It develops stock level-dependent **emergency** allocation policies that optimize substitution between blood types to increase the utilization of a limited stock of blood. It also develops stock age- and level-dependent red blood cell unit allocation policies for **normal** operating conditions. In addition to the novelty of the policies directing blood issuing versus the commonly studied blood replenishment, these policies are among the first to use stock age and stock levels to make substitution decisions, not as a recourse action, but as an active choice to reduce outdating and maximize service levels.

CHAPTER 2

REVIEW OF RELEVANT MODELING APPROACHES

2.1.Modeling Blood Inventory Management

The blood supply chain was described in the last chapter. This chapter looks at the same system through the lens of potential model- and data-based inventory management. It describes how researchers are addressing various points of intervention, that is, what problems are being tackled and what tools are being used or considered. It then focuses on the specific problems and tools being used or designed in this dissertation. Discussion of the research topics is organized by whether the study/tools (1) increase knowledge about the existing uncertainties, (2) create interventions to directly influence sources of uncertainty, or (3) enable interventions to respond to effects of uncertainties.

One of the goals of this dissertation is to maintain a high degree of fidelity between the modeling and reality, by avoiding unnecessary or unsupported simplifications and assumptions. In support of that goal, two major modeling approaches are used: simulation and Markov Decisions Processes (MDPs). These approaches accommodate the desired level of detail, while being adaptable to large problem sizes. In this chapter, there is a brief discussion of each of these approaches and how it is used in the reported literature. Finally, there is a recap of the detailed discrete event simulation model (DES) used in Chapters 3 and 4 of this dissertation.

2.1.1. Clinical research supporting blood availability assumptions.

Medically focused blood research can impact blood inventory management directly or it can improve quality of inputs to higher level modeling efforts. As an example of direct effect, disease prevalence research guides regulations on which donors are permitted donate blood and which must be deferred. Epidemiologists survey disease prevalence in sub-populations (Tounkara et al., 2009; Tserenpuntsag et al., 2010) and monitor for emerging diseases like Zika virus (Kuehnert, 2016). The work results in donor deferral regulations and guidance. This directly changes the size of the available donor pool. Disease prevalence information also serves as an input to models to develop disease screening procedures (see Xie, Bish, Bish, Slonim, & Stramer, 2012). In another example of direct effect, the RECESS study was a clinical trial in response to the growing belief (Flegel, Natanson, & Klein, 2014) that fresher blood results in better patient health outcomes, especially for cardiac patients. Results of the clinical trial (Steiner et al., 2015) did not confirm that hypothesis. If the trial outcome had been different or if decisions had been made without rigorous evidence, it could have resulted in regulation changes that would have made it more difficult to meet blood demand. This example also illustrates the importance of verifying assumptions before making changes to inventory control policies.

2.1.2. Interventions that directly impact supply and demand variability.

Blood banks regularly attempt to reduce donation process variability by influencing donor behavior through outreach: phone calls, advertising and public appeals for donations (Flavelle, 2020; Garcia, 2016; Reuters News Services, 2000; Zabarenko, 2014). These efforts are underlaid by research to understand potential impediments to donation (Asamoah-Akuoko, Hassall, Bates, & Ullum, 2017; Hawkins & Gillett, 2015; Suárez, Fernández-Montoya, Fernández, López-Berrio, & Cillero-Peñuela, 2004;

Titmuss, 1971). Much of the reported research focuses on cultural differences and beliefs which support or stigmatize blood donation.

On the demand side there is ongoing research to optimize transfusion protocols to improve patient outcomes. Changes to transfusion protocols affect inventory management systems by altering demand patterns. This line of inquiry has already resulted in a number of hospitals instituting patient blood management programs which aim to reduce the overall use of blood products. These programs do not directly reduce demand variability, but they directly reduce demand and would result in more slack in the inventory. In another direction, there is ongoing research indicating that faster administration of red blood cells, or administration red cells instead of plasma faster trauma, may lead to better patient outcomes (Brown et al., 2015; Powell Elizabeth et al., 2016), suggesting that red blood cells should be stocked on life flight helicopters, a challenging logistics problem.

Demand forecasting goes beyond examining current demand to attempt to predict future demand. This does not change demand but it does reduce uncertainty. Analogous to “preclinical” research, the direction of that research is to predict blood usage during specific procedures (Hayn et al., 2016) which rolls up into demand forecasting for a hospital, region or nation (Velindre NHS Trust, 2008). Forecasting is commonly perceived as a way to improve inventory performance (Fortsch & Khapalova, 2016); the research in the next chapter shows that that assumption is scenario dependent.

2.1.3. Inventory controls that respond to uncertainty

Many blood supply chain operations research (OR) solutions involve creating policies to mitigate the effects of uncertainty, rather than attempting to alter the system's stochastic inputs. This includes classic OR problems such as donor center process optimization, blood rotation optimization (Kendall & Lee, 1980), bloodmobile scheduling (Alfonso, Xie, Augusto, & Garraud, 2012; Cerveny, 1980), blood inventory routing (Hemmelmayr, Doerner, Hartl, & Savelsbergh, 2009, 2010) and, of course, classic inventory control.

In blood inventory control, important subproblems are crossmatching of blood units, replenishment (aka ordering), and allocation (aka issuing, picking, dispensing). Crossmatching describes both the medical process by which a unit of blood is determined to be compatible for a specific patient; and also the larger crossmatch and reservation process that incorporate this function into inventory management. Crossmatch and reserve policies have significant impacts on inventory control (Jagannathan & Sen, 1991; Perera, Hyam, Taylor, & Chapman, 2009). Optimal crossmatch procedures, both medical and operational, are ongoing research issues. The advent of Type and Screen (TS) technology has eliminated the need for crossmatch and reserve for the majority of patients (see Dillon, Oliveira, & Abbasi, 2017) but the technology has not been universally implemented.

Replenishment research is perhaps the most addressed inventory research topic. Replenishment policies are deeply tied to various inventory control issues including allocation policies. Allocation policies are addressed in Chapters 4 and 5 of this dissertation. Both chapters utilize active substitution of compatible blood types to improve inventory performance. Chapter 4 finds optimal stock level-dependent issuing policies for emergency conditions including type-based substitution but not including perishability. Chapter 5 finds improved stock level and age-based allocation policies for normal operations including perishability factors, type-based substitution and replenishment units of variable age. None of the existing literature incorporates all of these features.

2.2. Methods for Capturing System Complexity

Blood inventory management is complex. The highly variable supply and demand processes, plus perishability, plus substitutability, plus the high stakes of a stockout differentiate the problem from standard inventory control. It is known that some standard inventory management practices, such as JIT ("just in time") policies, are not applicable to blood, mainly because of the high stakes of stockout and the complex interactions in blood supply chains (Chapman, Hyam, & Hick, 2004). Most perishable inventory research does not consider substitutability and perishability simultaneously but one (non-blood) study which considered both found that use of standard perishable ordering policies could result in "pathological behavior" when used with substitutable goods (Deniz, Karaesmen, & Scheller-Wolf, 2010). Many common assumptions are incorrect and lead to suboptimal conclusions (Cattani, Jacobs, & Schoenfelder, 2011).

2.2.1. Simulation to model an existing process and see the effect of new policies

Nahmias (1982) cautioned that “Although many of the existing theoretical perishable inventory models discussed in this review have been suggested by blood banking problems, the actual problem, even at the hospital level, seems too complex to be adequately described by any single mathematical model.” Simulation is a modeling approach which can capture complex interactions and large amount detail. In that capacity simulation has been used extensively in blood research to understand the effect of new inventory polices and to compare multiple existing policies. In conjunction with dynamic programming, it is also used to develop the policies themselves.

Simulation allows understanding of large scale systems such as regional blood supply chains (Mustafee, Taylor, Katsaliaki, & Brailsford, 2009) or smaller scale but still complex systems. Simulation can be used to predict the effect of proposed system changes before they are implemented as in Blake and Hardy (2013) which looked at the impact of the proposed consolidation of two blood centers in Canada’s maritime provinces. Simulation permits exploration of how policies interact such as in Baesler et al. (2014) which used simulation to analyze different combinations of inventory policies in a blood supply chain or Costantino et al. (2014) which used simulation to explore the interaction of collaboration with coordination in a supply chain. This is similar to how simulation is used in Chapter 3. Chapter 3 does not suggest new policies; it uses simulation to describe the potential impact of the integration of demand forecasting under a variety of scenarios.

Simulation can also be used to compare policies newly developed with other methods such as in Janssen et al. (2018) which used mathematical programming to find ordering policies for perishable goods but used simulation to compare the policies. Van Donselaar and Broekmeulen used simulation to compare age dependent replenishment policies (2009). This is similar to the way simulation is used in Chapters 4 and 5. In the fourth chapter simulation is used to evaluate the efficacy of new policies. Simple Monte Carlo simulations as well as a more complex Arena simulation model are used. The simple simulations allow for speed while the increased detail of the Arena model adds confidence that the new policies would perform well in a real system. In the fifth chapter simulation is used both in the creation of new policies and in the testing of those policies.

The original research that motivated this dissertation used a detailed discrete event simulation (DES) Arena model to find the effect of a proposed blood shelf life reduction. Model inputs were derived from real data. This model is used in chapters 3 and 4 for the comparison of policies in an environment with more detail than is normally possible. An overview of the original project is also given later in this chapter.

Simulation Optimization

Simulation Optimization is a method for finding new policies that relies on simulation, enabling it to approximate complex systems and interactions but without guarantees of optimality. Duan and Liao used simulation optimization to get better replenishment ordering policies for a non-specific perishable product with a short lifetime (2013) and specifically for red blood cells (2014). Afshar and Haghani (2008) worked on a simulation-optimization framework for an emergency evacuation problem. The work by

Haijema et al. (2007) on improved ordering polices for platelets in began as a dynamic programming problem but when the state space proved too large it was solved as a simulation optimization problem.

2.2.2. Markov Decision Processes (MDP) and Approximate Dynamic Programming (ADP)

MDPs mirror the sequential decision making of a real allocation process and allow for complex system dynamics to be captured. They permit significant flexibility in system representation and solution methods. Dynamic programming methods have been used to examine the management of blood inventory with multiple sources of supply (Puranam, Novak, Lucas, & Fung, 2017) and to develop state-dependent emergency blood order policies (Johansen & Thorstenson, 2014). Haijema (2011) is a well-known study which uses dynamic programming to find improved issuing policies for perishable inventory product with short maximal shelf life (platelets).

Approximate Dynamic Programming (ADP) is used when exact dynamic programming becomes intractable, usually due to problem size. ADP has seen extensive use in operations research the last decade with applications including modeling a typical blood platelet bank (Abdulwahab & Wahab, 2014), multilocation inventory transshipment (Meissner & Senicheva, 2018) and multivehicle stochastic routing (Goodson, Ohlmann, & Thomas, 2013).

There are many different approximate dynamic programming methods. When modeling extremely complex MDP, simulation can be used with ADP for model-free ADP. Model-free ADP was selected for application in Chapter 5 because the problem is too large for exact dynamic programming and the transition probabilities are too complex to be captured directly. This approach works because simulation is able to capture complex system dynamics and bring them into the ADP model.

Simulation based ADP has not been used extensively in blood literature. One of the few healthcare related simulation based ADP models in the healthcare literature is Astaraky and Patrick (2015) on the scheduling of surgical suites. A broad overview of approximate dynamic programming can be found in Powell (W. Powell, 2011) and Bertsekas (2012, 2019b).

The model-free ADP approach of Chapter 5 uses a machine learning approximation architecture. A review of some advanced approximate policy iteration methods including machine learning/artificial intelligence architectures can be seen in (D. Bertsekas, 2011; D. P. Bertsekas, 2018) and (W. B. Powell, 2010).

Garbage in, garbage out

Heeding the adage ‘garbage in, garbage out”, significant effort was made to find reliable data sources to support choice of model inputs instead of making assumptions. Examples of inputs include blood type distributions in Punjab (Chandra & Gupta, 2012) historical influenza infection rates (Britten, 1932; Sharrar, 1969), relevant time horizons of emergency blood transfusions (E. K. Powell et al., 2016) and the number of units typically required by a victim of gun violence (Demario et al., 2018) among others.

2.3.Simulation to Show Effect of Proposed Shelf Life Changes[1]

In response to concerns about a possible reduction in allowable red cell shelf life Mayo Clinic sponsored work to determine how a maximum shelf life reduction would affect the ability to provide for patients and how it would change the amount of blood required to do so. An Arena model was developed to represent the transfusion department blood inventory management process at the Rochester hospital as closely as possible.

All blood units that enter the transfusion center at Mayo are recorded. From these records, data on blood type, age at receipt, date of receipt and whether it was transfused or outdated were collected. Transfused unit records were matched with de-identified patient records. From this information it was possible to identify the number of patients arriving each day and what their general characteristics (condition, gender, blood type, number of units transfused, time of stay) were. Analysis showed that the number of patient arrivals per day varies by day of the week. The historical number of units for each day was entered into Arena input analyzer which provided basic distribution definitions. Chi-square tests showed that the suggested distributions were sufficient. These distributions serve as the patient arrival distributions in the model.

In the simulation model each patient is assigned a blood type and membership in one of six condition categories depending on their age, gender and medical condition. This is necessary because, along with blood type, the patient's condition determines which units of blood the patient is permitted to receive. At the time of the study a requirement of younger units for cardiac patients was being consider. In the model, blood

[1] This section summarizes (Dumkrieger et al., 2014)

types and patient conditions are assigned according to historical percentages. “Patient condition” also determines how many units of blood the patient requires. Records were analyzed to determine the number of units of blood that each patient requires. Including detail of patients and number of units needed over time is important because while the majority of patients will only require one or two units some will require many times that amount. These demand “shocks” may have a different effect on inventory and need to be considered.

Blood units arrive at the hospital either via a stochastic process representing Mayo’s internal donor center, or the units arrive deterministically as if ordered from the Red Cross. The only historical information specific to the Red Cross is the age of the units they provide. It is assumed that all blood requested is received. The complex age distribution of new units is a unique feature in this research.

The actual inventory policy in place at the time of the study was used as a baseline. That policy was a “periodic review and order up to” policy with two levels, one for regular orders which occur three times per week and another level for extra orders which occur as needed. Current inventory and anticipated arrivals from the internal donation center (undergoing processing when the order is placed) are considered when determining order quantity. Each order is received the day after it is placed.

After a patient has arrived and been assigned the appropriate characteristics the existing inventory is searched to find the oldest compatible blood according to the specified allocation policy. The pattern of blood substitutions is one of the key elements that makes simulation the most appropriate tool for analyzing blood inventory problems.

A randomly determined percentage of the total units demanded are used on each day until there is no outstanding demand.

The model was internally validated. Producing the correct total number of units per year confirms that the patient generation function and the units/patient function are interacting correctly in the simulation model. Correct distributions of blood types, patient categories, age of donated units were also confirmed. One potential shortcoming of the model is that it did not consider temporary allocation of cross-matched blood in advance of a procedure. However, the advent of type and screen technology means that for patients without rare antigens crossmatch and temporary reserve is unnecessary (see Dillon, Oliveira, & Abbasi, 2017).

Seven scenarios were run. Each scenario has a different shelf-life ranging from the current 42 days down to 7 days. Scenario 6 has a 14-day shelf life limit for blood given to cardiac patients and a 42-day limit for blood given to other adults. Results are summarized in the table below.

Table 1

Results of Previous Research on Effect of Change in Maximum Shelf Life

		Red Cross Orders		Transfused			Outdated						
Scenario	Age Limit	# Units	Average age at Receipt from RC	# Units	Average Age at Receipt	% Matched Units from RC	# Units	% Total Units outdated	% Outdated Units from RC	Internal Units Average Age at Receipt	RC Units Average Age at Receipt	Overall Average Age at Receipt	Avg # Units in Inv.
1	7	76200	5.79	41,273	5.5	95%	63,773	61%	42%	2.6	5.9	4.5	720
2	14	54509	11.05	41,283	9	83%	42,072	50%	52%	2.5	12.4	7.2	1,374
3	21	22958	14.48	41,283	8.3	49%	10,521	20%	76%	2.7	17.7	6.3	1,698
4	28	13903	16.25	41,267	7.1	33%	1,482	3%	85%	2.7	20.9	5.5	1,767
5	35	12809	17.14	41,268	7.1	31%	387	1%	92%	2.8	27.2	4.8	1,793
6	mixed	12802	17.27	41,265	7.2	31%	384	1%	95%	2.8	20.6	3.6	1,813
7	42	12581	17.29	41,271	7.1	30%	157	0%	98%	2.8	23.2	3.3	1,812

The most important result of this work is that under a seven-day shelf life Mayo would need to request more than six times as much blood from the Red Cross as they currently do. The vast majority of these units would never be transfused and would go to waste. Even a more moderate halving of the shelf life to 21 days would nearly double the amount of blood Mayo would need.

In this simulation model no limits were placed on quantity ordered from the Red Cross; in this case Mayo would be able to provide for all patients even with 7 days shelf life. However, if there was a limit on the quantity, as there surely would be in real life, a tradeoff would have to be made between the number of units being outdated and the number of patients who will not served

This research showed that while it may be possible for the Red Cross to provide for Mayo’s needs under a shelf-life reduction scenario, it would not be possible for the Red Cross to provide for all of the hospitals it serves without major system changes.

The Arena model used in this analysis was used with modifications in Chapter 3 to compare the possible effects of demand forecasting and in Chapter 4 to compare the performance of issuing policies both under emergency conditions and under normal operations. The overall structure and level of detail of the model remain but some aspects were altered in some scenarios to allow for the comparisons to be made. Any changes are discussed in the appropriate section.

CHAPTER 3

USE OF SIMULATION AND FORECASTING MODELS TO EVALUATE THE BENEFIT OF FORECAST IMPROVEMENT ON BLOOD INVENTORY MANAGEMENT

3.1. Introduction

Management of blood inventory is very complex. Part of the complexity is due to the perishable nature of the product. Further contributing to the complexity is the stochasticity of demand for blood products and that serving this demand is mortally important.

Within the field of blood inventory management, one area that warrants further exploration is forecasting the demand for blood based on surgical schedules. It is not common practice for hospital blood banks to attempt to forecast demand based on surgical schedules or even to adjust their ordering based on the surgical schedule. It seems intuitive that by accurately forecasting demand, blood inventory management can be improved. This chapter evaluates whether forecasting improves inventory management and if it does, how much does the prediction accuracy impact the benefit.

Red blood cells (RBCs) are of critical importance and the World Health Organization (WHO) recommends hospitals keep 7 days of inventory on hand. Hospitals try to balance having sufficient resources on the shelf to serve their patients while minimizing wastage and associated costs.

Some blood units are required for critical emergency needs, such as for the treatment of trauma victims. The other portion is used in elective procedures which are scheduled in advance. Elective demand presents an opportunity to reduce system uncertainty through forecasting. Reducing system uncertainty would be expected to improve inventory performance. However, forecasting blood unit demand from scheduled procedures is quite complex as there are many different medical procedures that require blood and many factors, such as patient body weight, that affect the quantity of blood required, even for the same procedure.

Before undertaking arduous efforts to forecast blood usage based on surgical schedules it will be useful to know if the expected benefits will make these efforts worthwhile. This chapter uses detailed simulation models to clarify the effect of demand forecasting on inventory performance as measured by outdated units and unmet demand.

3.2. Literature on Forecasting Perishable Demand

3.2.1. Product Perishability.

Perishable inventory control is a well-researched topic of study and blood (both red cells and platelets) inventory is frequently the subject product. There are two main objectives in blood management: (1) minimizing outdating, and (2) maintaining service level.

Demand uncertainty is the most significant source of inventory inefficiency at the hospital level. Demand uncertainty can be attributed to the stochasticity and variability inherent in the demand stream. Reducing demand variability itself is impractical, but implementation of accurate forecasting may reduce the impact of the variability and produce better system performance (Fildes & Kingsman, 2011).

3.2.2. Demand Uncertainty and Forecasting Error.

Several authors have studied the value of improved demand forecasting and/or the effect of demand uncertainty in manufacturing systems. Fildes and Kingman (2011) found that in a simulated two-level Material Requirements Planning (MRP) model with lot sizing, improved forecast accuracy could substantially reduce unit cost, particularly in the face of high demand uncertainty. Others have considered demand uncertainty and forecasting in conjunction with process uncertainty (Germain, Claycomb, & Dröge, 2008) or environmental uncertainty (Wong, Boon-itt, & Wong, 2011) in the supply chain and found links between higher uncertainty and higher costs. Still others have studied the interaction of demand variability with transshipment and its spread and amplification through the supply chain (X. Chen, Hao, Li, & Yiu, 2012).

The potential value of forecasting in blood inventory management is insufficiently explored. Much of the existing perishable (and blood) inventory literature assumes a heavily "stylized" environment (Broekmeulen & van Donselaar, 2009), often assuming single products, lack of substitution, no lead time, uniform age at arrival to inventory, or only a two period shelf life. While these assumptions are necessary for the development of tractable models, it is not clear how implications on the benefits of forecasting hold up in the complex environment found in practice. RBC management involves eight products with complicated substitution patterns, a 42-day shelf life, complex age distributions of current inventory, and other nuances.

3.2.3. Prediction of blood demand.

Forecasting blood demand is not unusual at the regional (Fortsch & Khapalova,

2016) or national (Lestari, Anwar, Nugraha, & Azwar, 2017) levels, but these forecasts have long horizons and do not include local schedules as input. In a detailed discussion of the current processes for prediction of blood needs by the National Blood Service of England and North Wales (Chapman et al., 2004) it was noted that demand forecasting occurs annually with quarterly adjustments. These regional long-horizon forecasts are akin to tactical level forecasts whereas this chapter is concerned with the short-horizon hospital level forecasts for operational-level decision making.

The need for improved demand forecasting has been recognized by organizations responsible for such planning (Velindre NHS Trust, 2008, p. 117) but there is little evidence of operational level forecasting occurring in practice (W. P. Pierskalla, 2004, p. 122; Stanger, Wilding, Yates, & Cotton, 2012). Many of the models that have been developed predict blood usage for only a single, specific procedure type, such as “Major Head and Neck Surgery with Free-Flap Reconstruction” (Shah et al., 2010) or total hip replacement (Hayn et al., 2016) or are retrospective reporting of blood requirements (Demario et al., 2018). Some models predict whether or not transfusion will be needed (Shah et al., 2010) while others attempt to predict the number of units needed (Hayn et al., 2016). Forecasts of demand quantities would be a more useful outcome for the purposes of the scenarios considered in this chapter.

One notable reference, Guan et al. (2017), used historical data to develop a predictive demand and ordering model for platelets based on current patient data. The model estimated that the developed ordering policies could reduce outdating from 10.5% to 3.2 %. Platelets have a much shorter shelf life; however, the approach could be applicable to RBCs as well.

While the approach may work for RBCs it's not clear that the benefit will be worth the effort. Inghilleri (2010) defined three primary approaches to forecasting blood demand, with increasing complexity, and noted that more accurate methods require patient-specific information. Increased information collection comes with a cost. The following section describes a model of forecast error that will be used along with a detailed simulation model to investigate how various inventory characteristics and inventory management decisions interact with improvements in forecast accuracy. The detailed simulation model makes it possible to get a sense of the potential benefits of forecast improvement in a realistic environment.

3.3. Forecast and Simulation Models

3.3.1. Forecasting error and reorder quantities.

In this approach, the patient arrival stream to the hospital is generated randomly. Some portion of patients are deemed to be seeking elective procedures and therefore their demand is "forecastable". Prior to their simulated arrival at the hospital the quantity of RBC units demanded by each elective patient is recorded and a forecast of their demand is calculated. Forecasted demands are recorded and used to set orders. Orders arrive prior to the patients' arrivals. Note that units are not earmarked for specific patients, only the quantity is set. The forecasted demand is modeled by adding the forecast error term, ϵ, to the actual demand quantity. Patient level forecast error ϵ is assumed normally distributed with mean μ and standard deviation σ, i.e., $\epsilon \sim N(\mu, \sigma)$. For an unbiased forecast $\mu = 0$. The equations (3.1) – (3.6) below describe the forecasting model, using the following notation:

D_t: *Total demand at time t*
d_{ptx}: *demand for scheduled (elective) patient x at time t*
d_{ety}: *demand for emergent patient y at time t*
F_t: *Forecasted elective demand at time t*
ϵ: *forecast error*
pct.pred: *% demand that is scheduled (predictable)*
ss_e, ss_p: *safety stock percentage multipliers; emergent, elective*

$$D_t = \sum_x d_{ptx} + \sum_y d_{ety} \tag{3.1}$$

$$F_t = \sum_x \max\,(d_{stx} + \epsilon, 0) \tag{3.2}$$

$$\epsilon \sim N(\mu, \sigma) \tag{3.3}$$

$$\sum_x d_{ptx} = D_t * pred.pct \tag{3.4}$$

$$Inventory\ set\ points = Baseline\ set\ point * (1 - pct.pred) * (1 + ss_e) \tag{3.5}$$

$$Order_T = \ Inventory\ set\ points - \ Current\ Inventory_T + ss_p * \Sigma_{t \in (T-1,T]} F_t \tag{3.6}$$

where Order_T is the order quantity at period T, T=1, 2, …

In practice, the proportion of blood requirements going towards emergency procedures vs. elective procedures varies across hospitals and the model accommodates this with the forecastable demand percentage variable (pred.pct). Inventory set points cover only the non-forecastable demand and remain constant throughout a single scenario. When non-forecastable demand is present, new inventory target levels are baseline setpoint quantities multiplied by the percentage of non-forecastable demand plus a safety factor. Order quantities equal the forecasted quantity plus the order quantity from the non-forecastable portion along with any safety factors, (see equation 3.6). Safety stock multipliers are included separately for forecasted demand portion and for the unscheduled demand portion. Both positive and negative safety stock values are explored below.

3.3.2. Hospital simulation model.

The operations of a single hospital's blood bank were modeled using Arena simulation software (Rockwell Automation, 2016) as described in Chapter 1. The model has been verified and validated with historical data. This model features random streams of adult patients of several disease categories. Each patient demands one or more blood units over one or more days. For the analysis in this chapter supply is provided entirely by the regional blood center, that is, there are no donations coming directly to the hospital.

Baseline inventory target levels were set to meet the new demand with minimal outdating. In the preliminary scenarios (see Table 2 in next section) target inventory levels are constant and there is no demand forecasting.

Three sets of analysis were done. In the first analysis all demand is presumed forecastable and the forecast is assumed to be "perfect" with error mean and standard deviation equal to zero. This analysis assesses characteristics of the environment that affect the ability of a perfect forecast to result in perfect inventory performance.

The second set of analyses again assumes that all demand is predictable and that the forecast is unbiased but several error variances, along with several ordering safety factors are considered for two different demand levels. This analysis illustrates the impact of the error variance in conjunction with the safety factor on unmet demand and outdating.

The third analysis presents a designed experiment that seeks to evaluate jointly the impact of forecast error mean and variance, safety factors and percent of demand that is predictable for a single demand level. Total demand remains the same in each of the scenarios of the third analysis but some percentage of it is deemed to be from elective procedures.

3.4. Simulation Experiments

3.4.1. Preliminary experiments: Non-forecast related factors affecting forecast impact.

Real-world blood inventory systems are comprised of numerous environmental features in combination with multiple management decisions. Even in the presence of a perfect forecast the following features can prevent perfect inventory performance:

- Order timing and lead times
- Age of received replenishment units
- Forecast horizon
- Internal age-based queues
- Patient demand forecasted once, but demand extends over multiple days.

These factors are interrelated. If the time between orders is longer than the forecast horizon then the orders will not be in inventory when the demand occurs. If the forecast horizon is long and the age of arriving replenishment units is high, then the units may outdate before the intended recipient arrives. If there are internal age-based queues (such as set-aside of fresher units for pediatric patients) and there is substitution in the queues of older units before searching the queues of fresher units, the assignment of units to their "intended" patient may become suboptimal.

<u>Baseline (Perfect Performance).</u>

To establish a baseline with zero units outdating and zero unmet demand, the following assumptions were made:

1) Age of replenishment units upon receipt was set to 2, the approximate time required for processing newly donated units.

2) The forecasting horizon was set to 7 days; and

3) Orders are permitted (and occur) each day.

Under these conditions and with a perfect forecast, supply and demand match and there is no outdating and an average of 0.1 unit of unmet demand per simulated year (Table 2, row 1). Patient demand can occur across multiple days and that does not negatively impact the outcomes here, but under conditions of different distribution of replenishment unit ages or different forecast horizons, multiple-day patient demand could prevent perfect performance.

<u>Limited Order Frequency.</u>

Limiting order frequency to a scheduled three days per week results in 1.4 patients per year with unmet demand if the orders are scheduled on days 2,4 and 5 (Monday, Wednesday, Thursday) (see Table 2, row 2), but less than a single-patient unmet demand if the orders are schedule on days 2,4,6 (Table 2, row 3). Both these values are very small, but it illustrates the sensitivity of the problem to inventory management decisions.

Table 2

Exploratory simulation results showing effect of reorder schedule and replenishment unit age at delivery on annual average unmet demand and outdating.

Reorder Schedule, Day of Week	Unit Age at Delivery	Unmet demand (patients)	Unmet demand (units)	Outdated Units
1-7	2	0.1	0.1	0
2,4,5	2	1.4	4.1	0
2,4,6	2	0.3	0.4	0
2,4,6	28	0.3	0.4	0
2,4,6	35	114	416	390
2,4,6	Complex* (mean 18.01)	5	15.2	18.2

*Age distribution varies by blood type. For description of complex age distribution see (Dumkrieger, Huschka, & Stubbs, 2014).

<u>Increased Age of Supplied Units.</u>

Broekmeulen and van Donselaar (2009) addressed the stylized environments found in perishable inventory research and noted the importance of considering detailed age distributions. Rows 4, 5 and 6 of Table 2 show the results of adjusting the supply age. When the units delivered are 28 days (4 weeks) or younger, the unmet demand and outdating are not affected. A fixed age of 35 days (5 weeks) or younger results in a significant increase in outdating. When the supply has a complex, type-dependent age distribution, there is a mild increase in unmet demand and outdating compared to the scenario when the replenishment unit age is 28 days. This complex age distribution is from historical data and is used in subsequent experiments.

3.4.2. Effect of error variability under 100% predictable demand.

The second analysis also considers the case when the demand is entirely elective (i.e. it is forecastable) and tests combinations of the standard deviation of the forecast error (error-SD) and safety factors at two different demand levels. The first case, the high-demand case, is where demand is very large (approx. 8,100 patients, 39,600 units per year), and the second, low-demand case, is much lower demand; it has 10% of the demand of the first case. The mean of the forecast error is held to be 0 in all scenarios of this analysis. As in the latter rows of Table 2, three orders are placed per week, on days 2, 4, and 6. Lead time is 24 hours and forecast horizon is 7 days. The age of the units upon arrival are based on empirical data and vary by type. There are no neonate or pediatric patients. The mean units/patient is 4.88; with this value in mind a variety of error standard deviations between 0 and 4 were selected for testing. Several safety factors were tested: -1%, 0%,1%,10%; as a percent of forecasted demand and added to the forecast demand.

The case where the error mean and standard deviation are 0 represents perfect forecasting (Table 3, row 9). Unmet demand and outdating in that scenario are the results of what may be called the “friction” of the system

Table 3

Results of simulations for varying error-SD and forecastable demand safety factor (safety %). Row 9 shows results of simulation with 100% accurate forecast.

			HIGH DEMAND			LOW DEMAND		
	error -SD	safety %	outdated units	unserved patients	unmet units	outdated units	unserved patients	unmet units
1	0.0	-0.01	18.2	135.4	487.8	30.9	27.7	70.3
2	0.01	-0.01	19.0	131.8	487.5	28.6	61.6	188.0
3	0.1	-0.01	19.8	130.7	486.0	31.3	56.6	164.4
4	0.25	-0.01	20.0	125.0	482.3	29.7	51.2	151.7
5	0.5	-0.01	18.4	134.1	475.0	27.0	48.4	145.1
6	1.0	-0.01	18.8	122.9	461.6	28.0	49.2	143.8
7	2.0	-0.01	29.5	114.5	437.1	41.0	39.3	127.4
8	4.0	-0.01	111.8	89.5	353.7	141.1	22.6	73.8
9	**0.0**	**0.00**	**17.2**	**4.8**	**15.9**	**36.3**	**16.8**	**38.5**
10	0.01	0.00	46.6	38.3	136.0	27.7	38.7	106.3
11	0.1	0.00	39.3	36.8	131.5	27.5	39.2	106.0
12	0.25	0.00	44.6	38.8	128.6	25.8	37.8	103.9
13	0.5	0.00	49.4	36.3	131.9	29.2	38.4	106.7
14	1.0	0.00	60.8	34.5	124.7	31.4	36.4	107.9
15	2.0	0.00	77.2	26.1	95.5	45.9	31.8	96.4
16	4.0	0.00	225.0	22.8	78.1	162.5	16.2	53.9
17	0.0	0.01	293.6	0.0	0.0	31.6	14.5	33.3
18	0.01	0.01	292.7	0.0	0.0	33.8	15.0	35.3
19	0.1	0.01	276.7	0.0	0.0	32.1	23.5	58.1
20	0.25	0.01	241.2	0.0	0.0	30.6	29.0	70.0
21	0.5	0.01	247.9	0.0	0.0	32.3	26.6	72.8
22	1.0	0.01	268.2	0.0	0.0	35.6	26.4	72.3
23	2.0	0.01	310.1	0.1	0.1	52.6	20.3	63.5
24	4.0	0.01	494.1	0.3	0.7	188.1	11.2	38.3
25	0.0	0.10	3735.9	0.0	0.0	312.6	0.1	0.1
26	0.01	0.10	3729.2	0.0	0.0	306.2	0.1	0.1
27	0.1	0.10	3728.8	0.0	0.0	304.5	0.0	0.0
28	0.25	0.10	3730.8	0.0	0.0	304.3	0.1	0.1
29	0.5	0.10	3742.3	0.0	0.0	302.9	0.1	0.1
30	1.0	0.10	3756.1	0.0	0.0	309.3	0.1	0.1
31	2.0	0.10	3805.6	0.0	0.0	341.5	0.4	0.9
32	4.0	0.10	4018.2	0.0	0.0	521.6	0.5	2.0

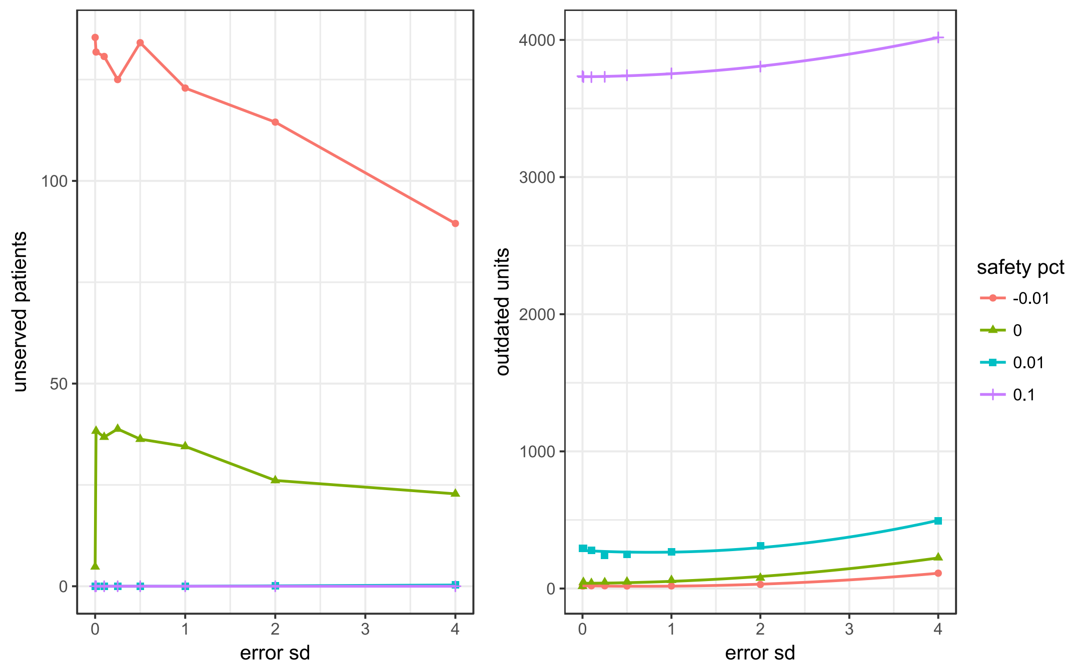

Figure 3. Plot of outcomes vs error SD for high demand case. Outdated units plot shows quadratic fit line.

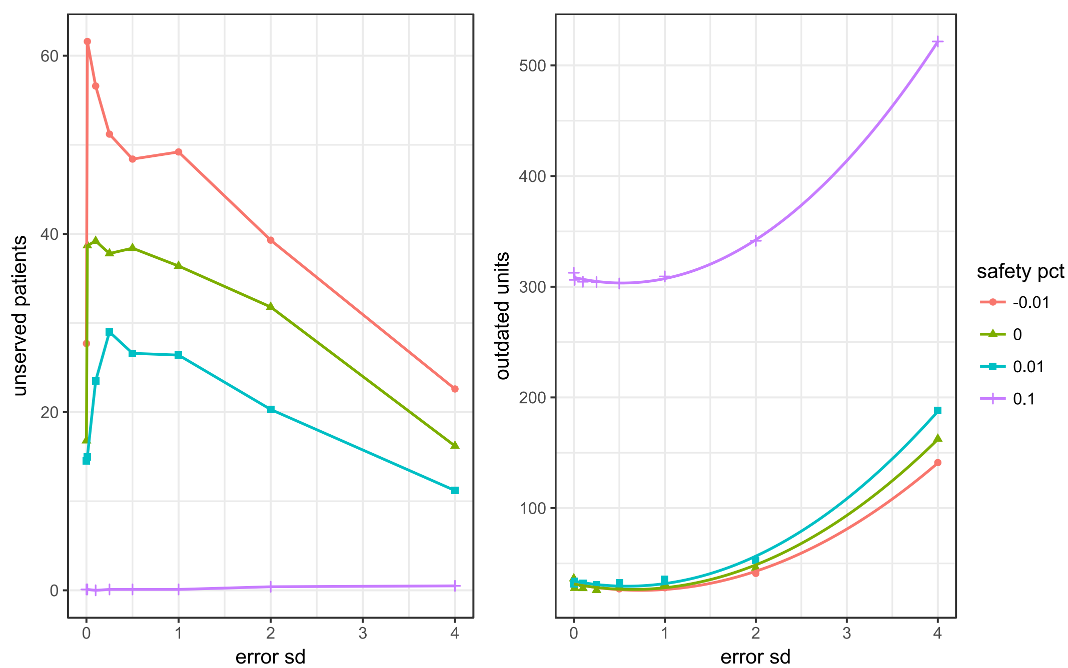

Figure 4. Plot of Outcomes Vs Error SD for Low Demand Case. Outdated units plot shows quadratic fit line.

There is a quadratic relationship between outdating units and forecast error-SD which is particularly strong in the low demand case (Figure 3 and Figure 4). In the low demand case, the R^2 of a quadratic polynomial fit to the results is above 0.99 for all % safety stock. The relationship is less clear for the high demand case with R^2 ranging from 0.96 to 0.99.

As seen in Table 2, Figure 3 and Figure 4, having high, unbiased forecast variability makes less difference in outcomes when demand is high. The effect is much greater for the lower demand case, as seen in the comparatively steep curves in Figure 4. Except for simulations with negative safety %, the low demand case has unmet demand equal to or greater in unmet demand than the high demand case. The increased "friction" in the lower demand case is seen in the proportionally higher unmet demand in the perfect forecast scenario (see Table 2, row 9).

It is important to note that the scale of unmet demand is very different than of outdating units. In the high demand case, the maximum unmet demand is only 1.2% of units demanded (1.7% of total patients) , and occurs when the error-SD equals 0 and the safety factor equals -10% (Table 2, row 1). In the low demand case the maximum unmet demand is 4.7% of demanded units (7.6% of patients) which occurs when the error-SD is 0.01 and the safety factor is -1% (Table 2, row2). These shortages may be small enough that they could in practice be overcome by a few emergency orders, but emergency orders were not allowed in this model.

Unmet demand initially increases with increasing error-SD but then decreases. The trend of decrease in unmet demand with increase in error SD may seem counterintuitive. The effect appears to be due to a flattening out of under-estimated forecasts. When the forecast overestimates, the resultant excess supply remains in inventory for a significant period of time and is available to serve demand whenever the forecast underestimates. Pooling demand from multiple patients means the sample bias will tend toward zero but because the minimum order is 0 the bias shifts towards the positive, shifting the average ordered quantity higher than the forecasts would have recommended. Younger age of units at delivery would be expected to increase this effect.

Units demanded per patient are drawn from the same distribution in the high and low demand cases, but because there are more patients in the high demand case, the overall demand variation is less. In the high demand case, the coefficient of variation for daily demand is 0.38 (weekly: 0.11). For the low demand case the coefficient of demand for daily demand is 0.71 (weekly 0.25). The distributions for the high demand case are more normal and symmetric (see Figure 5) while the demand is more skewed in the low demand case (see Figure 6). The higher demand variability in the low demand case translates to higher variability in order size (high demand order size coefficient of variation: 0.19 – 0.22; low demand coefficient of variation: 0.44-0.52) and the risk-pooling effect is diminished in the low demand case.

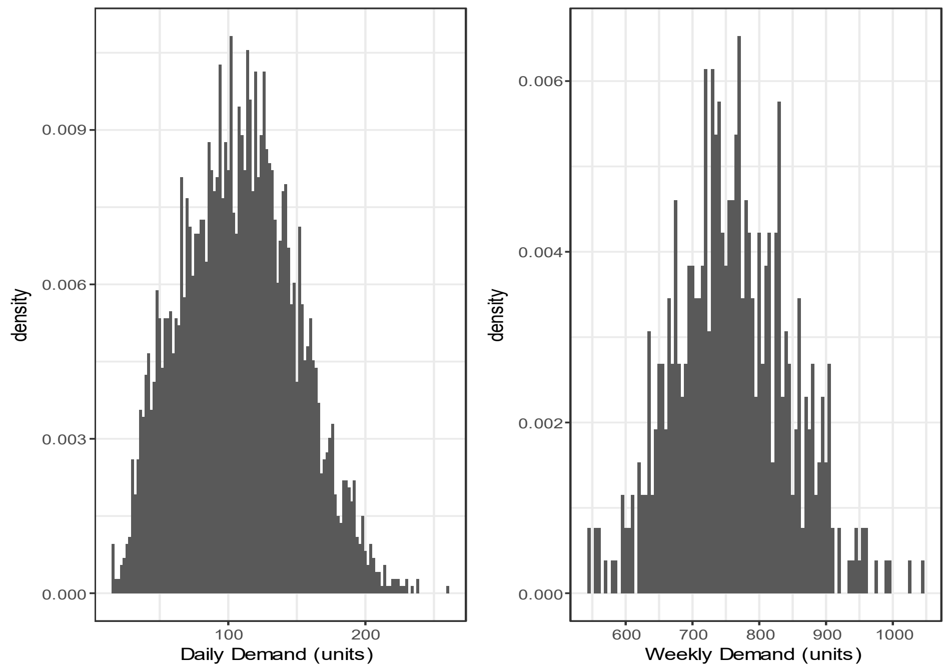

Figure 5. Daily and Weekly Demand Distributions for High Demand Case

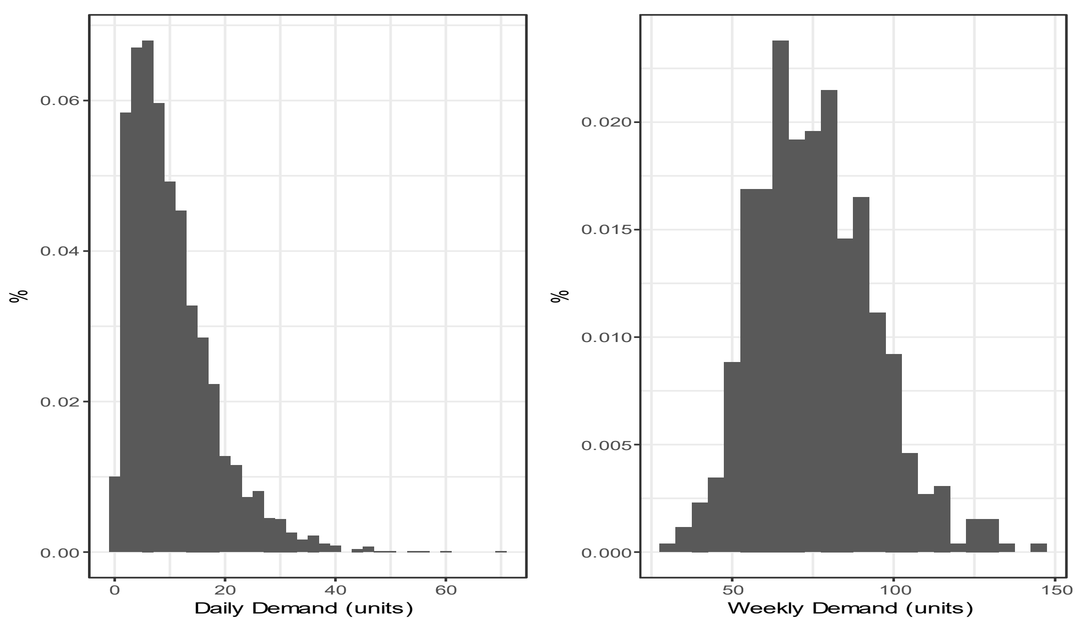

Figure 6. Daily and Weekly Demand Distributions for Low Demand Case.

Further Discussion of the Simulation Results.

High-demand hospitals will experience less variation in total demand even with the same unit-per-patient demand distribution. They will experience less negative impact from forecast error variance increase and less benefit from forecast improvement, when the forecast is unbiased.

3.4.3. High demand case, designed experiment.

Not all hospitals face entirely forecastable demand. In the experiment discussed in this section the demand quantity is fixed to the high demand level but the percentage of the demand, which is elective, and therefore forecastable, varies. Five variables in the experiment were that were allowed to fluctuate are:

$$forecast.pct \in [0.1, .9]$$
$$ss_p, ss_e \in [-1\%, 1\%]$$
$$\mu_{error} \in [-1,1]$$
$$\sigma_{error} \in [0,4]$$

Forecast.pct is the percent of demand that is forecastable, ss_p and ss_e are the safety factors of forecastable and emergency demand respectively, μ_{error} is the mean of the forecast error and σ_{error} is the standard deviation of the forecast error.

A central composite design with two repeated center runs was created in JMP (SAS Institute Inc., 2018). Simulations were run according to the design matrix for ten simulated years. Results for responses of outdated units and unserved patients are shown in Table 4.

Table 4

Coefficient estimates and variable significance from analysis of variance on results of experiment on effect of percent of demand that is predictable, forecast error mean and SD, and safety percentages on high demand case. Two analyses were performed, one with outdated units as the response variable and one with unserved patients as the response variable. Variables centered prior to Anova.

	Outdated Units	*Unserved Patients*
Intercept	-465.4	-902.4
forecast.pct (0.1,.9)	**5867.9***	**4367.6***
$\sigma_{error}(0,4)$	280.5	-102.6
ss_e	128.1	-1.2
ss_p	1167.9	-180.2
$\boldsymbol{\mu_{error}}$	**7633.2***	**-4353.4***
$forecast.pct * \sigma_{error}$	266.9	-115.4
$forecast.pct * ss_e$	-50.1	-1.3
$\sigma_{error} * ss_e$	-1313.9	-202.7
$forecast.pct * ss_p$	1292.4	-202.7
$\sigma_{error} * ss_p$	-44.2	-1.3
$ss_e * forecast.pct$	-252.9	-115.4
$\boldsymbol{forecast.pct * \mu_{error}}$	**8518.3***	**-4897.4***
$\sigma_{error} * \mu_{error}$	282.9	115.4
$ss_e * \mu_{error}$	37.2	1.3
$ss_p * \mu_{error}$	1295.4	202.7
$forecast.pct^2$	2863.7	1256.6
σ_{error}^2	2171.2	1128.6
ss_e^2	1917.2	1128.6
ss_p^2	1993.7	1128.6
$\mu_{error}{}^2$	1974.7	1129.1
R^2	0.93	0.94

'' indicates significance p<=0.05*

Forecastable demand percent (forecast.pct), the forecast error mean, and the interaction of these two factors are the only significant terms ($p<=0.05$) in this analysis.

Part of the difficulty in improving outdating and unmet demand when only part of the demand is forecastable stems from the problems in determining how to order for the forecastable and non-forecastable demand portions together. Most facilities operate with an order-up-to policy. When the entirety of the demand is forecastable the solution is straightforward. In this analysis the order quantities were set according to equation 3.6 when there was unforecastable demand, but a better method certainly exists.

Because only one large hospital setup was considered in this section, the generalizability of the results of this section to other configurations is limited. As seen in subsection 3.4.2 a small hospital would be expected to be more strongly affected by forecast error. The model used in the experiment in this section s was of a large facility with many different patients' demands per day. This results in less variability in total demand, which makes forecasting less valuable.

3.5. Discussion of Simulation Findings

There are many factors which impact the effect of forecasts on inventory performance. Beyond the accuracy of the forecast, other factors, such as forecast horizon, are within the control of the hospital, while others, such as age of units at delivery, may not be controlled by the hospital.

A perfect forecast, under otherwise realistic conditions, allows near perfect inventory performance when the entire demand is forecastable, but application to scenarios where only part of the demand is forecastable is more difficult.

The benefits of forecasting are muted when forecastable and emergency demand are present simultaneously and both demand types are served by the same hospital inventory and reordering is based on inventory on hand. The benefit of forecasting is greater for small hospitals and when the average age of units delivered is relatively older.

These analyses do not include inventory and management costs as they are notoriously difficult to quantify in blood inventory management and are highly variable across institutions. Individual facilities must consider these results based on their own individual costs.

Another concern when considering forecasting demand based on surgical schedules is the uncertainty surrounding the feasibility of creating forecasts of sufficient quality. Reported procedure-specific forecasting efforts have not been of high accuracy (Hayn et al., 2016) but this may change if or new forecasting methods are developed. The cost of developing accurate forecasts can be expected to vary by facility, for example based on the availability of data on which to make prediction models or on the variety of procedures performed. As with the inventory management costs, facilities need to consider the potential costs of improved forecasting for their specific institution.

CHAPTER 4

MDP MODEL TO FIND OPTIMAL ALLOCATION UNDER EMERGENCY

4.1.Introduction and Problem Definition:

The majority of studies on blood management assume the perspective of a robust blood supply system, such as found in the U.S., where the availability of blood for transfusion is taken as a given. Well-developed volunteer donor networks and supply chains mean that patients can reasonably expect that, should they need it, blood will be available, and the transfusion will be safe. Correspondingly, the analytic literature on blood management focuses on minimizing outdating while maintaining a satisfactory level of service. Common assumptions include the widespread availability of blood at a system level, that it can be reliably ordered and received, and that the costs of being unable to instantaneously meet the demand are not substantial. The ability to backorder demand is also frequently assumed. In places like the U.S., the assumption of the ability to delay demand without critical consequences is generally reasonable. This is partially due to the prevalence of elective (non-emergency) procedures which can be rescheduled if supply runs short. However, particularly in the developing world, shortages are common and delayed satisfaction of demand can result in patient death. (Schantz-Dunn & Nour, 2011).

Insufficient supply may be the result of chronic shortage or a temporary state due to an emergency such as a disaster that increases demand while reducing or eliminating supply. It is this second condition under which the U.S. is most likely to face a shortage. A less frequently cited Red Cross statement is “the blood used in an emergency is already on the shelves before the event occurs" (American Red Cross, 2016). When demand spikes because of an emergency, shortages can occur, and it is critical to allocate the units on the shelf in the most efficient way possible.

This chapter exploits the substitutability of blood types to find the best issuing policy to assign blood units to patients for the optimal utilization of a limited supply. Red blood cells have types. Each patient is of a single blood type and will be able to safely receive units of their own type or of a subset of the other substitute types. Most studies of blood management, even if some acknowledge substitutable nature of blood units, fail to consider the opportunity substitutability represents. Embracing substitutability is an opportunity to improve utilization.

Explicitly considering the substitutability of blood types, this chapter addresses the problem of how to best allocate the available blood units from the inventory at the beginning of an emergency event to a stream of patients arriving for treatment after the onset of the emergency. Because of the immediacy of the need, this chapter does not initially consider the perishability of blood products. However, it is demonstrated later in the chapter, through simulation tests, that the resultant policy is still effective in the face of perishability. In this problem, the blood type of the arriving patients is not known prior to arrival and patients are served as they arrive. The number of patients arriving is not limited; the intent is to serve as many patients consecutively as possible before encountering a patient that cannot be matched with a compatible unit.

4.2. Literature Review on Blood Allocation and Emergency Response

There are four general problems in inventory management: ordering policy, issuing policy, disposal policy and pricing policy (Shen, Dessouky, & Ordonez, 2011). For a review of inventory management of blood products see Beliën & Forcé (2012), Chapman, Hyam & Hick (2004); Prastacos (1984), or for review of perishable inventory, which often includes blood, see Bakker, Riezebos & Teunter (2012), Goyal & Giri (2001) and Nahmias (1982). This chapter is focusses on policies for issuing units of blood by type from available inventories to arriving patients.

4.2.1. Emergency Context

The context most directly addressed in this chapter is that of a limited supply in an emergency response situation where additional supply is not immediately available. A study of the effect of several manmade disasters (plane crashes, building collapse, terrorist attacks) on the U.S. blood supply indicated that in those events sufficient blood was available for emergency procedures and that public response to disaster leads to excess donations, most of which end up going to waste (Schmidt, 2002). However, the attacks studied occurred in densely populated areas with large standing inventories and ready access to additional supply. This research addresses the ability to utilize the blood on the shelves with greater efficiency, as would be necessary when increased demand is seen in an isolated area or where supply was already strained.

New shortages can arise from changes to supply. Several authors have studied the potential for disease to decrease the blood supply. Semanza & Domanovic (2013) studied the effect on the blood supply of increased infectious disease prevalence due to global warming. Others looked at the effect of pandemic influenza (Kamp, Heiden, Henseler, & Seitz, 2010; Zou, 2006) and dengue fever (Arellanos-Soto et al., 2015; Teo, Ng, & Lam, 2009) specifically. The above works focus on describing the effect of a disaster and not on coping with a disaster when it occurs. One exception to this was McQuilten et al. (2014) who looked at categorizing the urgency of procedures requiring blood under a prolonged shortage and the effect of withholding units to non-urgent procedures.

In addition to vulnerability due to demand surges in isolated areas or supply decrease, recent consolidation of regional blood centers in the U.S. has led to a situation where the disruption of operations of a single blood center, via whatever mechanism, could have significant impact. Following 9/11, an interorganizational task force found that disruption of the blood supply chain, not lack of supply, posed the greatest risk following a disaster or terrorist attack (AABB, 2008). For regions already experiencing chronic shortage, without a robust blood supply network, no special circumstances are necessary to make this work relevant.

4.2.2. Allocation Problems.

The issuing, or allocation, problem for blood focuses on the allocation of blood to hospitals (Sapountzis, 1989) or from a hospital to an individual patient (Nahmias & Pierskalla, 1973). The allocation of blood units to individuals is in fact the major focus of this chapter.

Blood research most often studies ordering (replenishment) policies and assumes a pre-specified issuing policy such as first-n-first-out (FIFO) or last-in-first-out (LIFO). When allocation policies are studied it is most often to compare several different FIFO or

LIFO allocation policies to each other or to compare the performance of one policy under multiple different circumstances. Pierskalla & Roach (1972) studied FIFO issuing polices for blood products. Goh, Greenberg & Matusuo (1993) compared two FIFO allocation policies for perishable goods. Parlar, Perry & Stadje (2010) considered FIFO and LIFO allocation policies for perishable inventory. These efforts focused on issuing policies but did not include any substitution aspect.

There are few allocation policies that take substitutability into account. Efforts to find allocation policies which include substitution most often allow substitution between demand classes, usually age-based, and not between the blood groups. Nahmias & Pierskalla (1973) studied LIFO and FIFO issuing polices under a problem where demand existed for units of a specific age. Units of younger age were allowed to be given to unfilled demand for units of older age. More recently, Deniz, Karaesmen & Scheller-Wolf (2010) found that use of the standard replenishment policy may produce pathological results when used with an allocation policy that allows age-based substitution. One recent study that did focus on an allocation policy and allowed for substitution was by Civelek, Karaesmen & Scheller-Wolf (2015). Here an item of specific age was demanded and substitution with products of a different age was allowed. In this case substitution was a recourse action, meaning that it was undertaken only when units of desired age were not available.

Substitution is not a recourse action in the policy formulated in this chapter; substitution is an active choice made to improve overall utilization of limited supply. Substitution among compatible blood types is mentioned as a concept in a few papers but it is not allowed by the issuing policy used (Kopach, 2004; Kopach, Balcıoğlu, & Carter, 2008) or no attempt was made to deliberately substitute or otherwise utilize substitution as a tool (Katsaliaki & Brailsford, 2006). Administering compatible but "mismatched" units that are compatible is medically sound and occurs frequently in practice (Yazdanbakhsh & Nandi, 2016).

4.2.3. Other related research

Netessine & Rudi (2003) derived a stocking policy for non-perishable goods under consumer driven substitution. Ishii & Nose (1996) studied replenishment policies under substitution between priority classes of demand, assuming a FIFO issuing policy. Sainathan (2013) studied pricing and order quantity for perishable good while allowing substitution between two age classed demand groups, but did not give an issuing policy. Kok & Fisher (2007) and Mahajan & Van Ryszin (2001) studied substitution in assortment planning.

4.3. Exact MDP Model Formulation

A Markov Decision Process (MDP) model is developed in this section for allocating current inventory blood units to patients. The MDP determines the optimal type to administer to a patient if the patient's own type is not available. In this model, patients arrive one by one and are served a single unit of compatible blood. Some previous studies have shown this assumption of dispensing only single units does not significantly degrade solution quality or generalizability (Sarhangian, Abouee-Mehrizi, Baron, & Berman, 2017). The objective function in the model maximizes discounted total reward. Using a discount rate despite the short time horizon encourages the allocation of units as soon as possible in an attempt to delay as long as possible the arrival of the first unserved patient while maintaining a 100% service level. This allows maximum time for replenishment since the restocking time may be unpredictable in an emergency.

<u>States of MDP</u>

Each $(I, \mathfrak{B})$ state represents the combination of inventory position, I, and patient blood type, $\mathfrak{B}$. Using the most common blood type classification system (ABO/Rh) there would be eight common blood types (A+/-, B+/-, AB+/-, O+/-), when considering Rh-positive and Rh-negative distinctions. Due to problem size, this model ignores the Rh factor and focuses only on matching ABO types, that is:

$$\mathfrak{B} = \{A, B, AB, O\}$$

For Rh-negative patients the four-type model allows all appropriate substitutions but for Rh-positive patients this assumption limits the possible substitutions since Rh-positive patients can accept corresponding Rh-negative blood type. Later it is shown how these results can be applied for all Rh types. Some studies have shown that Rh positive/negative can be ignored in trauma response without significant adverse effect (Flommersfeld et al., 2018)

The MDP formulation creates very large state spaces very quickly. The size of the state space depends on the starting inventory level of each type. Let $M_b$be the starting inventory of blood type b. Then the number of states is:

$$|S| = |\mathfrak{B}| \times \prod_{b \in \mathfrak{B}} (M_b + 1) \tag{4.1}$$

For example, given a starting inventory of 20 units of each blood type there are over

700,000 states. Large hospital inventories can include hundreds of units of each type. Inclusion of patient blood type as part of the state definition prevents the existence of absorbing states in the MDP, but the zero-inventory states are an absorbing set.

Decisions in MDP.

When a patient arrives requiring a unit of compatible blood, the decision may be made to provide any of the available compatible units. If no unit is provided the patient becomes "unserved". The actions available in state s, $a_{(s)}$ are:

$$a_{(s)} = \{0\} \cup \{\, b: b \in \mathfrak{B}, I_b > 0\} \tag{4.2}$$

where I_b is the current inventory of type b available and 0 is the do-nothing action.

Transition probabilities in the MDP Model.

This model employs two categories of transition probabilities: transitions between inventory positions and transitions between patients' blood types. Transition between inventory positions is deterministic based on the action chosen. In this model it is assumed each patient requires a single unit of blood. If the do-nothing action is chosen, the inventory position remains the same; otherwise inventory of the provided blood type decreases by 1.

The blood type of an arriving patient is random, and the blood type of the next patient is independent of the blood type of the current patient. The probability of a new patient having a specific blood type is constant throughout and is based on the blood type distribution of the population under study.

Rewards in the MDP Model.

Allowable blood type substitutions are shown in Figure 7. Serving a patient with a compatible unit results in a positive reward. This model assumes equal rewards. Not serving gives zero reward and serving an incompatible unit results in a negative equal reward. Rewards are denoted r(s, a).

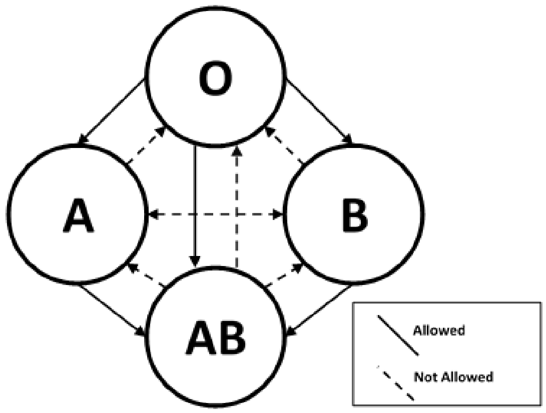

Figure 7. ABO Allowable Blood Type Substitution

<u>MDP Objective Function.</u>

The objective is to maximize the expected value of current and future rewards over all available actions. This value is computed via stochastic dynamic programming:

$$v^n(s) = \max_{a \in A}\{r(s,a) + \lambda \sum_{s' \in S} p\big(s' \big| (S,a)\big) v^{n-1}(s')\} \tag{4.3}$$

This function includes a discount factor, λ, which makes providing a unit for the current patient now more important than providing a unit for a later patient.

4.4. Computational Results and Analysis

The initial solution to this problem was found using the Gauss-Seidel variation of the Value Iteration algorithm, programmed in Python. See Powell (W. Powell, 2011) for details on this approach.

4.4.1. Results: Optimal policy for U.S. blood type distribution

The initial model sets transition probabilities according to the mean U.S. blood type distribution; see Table 5. At each transition the probability of moving to a state where the new patient blood type is A is 40% and the probability of moving to a state where the new patient blood type is B is 11% and so on.

Table 5

US Population ABO Type Distribution

A	B	AB	O
40%	11%	4%	45%

The initial I included all potential inventory positions from zero to 20 units of each type on hand. The resultant policy recommends the unit type to be administered to the patient in each $(I, \mathfrak{B})$ state

As expected, the policy is to give each patient type its own blood type until it is exhausted, then a compatible alternative if one exits. The only option for O-type patients is to receive O-type blood; for A and B patients O is the only option once their respective inventories are depleted. In this model, the only non-obvious substitution choice is what to give to AB patients once AB is depleted. The answer depends on the inventory of A, B and O on hand. Figure 8 and Figure 9 show the recommended actions when the current patient is type AB and there are no AB units available.

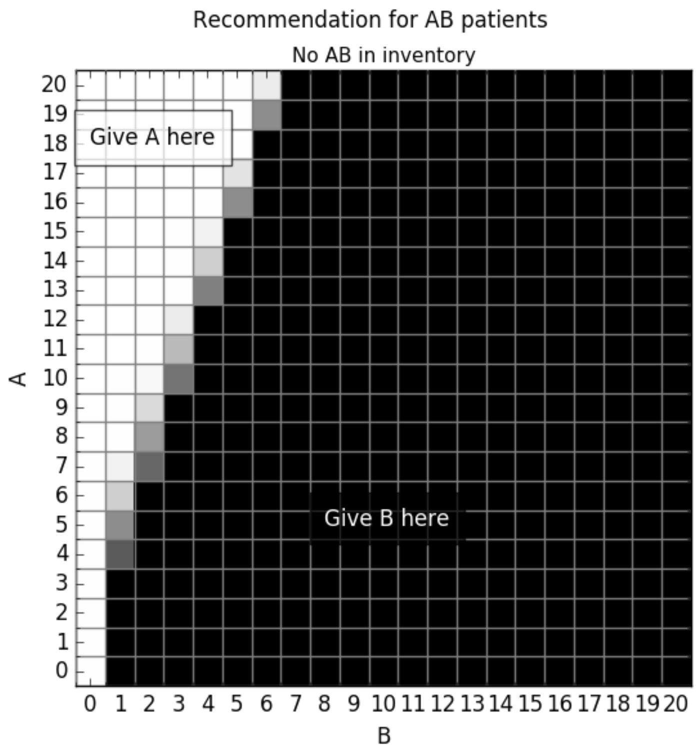

Figure 8. Recommendations for AB Patients When There is No AB Inventory (Projection)*.

**Note: Quantity of A inventory is on the vertical axis and B inventory quantity is on the horizontal axis. This plot is for all inventory levels of O. B is recommended in the black areas independent of O inventory. In the white areas A is recommended independent of O inventory. Gray areas indicate that the recommendation depends on O inventory. Darker gray areas indicate areas where B is given at lower levels of O inventory; lighter gray indicates where A is given at lower levels of O inventory. Levels of grey shades are given only for visualization; exact percentages are given in the results of the value iteration algorithm.*

Expounding on when to give A units and when to give B units in the gray areas in Figure 8 above, optimal decisions from the value iteration algorithm can also be represented in the 3-dimensional representation shown in Figure 9. Though relative proportions of A and B are important, in general B is recommended more frequently as O inventory increases. Note that the boundary between the “allocate A” region and the “allocate B” region is not linear, but close to linear. Later (in Figures 17–19) this boundary is approximated as a linear plane.

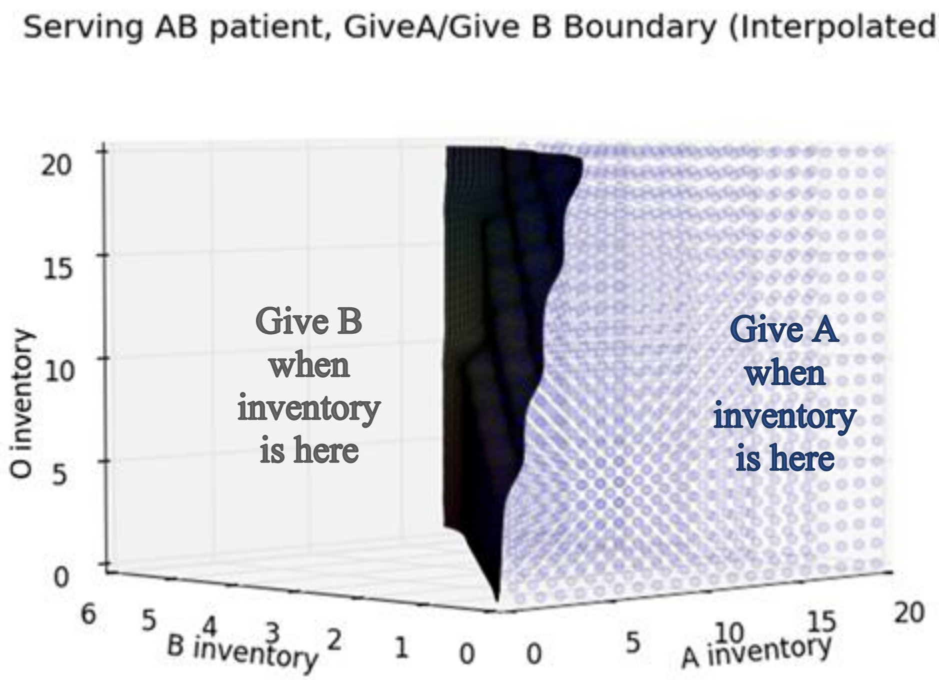

Figure 9. Serving AB Patient, Give A/Give B Boundary*
** Note: Dots indicate inventory positions where A is recommended to AB patients when there is no AB blood in inventory. The other side of the dark plane is where B is given to AB patient.*

After the MDP was solved, a simple Python-based simulation was created to compare the “optimal” policy to a “standard” or “default” policy. The default policy is a fixed priority policy. In the “default” policy each patient receives a unit of its own type unless that type is out of stock. For AB patients, once AB inventory is exhausted, A is allocated until it is exhausted, then B is used.

Given a specific starting inventory, a series of patients were generated and assigned blood types according to the U.S. average blood type distribution. In the control group, the simulated patients were allocated blood units via the default policy. For the comparison group, the same stream of patients was allocated blood via the optimal policy generated by the MDP model.

5,000 simulation runs were performed with initial inventory of 20 units of each type except AB. Initial inventory for AB was zero. The patient-by-patient unit allocations were very similar for the standard and optimal methods. This is as expected because AB is a small percentage of the US population and standard and optimal AB policies differ only slightly.

The patient number of the first unserved patient was found in each simulation. Under the standard allocation policy an average of 42.87 patients were served before the first unserved patient while under the optimal AB policy an average of 43.78 were served. A paired t-test shows this difference (0.91 patients) is statistically significant ($p<0.001$). Bootstrapped mean confidence intervals are shown in Figure 10.

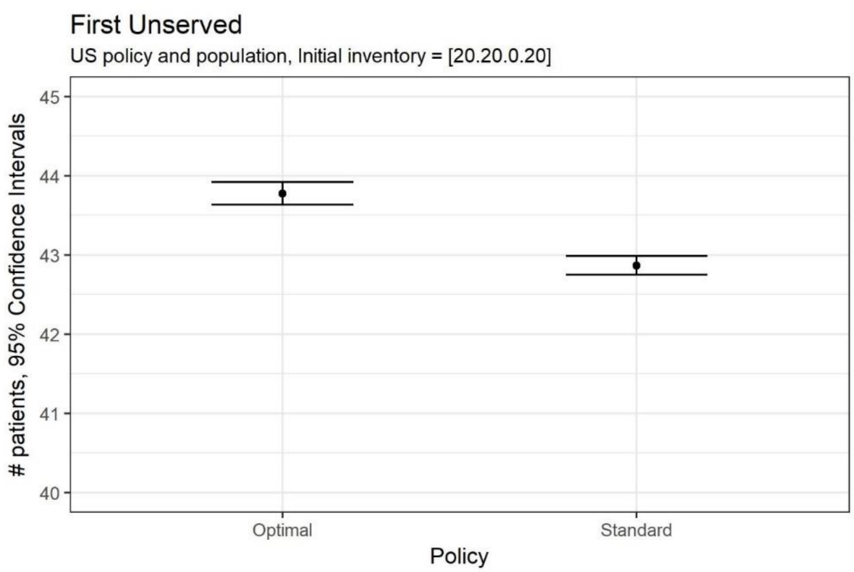

Figure 10. First Unserved Patient: US policy and population. Patient number of the first unserved patient, optimal vs default AB policy, based on U.S. population average for initial inventory [20A, 20B,.0AB,20O]

The first unserved patient metric addresses the performance of a policy under a 100% service level objective. For comparison purposes the patient number of the third unserved patient was found. The average patient number of the third unserved patient was 45.96 under the default policy and 47.12 under the optimal policy ($p<0.001$), a difference of 1.17. Bootstrapped confidence intervals on the mean are shown in Figure 11.

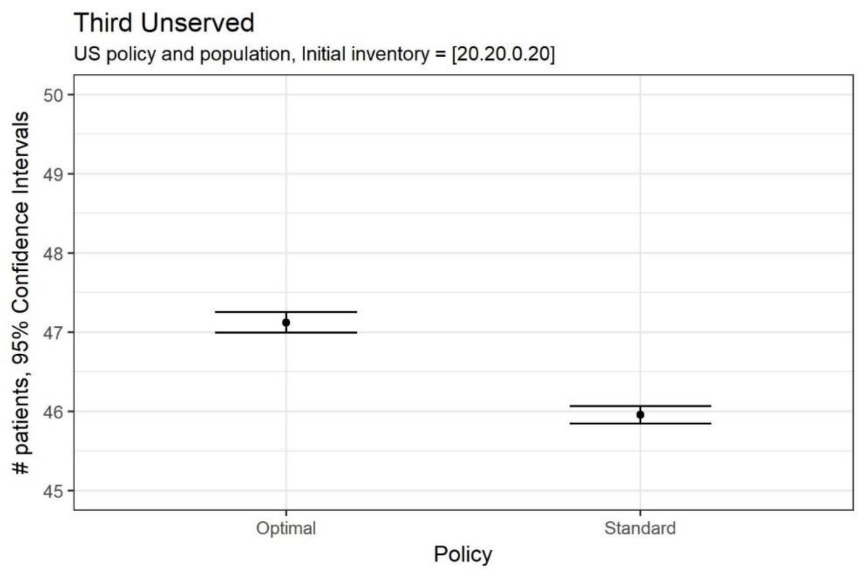

Figure 11. Third Unserved Patient. Patient number of the third unserved patient: optimal vs standard AB policy, based on U.S. population average for initial inventory [20 A, 20 B, 0 AB, 20 O]

Another metric of interest is the utilization of inventory under a specific policy. In these simulations 60 units of inventory in total were available for the 80 patients that were generated. The percentage of the first 60 patients served is a measure of utilization. If units were perfectly matched each of the first 60 patients would be served and utilization of the 60 units of inventory would be 100%. Perfect allocation may not be possible. To find the theoretical minimum number of unserved patients in the first 60, given a specific inventory, patient types and unit types were assumed known a priori, before any patient is served, and a maximum bipartite matching between patients and unit types was made. Patients not part of a matching would be unserved. Across all simulations the average theoretical minimum number of unserved patients was 11.26. On average 13.66 f the first 60 patients were unserved using the default policy and 12.00 patients were unserved under the optimal policy (Figure 12). Thus, the optimal policy does utilize on-hand inventory more efficiently than the standard policy.

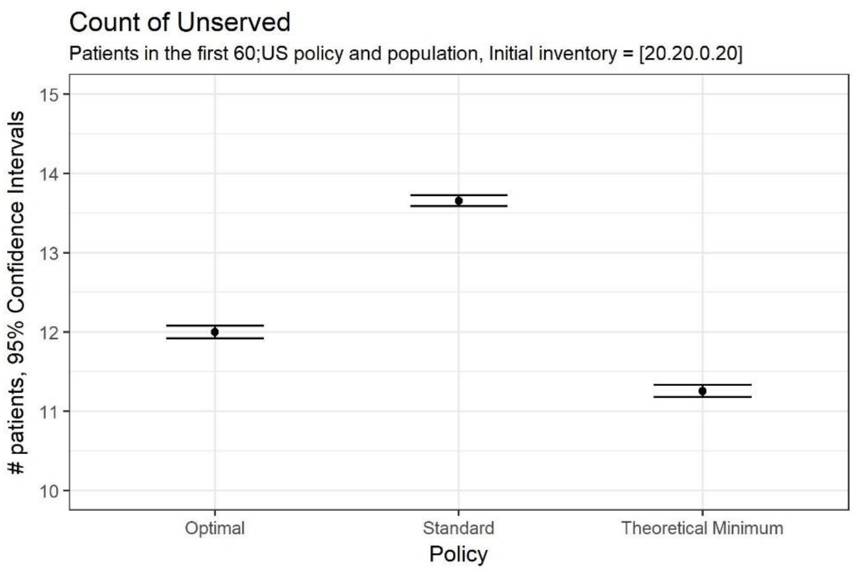

Figure 12. Count of Unserved Patients in First 60 Patients; Comparisons with Theoretical Minimums

4.4.2. Results: Effect of population distribution

In the previous simulations the MDP transition probabilities matched the patient population of the simulations. Blood type distributions change with demographics and therefore vary by region. To investigate the sensitivity of the policy to the underlying recipient population, the MDP was solved using three additional sets of population data to estimate transition probabilities (Figure 13.). Each population-specific MDP policy was simulated on multiple patient populations. A significant difference in performance would be a strong motivation to generate optimal policies based on the population served by a given facility. It would also be an indicator of the sensitivity of the policies to misspecification of modeled transition probabilities. Such inaccuracies may arise from uncertainty about donor rates or local demographics or from different donation rates by blood type.

Chicago Illinois, U.S. was chosen because of the increased prevalence of type O blood in the population but overall similarity to the U.S. average. Hungary (Rex-Kiss, Szabó, Szabó, & Hartmann, 1973) and Lucknow, India (Chandra & Gupta, 2012) were chosen for comparison to the U.S. These two populations have larger percentages with AB blood. Hungary has a much lower O percentage and higher A percentage than the U.S. The India population has a much higher B percentage than the any of the other populations considered. The MDP was solved using each population's transition probabilities resulting in a population-specific policy. Each population-specific policy was then applied in simulation to each population.

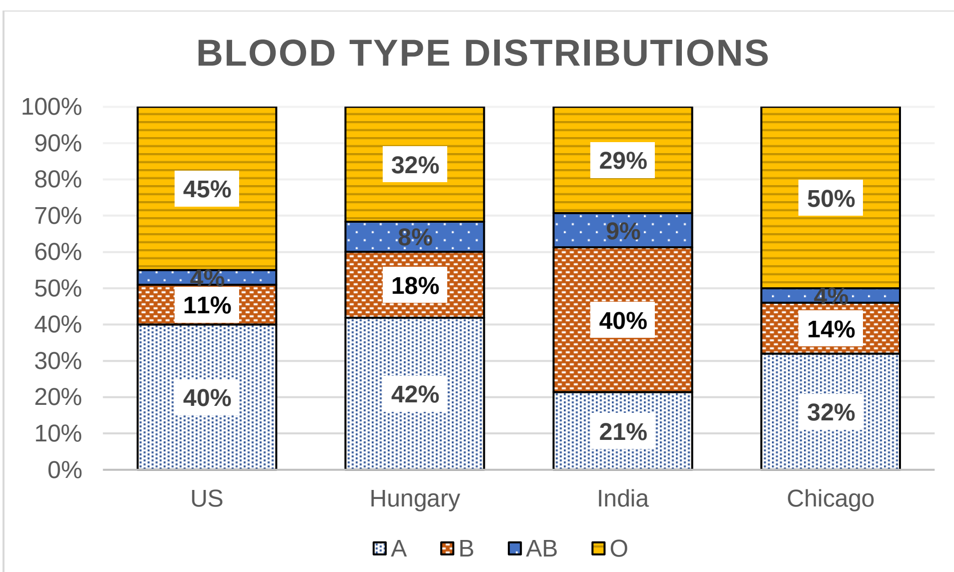

Figure 13. Blood Type Distributions: U.S., Hungary, India, Chicago

Figure 13 and Table 6 show the results of sixteen simulations. Each location-optimal policy was applied to each of the simulated populations. The patient stream was held constant across policies for each simulated population. Table 6 shows the average first unserved patient for each run. When the "Chicago optimal" and "U.S. Average optimal" policies were applied to the simulated Chicago population (or to the simulated U.S. Average population) the means are approximately equal. This is because the policies are only slightly different and the proportion of the population with AB is so small that the MDP states where the policies differ are rarely encountered.

Each of the optimal policies performed better than the standard policy when applied to the population for which the policy was created. In most cases, any of the optimal policies is better than the standard policy, even if it was generated based on transition probabilities representing different populations. However, this is not true in the case of the India population. Only the "India optimal" policy performed better than the standard policy when applied to an India population distribution; the other policies do not perform as well as the standard policy. This is because the India blood type distribution features a much larger proportion of B blood type than the other blood type distributions. It is interesting that application to the simulated India population results in significant differences between the other polices. When applied to the simulated India population the Hungary optimal policy performed better than the Chicago optimal policy which performed better than the U.S. optimal policy. This ranking correlates with the ranking of the locations by percentage of B blood in the population (i.e. Hungary has the highest percentage of B blood of the three).

Table 6.

First Unserved patient: Standard, U.S. Optimal, Hungary Optimal, India Optimal and Chicago Optimal Policies applied to U.S., Hungary, India and Chicago Average populations

		Transition Population			
		Hungary	India	US	Chicago
POLICY	Standard	50.09	56.29	42.73	40.47
	Hungary Optimal	53.70	53.84	43.63	40.76
	India Optimal	51.65	56.33	43.18	40.62
	US Optimal	53.70	53.31	43.63	40.76
	Chicago Optimal	53.70	53.66	43.63	40.76

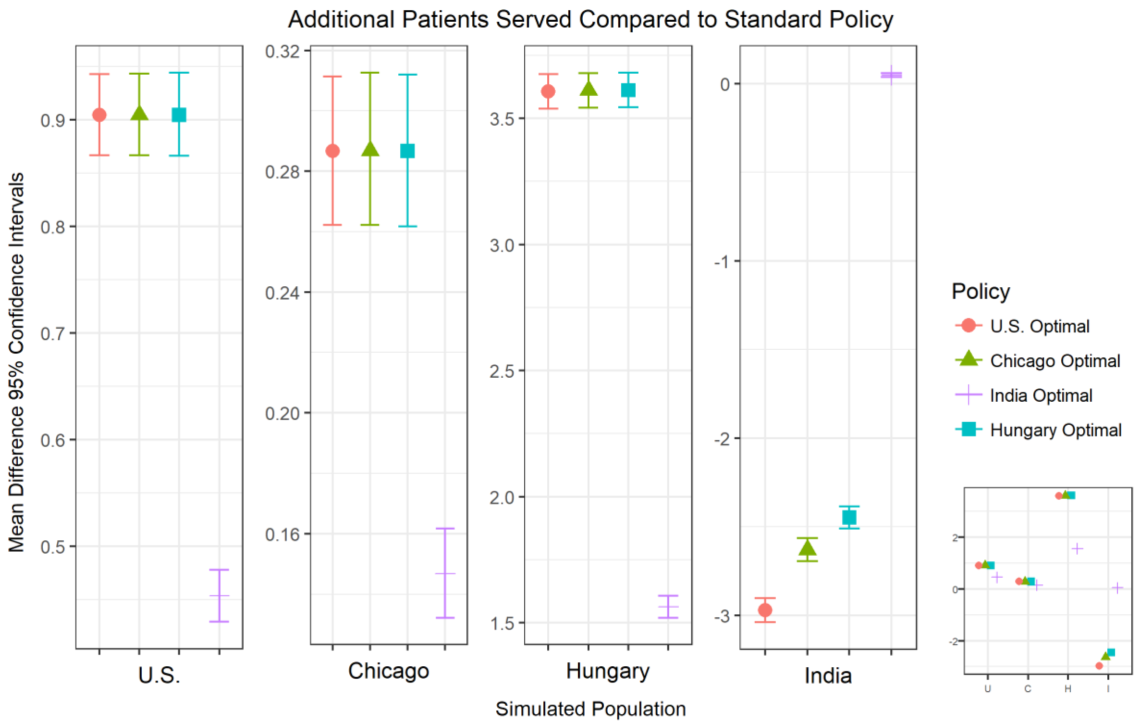

Figure 14. Pairwise Comparison of Location-specific Polices with Simulated Patient Populations. *Note: Vertical scales are not identical on the overall plot. (Inset shows all points on same scale, for reference.)*

Large differences in the population blood distribution do benefit from having tailored polices, but when the blood distributions of the populations are similar the benefit from custom policies is minimal. The question of "how different" is different enough to require a custom policy is unanswered.

4.4.3. Results: Effect of initial inventory.

All previously discussed simulations started with an initial inventory of [20 A, 20 B, 0 AB, 20 O]. Assuming zero AB inventory on hand is not unrealistic, as some hospitals do not order AB inventory. Nevertheless, it was worth further investigation to determine the effect of different starting inventories on the performance of the optimal policy and the interaction of starting inventories with the patient blood type distribution and policy population blood type distribution. In these simulations initial inventory was chosen so that the proportions of the types in initial inventory reflect the relative prevalence of the blood types in the population. The initial inventory for these simulations reflect the U.S. population average. There were 18 A; 5 B; 2 AB and 20 O units in initial inventory.

Simulations were run for multiples of this initial inventory position and for the same inventory levels with and without any initial AB inventory. Interval plots for each of these are shown in Figure 15 and Figure 16 below. The effect of the optimal policy is much greater in the case where no AB inventory is initially present (as expected) and the benefit increases linearly with inventory volume. When proportional amount of AB inventory is present the effect is decreased, and the benefit does not follow a linear growth pattern.

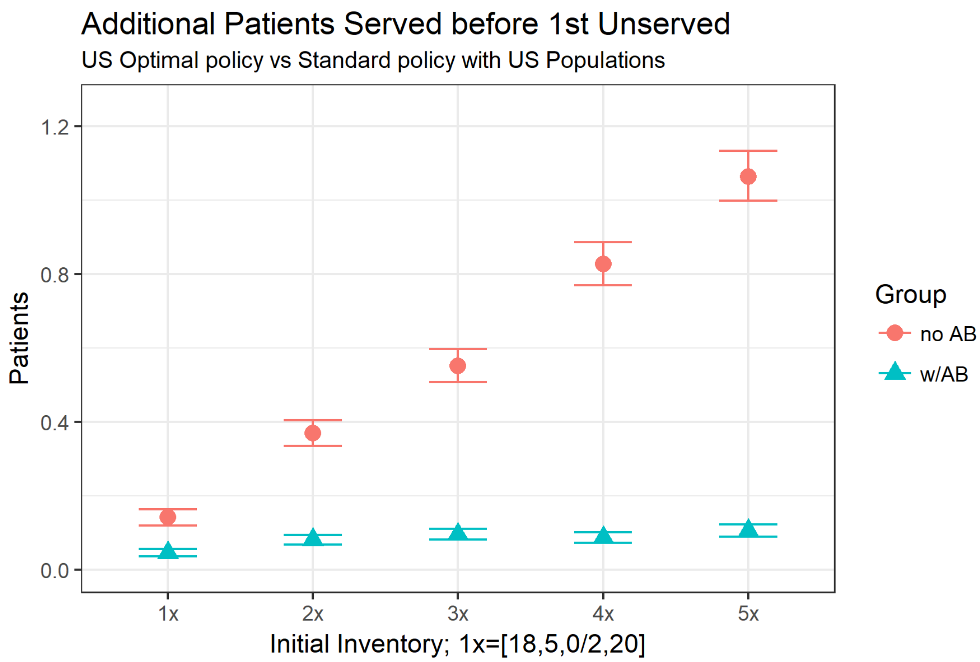

Figure 15. Effect of Inventory Level on Policy Improvement.

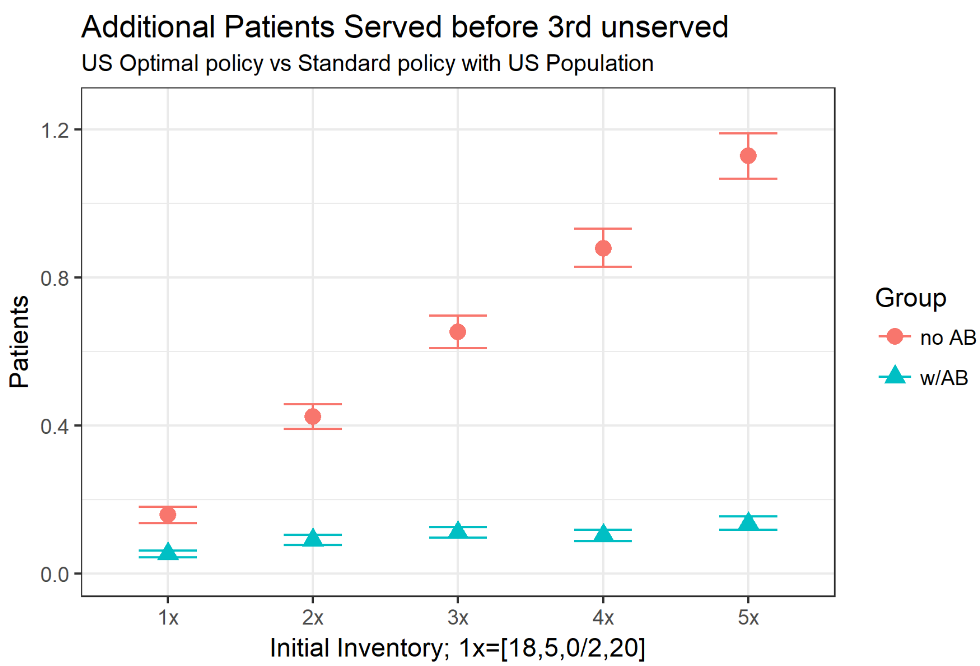

Figure 16. Effect of Inventory Level on Policy Improvement

4.4.4. Results: Effects of timing the application of optimal policy.

The huge state space necessary for even moderate inventories is the primary challenge of this model. This prompted consideration of the effect of the timing of the application of the optimal policy. If the policy is primarily beneficial at low inventory levels, perhaps determining optimal policies for larger inventory positions is unnecessary.

To investigate this idea additional simulations were performed, starting the optimal policy either (1) at the beginning of the simulation or (2) when there were no more than 10 units in inventory for any one type. Initial inventory was [20 A, 20 B, 0 AB, 20 O]. Results are shown in Table 7 and Figure 17. Though both timings are superior to always using the standard policy, the effect of using the optimal policy earlier is greater than applying the optimal policy only when the maximum inventory for any blood type dropped to ten or fewer units. The size of the effect varies dramatically by population. For example, more patients are served in either Hungarian scenario than in either U.S. average scenario. 3.5 additional patients were served before the first unserved patient through earlier application of the optimal policy for the Hungary case. For the U.S. average case an additional 0.9 patients were served.

Table 7

Additional patients served through earlier application of optimal policy, compared to using standard policy until the inventory of any type reaches 10 or fewer units and then allocating according to optimal policy. (Note: sd = standard deviation.)

Population specific Optimal Policy	Optimal Policy Always Applied				Optimal Policy Delayed				Additional Patient Served by not delaying.
	First Unserved		3rd Unserved		First Unserved		3rd Unserved		
	mean	sd	mean	sd	mean	sd	mean	sd	
Chicago	40.8	5.9	44.5	5.9	40.5	5.5	44.1	5.4	0.3
Hungary	53.7	4.4	56.6	4.1	50.2	4.0	53.0	4.0	3.5
India	56.3	4.0	59.1	3.5	56.3	4.0	59.1	3.5	0.0
US Average	43.6	5.1	47.0	4.7	42.7	4.4	45.9	4.0	0.9

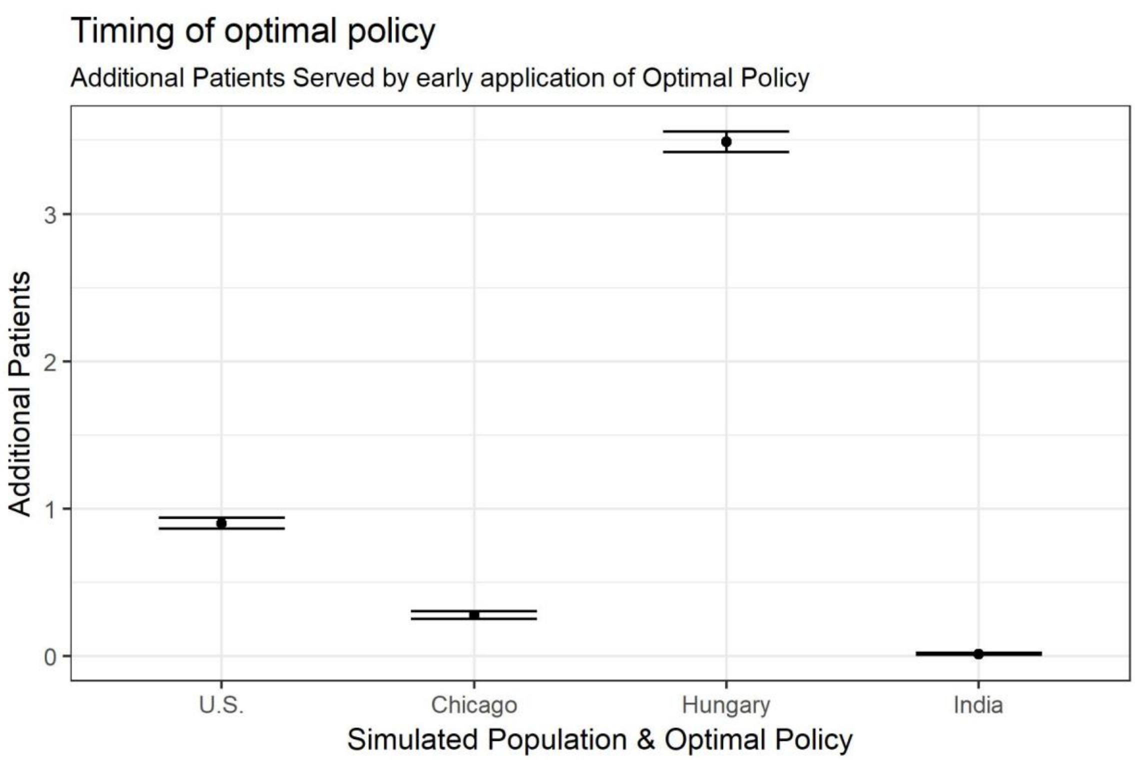

Figure 17. Timing of Optimal Policy. Additional patients served through earlier application of optimal policy, compared to allocating *using* standard policy until the inventory of any type reaches 10 or fewer, *then* allocating according to optimal policy

4.4.5. Results: Rh-type specific policies.

In the policies described above the transition probabilities were calculated as the percentages of each ABO patient type, ignoring Rh type. This may skew results as it is overstating the availability of O, the universal donor type in this Rh-agnostic model. 43% of the US population has blood type O (either O- or O+), however the true universal donor, O-, is found in only 8% of U.S. population.

To determine impact of this simplification and to explore options for generating an 8-type policy, the MDP model was next solved using transition probabilities based on only Rh-negative types. Because such a large percentage of the population has Rh-positive types, the overall ABO distribution and the ABO/Rh-positive distribution look very similar. However, Rh-negative distribution skews more toward A and less towards O. Blood type percentages by ABO type for US and Chicago overall average and Rh-negative subset only are shown in Table 8.

To visualize the differences between the policies an approximate separating hyperplane was fit to the set of boundary points between the inventory positions where A is recommended and the inventory positions where B is recommended when there is no AB inventory on hand and the patient is AB. These planes can be seen in Figure 18. The population that the policy is designed for impacts the plane more than the Rh factor. In both locations the Rh-positive policy is “steeper” than the Rh-agnostic policy, meaning the effect of O inventory levels is more pronounced.

Table 8

US and Chicago type distribution, all and negative

Population	A	B	AB	O
Chicago Rh-Negative	34%	12%	5%	49%
Chicago All	32%	14%	4%	50%
US Rh-Negative	42%	11%	5%	42%
US All	40%	11%	4%	45%

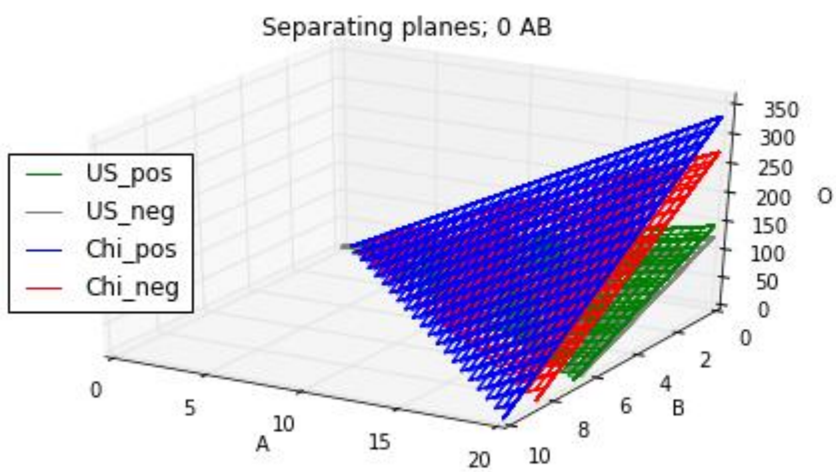

Figure 18. Approximate Separating Planes Between Inventory Positions where A is Recommended (below the planes) and Areas Where B is Allocated (above the planes), for AB Patients When There is No AB Inventory On Hand. From top to bottom are the planes derived from Chicago Rh positive optimal policy, Chicago Rh negative optimal policy, US Rh positive optimal policy, US Rh negative optimal policy.

The Chicago models also requires more B inventory on hand before B is recommended (areas above the planes). A visualization of the intersection of this plane with the Z = 0 plane is shown in Figure 19. Figure 20 shows the separating hyperplanes for the U.S., India and Hungarian average populations.

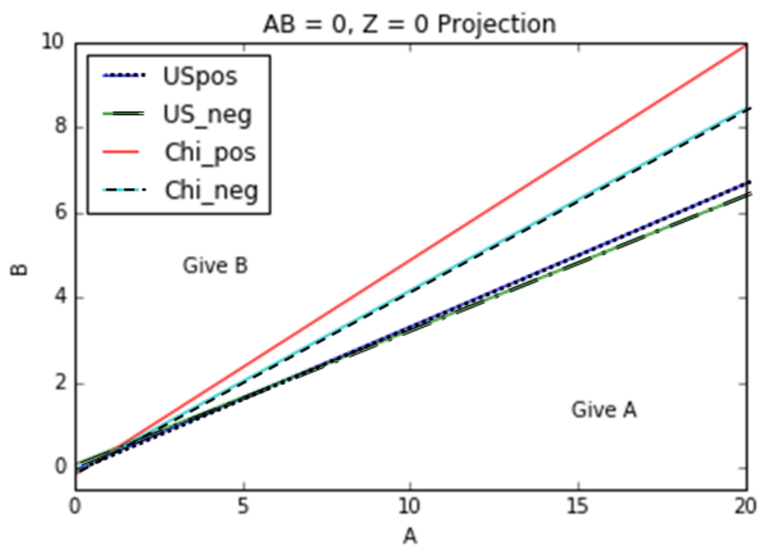

Figure 19. Policy Recommendation Boundary Planes Intersection with Zero O Inventory Plane

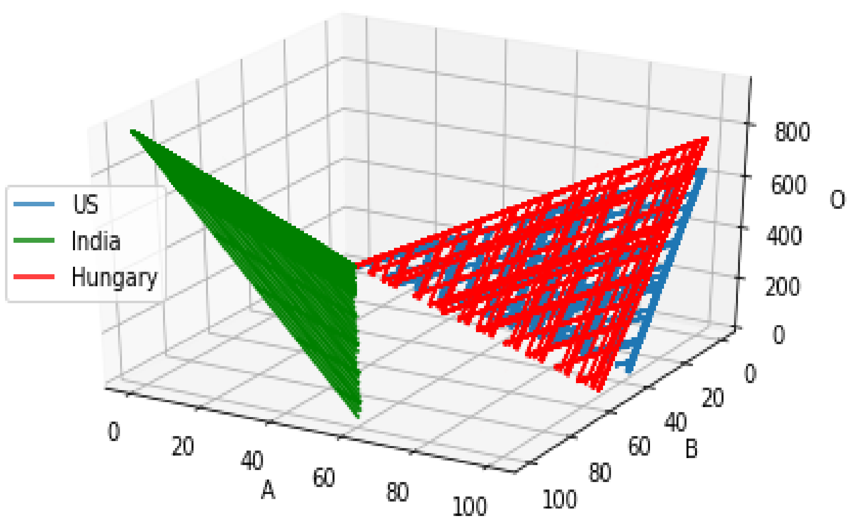

Figure 20. Extended Projections of planes separating "Give A" and "Give B" areas for U.S., India and Hungary averages. A is given on the side of the plane where quantity of A >0.

4.5. Detailed Simulation Case Studies

The preceding sections have analyzed the effectiveness of the MDP-derived policies using a simplified simulation and elucidated the necessity and benefits of tailoring the policies to the specific populations being served. This section presents two discrete event simulation (DES) case studies. In the first case study a single hospital faces a mass casualty incident which causes a demand spike. The second case study simulates the effect of a pandemic influenza outbreak, causing a moderately sustained reduction in the availability of new supply.

The numerical analysis of the previous sections was developed to evaluate a specific situation wherein inventory replenishment was prevented for the duration of an emergency event. The duration of the event was not specified, and it was assumed that the time frame was short enough that outdating would not occur. The objective was to determine which of the considered polices would stave off unserved patients for the longest time.

In this section the allocation policies are compared via simulation created in Arena (Rockwell Automation, 2016). See Chapter 2 for more detail on the simulation model. Unlike the previous analysis, this detailed simulation model includes blood units aging and outdating, replacement units arriving, and multiple units being demanded by some patients. It models the operations of a single hospital receiving supply units of varying age and using them to satisfy the demands of a stream of random patients. Now each patient demands one or more units over one or more days. There is no crossmatch reservation. There are no neonate or pediatric patients requiring partial units. Each unit demanded is considered separately but all demand from a single patient in one day is addressed concurrently. The patient blood type distribution is based on the U.S. average. Eight blood types are included.

The optimal policies previously discussed were based on only four blood types and solved over a limited range of inventory. To confront the inventory amount limitation separating hyperplanes were used to approximate the exact MDP recommendation. The planes allow the policy to be approximated at any inventory. Two optimal policies from the four-blood type approximation, the "U.S. negative" and the "U.S. positive" policies were used to address the eight blood types in the DES model. The U.S. negative approximation was used when satisfying demand by Rh-negative patients. Demand of Rh-positive patients was addressed by first attempting to satisfy demands with the U.S. positive policy approximation and if no compatible positive unit is available then the US negative policy approximation is used to select units. For the sake of brevity, this approach will be referred to as the "P1" policy. This is not an optimal overall policy but allows examination of the potential benefits of the approach.

The MDP-derived optimal policies were developed to address emergency situations. The DES model allows emergency events to be created. In these simulations there are two emergencies per year to allow sufficient recovery time between events. The base demand and the timing of emergency events are held constant across each scenario and replication. Two separate detailed scenarios were conducted where the type of emergency event was either (1) Pandemic Flu, or (2) Mass Casualty Incident (MCI). The event types are modeled separately and are discussed separately in the following sections.

4.5.1. Scenario 1: Mass Casualty Incident (MCI)

Mass casualty incidents are becoming increasingly common. The World Health Organization defines a mass casualty incident as an event “which generates more patients at one time than locally available resources can manage using routine procedures” (World Health Organization, 2007) and notes that context is important. In large metropolitan areas only very large events (e.g. hurricanes, nuclear events) are disruptive enough to constitute a mass casualty incident whereas in smaller or more isolated communities the crash of a single bus may sufficiently tax local resources as to constitute a mass casualty incident.

The hospital simulated in the MCI case study uses 96 units of RBC per day on average. This places the simulated hospital in the high demand category (Bloch, Cohn, Bruhn, Hirschler, & Nguyen, 2014). Like most U.S. hospitals supply is provided by a third party responding to orders placed by the hospital. Orders are placed three times per week and arrive 24 hours after the order is placed. When hospital inventory drops below 70%, urgent orders are placed which arrive 23 hours later. Total inventory is 7 days of supply as recommended by the AABB Disaster Operations Handbook (AABB, 2008) .

In the MCI scenario simulated, the emergency event triggers the creation of a set number of emergency patients. These emergency patients arrive at the hospital and request blood over the 30 minutes to 1.5 hours following the start of the emergency event.

Six scenarios were simulated (See Table 9).

- Scenario S-1: 0 emergency patients (Baseline)
- Scenario S-2: 48 emergency patients, demanding 1 unit each
- Scenario S-3: 96 emergency patients, demanding 1 unit each
- Scenario S-4: 192 emergency patients, demanding 1 unit each
- Scenario S-5: 238 emergency patients, demanding 1 unit each
- Scenario S-6: 96 emergency patients, each demanding multiple units, with number of units distributed according to the Triangular distribution with mean of 3 and range of 1 to 5.

In S-6, the expected number of units required per hospital admission is three based on expectations during a disaster (AABB Disaster Operations Handbook , 2008) and the reported average units required per gunshot victim at a trauma center (Demario et al., 2018).

Each scenario was simulated using the default allocation policy and P1 for a total of 12 runs. Each run lasted for 20 simulated years with 2 emergency events per year resulting in 40 data points per simulation run.

4.5.1.1. MCI Case Results

The results of the MCI simulations are shown in Table 9.

Table 9

Results of Mass Casualty Incident DES simulations. (Table shows the unserved demand (unserved patients and unserved units) and outdated units by scenario. % improvement: % improvement of P1 policy over standard policy for each scenario. "Emer" = emergency)

Scenario	Emer. patient per event	Units per emer. patient	Allocation Policy	Unserved Patients		Unserved Units		Outdated Units	
				Patients / year	% improve	Units/ year	% improve	Units/ year	% improve
S-1 (Baseline)	0	NA	Standard *P1*	13.15 11.8	10%	52.4 49.8	5%	106.25 98.7	7%
S-2	48	1	Standard *P1*	12.4 12.65	-2%	52.95 49.65	6%	102.8 102.25	1%
S-3	96	1	Standard *P1*	11.9 11.35	5%	53.3 48.45	9%	106.45 97.6	8%
S-4	192	1	Standard *P1*	16.4 14.4	12%	67.3 55.25	18%	109.3 101	8%
S-5	288	1	Standard *P1*	25.05 28.4	-13%	82.55 84.6	-2%	106.05 100.9	5%
S-6	96	TRIANG [1,3,5]	Standard *P1*	18 15.7	13%	70.8 65.45	8%	101.65 96.2	5%

Policy P1 policy resulted in fewer unserved units in each scenario except for S-5. P1 also resulted in fewer outdating units over the simulated time course. S-6 is of note as the only scenario where each emergency patient demanded multiple units. On average each emergency patient in S-6 demanded 2.5 units for a total average emergency demand of 240 units per event. P1 performed very well compared to the standard policy for this scenario in particular. In contrast, P1 failed to outperform the standard policy in S-5, which featured a similar total number of emergency units demanded but where each patient required only one unit. This suggests that the assumption of a single unit per person in the MDP model did not significantly diminish performance in realistic situations.

4.5.1.2. MCI Case Discussion

Allocation policy P1 approximates an optimal allocation policy. It was not developed to cope with outdating, with possible demands of multiple units per person, or with resupply. Despite that, the simulation results show fewer unserved patients and outdated units during both the standard operations and during simulated Mass Casualty Incidents. The simulated hospital has very high demand, with high levels of on hand inventory, making it strong again sudden demand spikes. A policy like P1 would be expected to be more applicable and more beneficial in a low-demand scenario. The simulation results support the usefulness of this analytical approach and the resultant policies both during emergencies and during normal operations.

The choices made for replenishment policy, lead time, unit age distribution, etc. represent a single hospital configuration. P1 may perform differently on other configurations.

4.5.2. ***Scenario 2: Pandemic Influenza***

Pandemic influenza has been a common and deadly event throughout human history. The "Spanish flu" of 1918 killed 3-5% of the total world population. Pandemics have occurred repeatedly since that time and will almost assuredly happen again (Jester, Uyeki, & Jernigan, 2018). In fact, at the time of writing this dissertation, COVID-19 is a worldwide pandemic (Boseley, 2020). Pandemics are a concern for many emergency planning organizations including the World Health Organization (WHO), American Association of Blood Bankers (AABB) and the US Department of homeland Security (DHS). It is believed that the demand for blood may decrease during a pandemic due to cancellation of elective procedures; however chronic (e.g. hemolytic disorder) and emergent needs (e.g. trauma, maternal hemorrhage) will persist. It is suggested that reductions in demand of 10-50 % (AABB Interorganizational Task Force on Pandemic Influenza and the Blood Supply, 2009) are possible but predicting the impact of pandemic influenza on the demand for blood is “is extremely difficult to quantify in advance (World Health Organization, 2011).”

The unavailability of suitable donors and workforce absenteeism at the collection or processing sites represent a source of a serious impact on the blood supply during a pandemic. The Covid-19 pandemic has led to decrease in blood donations and “staggering” drop in blood supplies (American Red Cross, 2020; Bekiempis, 2020; Flavelle, 2020; Marks, 2020). The decrease in donations has been tempered by reduction in elective procedures but as of the time of this writing the effect of the ongoing pandemic in terms of blood availability remains to be seen (Flavelle, 2020; Lardieri, 2020).

Events such as biological or physical attacks on processing facilities can be expected to have similar effects and decrease supply rather than increase demand. In this pandemic influenza simulation case study, the simulation model assumes there is no regional blood center to fill orders. Supply comes directly from a population of donors which has the same blood type distribution as the population being served. This allows the simulated hospital to function as a proxy for a region experiencing a pandemic.

A notable feature of pandemics is that their impact is not constant over time or geography. A pandemic may last for years and within that time regions will experience "waves": periods of increased severity and impact. These waves last for weeks or months (Jester et al., 2018; Laughland, 2020).

In this simulation case, the emergency event triggers a reduction in the available supply by a specified percent. This reduction lasts for the length of the "wave" which was set to 28 days in this model. This duration would reflect a very local perspective or the height of severity in a community.

The standard policy and the revised policy were used to determine unit allocation under 4 different scenarios:

- Scenario F-1:0% supply reduction (baseline)
- Scenario F-2: 30 % supply reduction
- Scenario F-3: 50 % supply reduction
- Scenario F-4: 70% supply reduction

70% reduction in supply in scenario F-4 may seem high but during the autumn wave of the 1918 pandemic some communities reported incidence rates as high as 75% (Britten, 1932) while the incidence over the larger region was lower.

Each scenario was simulated using the default policy and policy P-1 for each scenario for a total of 8 runs.

4.5.2.1. Pandemic Case Results.

The results (Table 10) show reduced unmet demand and fewer outdating units using allocation policy P1 under each scenario. P-1 improves on the standard policy for unserved patients, unserved units and outdated units in all scenarios.

Table 10

Results of Pandemic Influenza DES simulations. Table shows average annual unmet demand (patients unserved and units unserved) and outdated units by scenario. Note: % improvement: P1 improvement over Standard policy for each scenario

Scenario	% Reduction in Supply	Policy	Unserved Patients		Unserved Units		Outdated Units	
			Patients/ year	% improvement	Units/ year	% improvement	Units/ year	% improvement
F-1	0 (no emergency)	Standard	2.2	4.7%	8.7	0.6%	2364.3	0.0%
		P-1	2.1		8.7		2364.0	
F-2	30	Standard	16.7	9.9%	65.0	7.6%	769.5	0.8%
		P-1	15.1		60.0		763.3	
F-3	50	Standard	152.1	2.2%	575.0	2.7%	218.8	7.4%
		P-1	148.7		559.5		202.6	
F-4	70	Standard	421.0	0.5%	1647.1	1.3%	131.8	10.7%
		P-1	419.1		1625.9		117.8	

As shown in Figure 21 the average number of patients with unmet demand increases throughout the duration of the pandemic and recovers quickly, but not immediately, after incoming supply quantities return to normal.

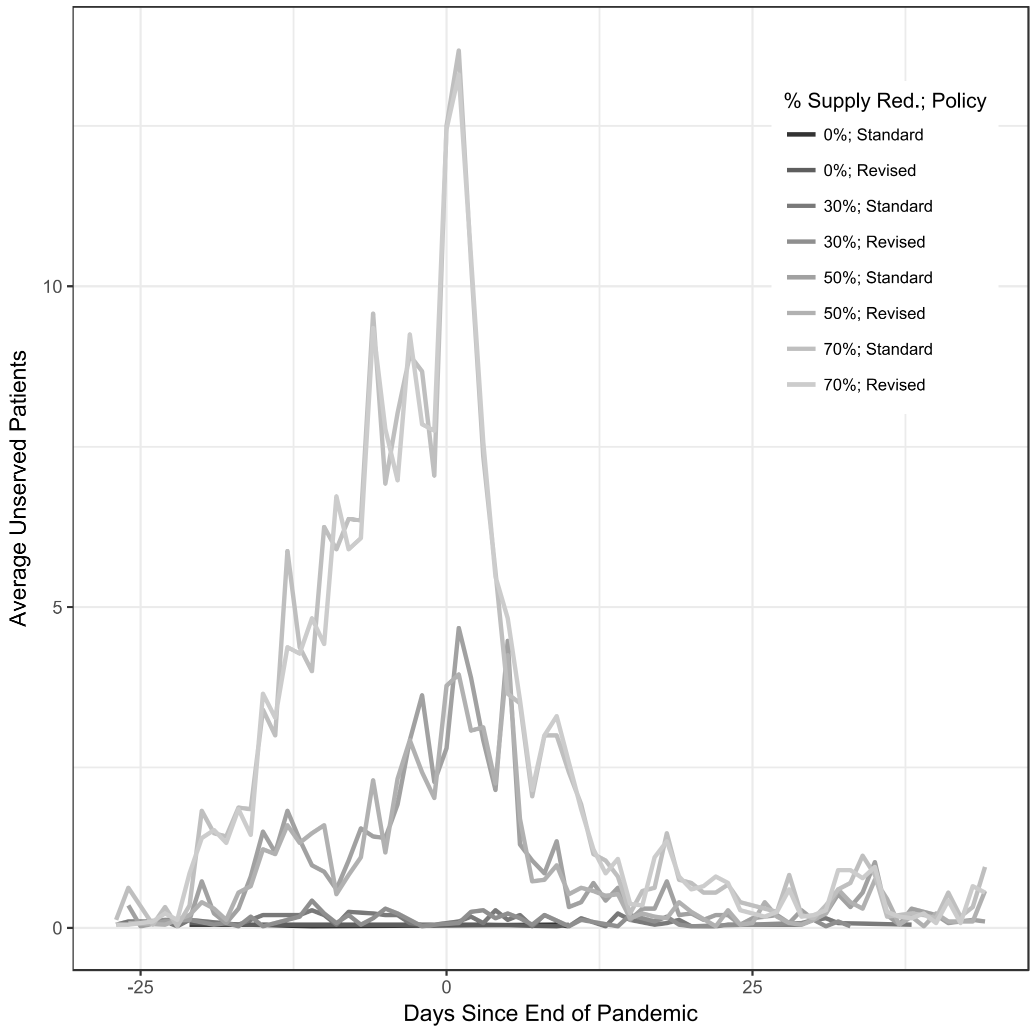

Figure 21. Unmet Demand Over The Course of a Pandemic. Shows the average number of patients with unserved demand over the duration of the pandemic and in the days following. Day 0 is the last day of the pandemic (last day of reduced supply).

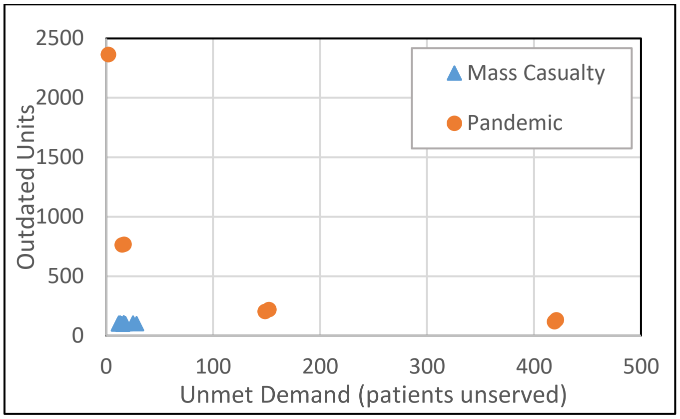

(a)

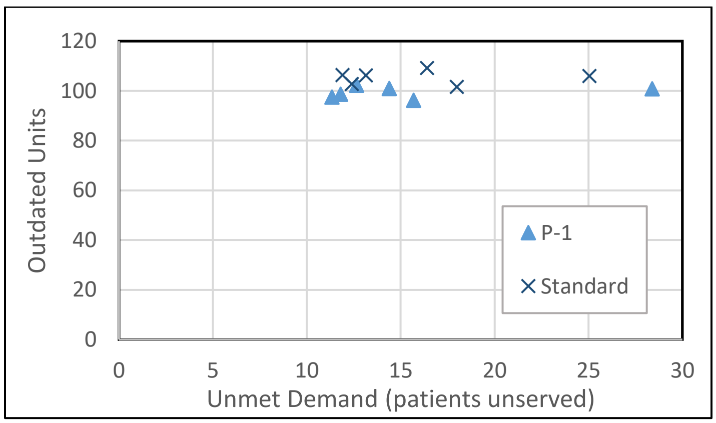

(b)

Figure 22: Unmet Demand and Outdated Units: (a) Unmet demand and outdated units for all scenarios; (b) closeup of unmet demand and outdated units, showing mass casualty scenarios only.

4.5.2.2. Pandemic Case Discussion.

The Pandemic Flu case study clearly illustrates the tradeoff between outdating unit and unmet demand, something that was not visible in the Mass Casualty case study (Figure 22). The Pandemic Flu results also illustrate the pattern of unmet demand over the duration of the influenza wave and the time required for operations to recover. In the simulation model, donor availability returned to 100% at the moment the emergency event ended but, as illustrated in Figure 21, recovery still took time.

4.5.3. *Discussion of the Simulation Case Studies*

The results of the two simulation studies support the usefulness of the individualized, optimal MDP-derived allocation policies. Though the optimal policies were not derived to address eight blood types, outdating units or arrival of new units they still provide benefit both during the emergencies and during normal operations when these factors are simulated. In addition to reducing the number of unserved patients the approximate policy, P1, reduced the number of outdating units compared to the standard policy.

The Pandemic simulation case study illustrates the clear tradeoff between outdating units and unmet supply and increasing unmet demand with increasing emergency severity (increasing reduction in supply availability). The MCI did not. This is likely due to the heterogeneously aged units arriving from the regional blood center and the ability to order additional units of specific type in the MCI study as opposed to the homogeneity of age and limitation on blood availability under the Pandemic Flu case study. It is a reminder of the importance of these variables in the study of perishable inventory.

Transshipment is often proposed as a solution to a shortage and it is useful when geography and time permit. However, it is not obvious how long transfers take to organize or that they will always be possible. Furthermore, while movement of units between hospitals is organized by blood centers direct transshipment is not the norm in the U.S. or the UK. Regional blood centers which organize the movement of blood in a region may be located several hours from the hospital in question under the best of condition. In moments of crisis the events that create additional demand (e.g., inclement weather leading to vehicular accidents and impassable roadways) may hamper new collection (Agnew, 2012) or resupply from other areas (Igarashi et al., 2018), as was the case following 2005 Hurricane Katrina in the United States (Stephan, 2008). When dealing with Red Blood units, time is of the essence. Half of RBC units given to trauma patients are administered in the first two hours (Stanworth et al., 2016) and trauma patient mortality has been shown to increase for every ten minutes that blood transfusion is delayed (E. K. Powell et al., 2016). More broadly, RBCs are demanded most often in the first 24 hours following a disaster and 62-74% of the first day's usage happens in the first 4 hours (Glasgow, Davenport, Perkins, Tai, & Brohi, 2013). Therefore, it is important to better utilize the supply on hand than rely on outside support. In the case of a pandemic or a mass casualty disaster with regional or national impact, resupply may simply not be available.

4.6.Conclusions and Potential Future Research

This research demonstrates that optimally using compatible blood units can increase the number of patients successfully served before the first unserved patient in emergencies. In such situations, this additional time may allow for more units to be delivered, preventing patients from going unserved and the attendant potential mortality. This chapter presents a method for finding such policies.

Comparison of populations from several regional locations showed that location optimal policies increase the number of patients served before the first unserved patient compared to policies derived for locations with very different populations with different blood type distributions (e.g. Hungary and India) but that location optimal policies for locations with similar populations (e.g., U.S. Average and Chicago) perform the same unless extreme events are encountered.

The results of this chapter also show that greater benefit occurs when optimal policies are applied continuously rather than only at low inventory levels, suggesting there is motivation to extend these policies to higher inventory states.

Though the improvements due to the optimal policies are small for the individual analyses presented here, the aggregate impact of implementing such policies in the thousands of hospitals across the U.S. would be much larger.

This research presents an innovative computational approach for extending the MDP-generated policy. By approximating the policies by using the separating linear hyperplanes developed for low inventory levels, the MDP solution can be extended to higher inventory levels, overcoming the large state space size requirement.

This is the first research to specifically address optimal allocation policies tailored to the demographics of the population being served. Inventory literature and disaster preparation focus on preparation at the regional level while the policy presented here is something that individual hospitals can use to support their own performance at no additional cost or taxing of outside resources. It should be noted that the few existing empirical studies of blood management during crisis usually apply to large medical centers in urban areas. It takes a much larger event to constitute a mass casualty incident in those situations, since large medical centers have the advantage of large standing inventories. This research is especially useful for small or isolated communalities.

4.6.1. Limitations and Potential Future Work

In the development of the optimal policies it was assumed that each patient demands a single unit. Though that simplification is supported by prior studies, the developed model could be adapted to include multiple units per patient. It should be noted however that the MCI case study did show that the policies developed with the single-unit assumption performed well in a simulation where each emergency patient demanded one or more units.

The policies discussed here were not developed for all eight blood types. The developed policies were compared with only one standard policy, and not every hospital uses the same allocation policy. The simulation case studies also describe only a single hospital configuration. The chosen “default” policy is typical of those used in American hospitals (Armstrong, Wilkinson, & Smart, 2008) but facilities in other locations may have other standard policies. Use of a single “standard” policy allowed for comparison among populations. More research is necessary to generalize the empirical findings from the simulations.

Modeling positive and negative Rh-factor groups separately was a necessary simplification for computational tractability. Including them in future research could result in an improved policy for eight blood types. Also, due to the scope of the emergencies considered, the MDP policies did not consider blood perishability. Although the DES case studies did include perishability and the MDP policies performed well, enhanced MDP policies that consider perishability may perform even better.

In the next chapter the model includes all eight blood types and includes perishability. However, the goal in that chapter is to improve performance under normal operating conditions, not only emergencies where objective is to most effectively use available blood unit inventory.

CHAPTER 5

APPROXIMATE MDP MODEL TO FIND OPTIMAL ALLOCATION UNDER NORMAL OPERATIONS

5.1.Introduction and Problem Definition

This chapter presents approximate dynamic programming (ADP) models that incorporate inventory age information to improve inventory performance under normal operations. In these models inventory ages, outdates, and is replenished. All eight ABO/Rh blood types are included.

The model of Chapter 4 was created to find improved policies for short-term emergency situations. With that objective in mind, it was not necessary to include temporal dynamics in the model. In this chapter the objective is specifically to improve performance during normal operations. The models developed here extend the active substitution idea of Chapter 4 while incorporating temporal dynamics. Approximate dynamic programming (ADP) solution approaches are used to address the explosion in problem size resulting from inclusion of inventory age information.

Several ADP models were tried. Rollout was successful only in the low demand case. Subsequently, approximate policy iteration (API) with Q-factor approximation was tried but it was also unsuccessful. Finally, API with policy improvement through rollout and approximation in policy space was used. In this chapter the process for finding improved policies is described for each ADP model. Performance is demonstrated for two hospital “scenarios” of vastly different daily demand, each based on real data. Results show statistically significant improvements, in comparison to the default policy, using rollout approach for the low demand (Cobre Valley) scenario and the API approach for both the low demand and high demand (Mayo Clinic) scenarios.

The following section provides a brief overview of literature related to blood inventory management with perishability and substitutability followed by background on the ADP methods used in this chapter. The chapter then describes the development of the rollout approach and its implementation followed by the development of the API policies. Presentation of results in this chapter begins with comparison of multiple alternate fixed priority policies to demonstrate that the default policy is a reasonable policy. None of the other fixed priority policies considered are shown to be significantly better than the default policy. Results for each of each API approach are presented for each scenario. Comparisons are made between the default and API generated policies under the base conditions and sensitivity analysis repeats the comparisons under altered demand and age of replenishment units distributions.

5.2. Blood Inventory Management with Perishability and Substitutability Background

A number of reviews of blood/perishable inventory management research have been published. The lack of allocation policy research is notable throughout.

In Prastacos' (1984) *Blood inventory management: An Overview of Theory and Practice* all issuing discussion is of FIFO or LIFO policies. Pierskalla (2004) presented a review of supply chain management of blood banks. In that thorough overview the discussion of optimal issuing policies only includes FIFO or LIFO policies. Beliën & Forcé (2012) offered an updated review of blood supply chain management research.

Only 11 of the nearly 100 papers cover issuing polices and many of those only consider FIFO and LIFO policies. Osario et al. (2015) presented a taxonomic framework for quantitative blood models and noted that, with few exceptions, the articles listed use FIFO as the issuing policy. The more recent Janssen et al. (2016) review of deteriorating inventory from 2012 – 2015 only listed 3 papers on blood allocation. Overall, very little research has been done on allocation polices for blood beyond the consideration of FIFO or LIFO policies. Furthermore, most allocation research does not consider the potential for substitution and when substitution is considered it is most often between age-based groups. This dissertation research is predicated on the belief that substitution can be actively employed to outperform simple FIFO and simple LIFO allocation polices.

Pierskalla & Roach (1972) did consider optimal issuing policies for a perishable blood product. In that study, stock was grouped according to age and substitution was permitted between age groups but there was no consideration of blood types.

Nahmias (1977) presented an ordering policy for a perishable good (but excluding substitution possibilities) assuming a FIFO "depletion" policy. The replenishment units outdated in the same order that they entered inventory. This assumption about the age of the replenishment units, or an even more restrictive assumption that all replenishment units are new, is common. The frequent assumption of fixed lifetime was noted by Nahmias (1982) in his perishable inventory review which also noted that issuing is virtually ignored.

Deniz et al. (2010) addressed joint replenishment and issuing policies for an unspecified perishable good with age-based demand classes. Substitution between age classes was a potential recourse action. They studied four issuance policies defined by the type of substitution permitted (no substitution, upward substitution, downward substitution and full substitution). These types of substitution are typical. The substitution permitted between blood types is closest to upward or downward substitution although it branches in a way that age-based substitution does not.

Haijema (2011) addressed the issuing policy for perishables with a short fixed life; specifically platelets. It was assumed that there were two events of issuing of batches per day. Substitution between age categories but not between blood types was allowed. All replenishment units were the same age upon arrival. The model in this chapter does not assume batch demand and does not restrict the age of the replenishment units other than requiring they be younger than the maximum shelf life.

Haijema (2013) found a new stock level-dependent ordering policy for a perishable product with a short shelf life using dynamic programming and simulation. Stock level-dependent replenishment policies are common, but stock level-dependent allocation policies are much less common. Both the policy in Chapter 4 and in this chapter are stock level-dependent. The policies developed in this chapter are also dependent on the age of the stock.

Haijema (2014) also describes stock age-dependent combination ordering, issuing and disposal policies found through stochastic dynamic programming. The MDP optimizes over the ordering and disposal policies; not over the issuing policy. That is, several issuing policies are considered in order to optimize the combination of policies but the issuing policies themselves may not be optimal themselves. The inventoried products in the model have a fixed maximum shelf life. ABO/Rh type compatibility type is not considered. Though the policy is age-dependent, substitutions take place in order to maximize age-dependent utility, rather than using the age information to reduce outdating or maintain service level. Similarly, Coelho & Laporte (2014) and Civelek et al. (2015) found replenishment policies for perishable goods in the presence of age-based demand preferences.

Three of the models most directly related to the dissertation research are in Duan and Liao (2014), Abbasi & Hosseinifard (2014) and Dillon et al. (2017).

Duan and Liao (2014) presented a model of red blood cell inventory control. They used a meta-heuristic to find optimal ordering policies for a two-echelon blood supply chain. The model included ABO/Rh substitution but only as a recourse action when the direct match was missing. The substitution pattern was fixed priority meaning that it was not age- or stock level-dependent. The model permitted replenishment units of varying ages but did not permit "leapfrogging" units. Leapfrogging is permitted in the model in this chapter. Duan also assumed batch issuing.

Abbasi & Hosseinifard presented a modified FIFO issuing policy for RBCs when replenishment is not controllable. The policy included age-based groups. Units were issued FIFO within groups and LIFO between groups. There was no consideration of blood type compatibility or blood type substitution. Uncontrolled replenishment would be the case when a blood bank only receives units directly from donors. A model with only uncontrolled replenishment could be used to approximate a region, as in the pandemic influenza case study of Chapter 4, but may not be representative of the majority of U.S.

hospitals; for example only 10.6% of (responding) U.S. hospitals reported collecting blood in 2011 (Whitaker, 2011).

Dillon et al. (2017) critiqued queuing-based perishable inventory models that make strong assumptions about the supply and demand processes. They also developed a two-stage stochastic programming model for RBC replenishment and allocation decisions. One of the tested allocation policies permitted recourse substitution between blood types with a fixed priority. The model assumed batch demand and that all replenishment units were same age and new.

The assumption that all units are new or the same age is inaccurate in reality. The consideration of accurate age information is critical for accurate modeling and its inclusion can lead to cost improvements (Broekmeulen & van Donselaar, 2009; Gürler & Özkaya, 2008). The models in this chapter use blood type-specific mixture distributions based on historical data to provide the age of replenishment units. See Appendix A for a discussion of this issue.

The lack of reported research in substitution between blood type groups is likely due, at least partially, to the inventory research focus on platelets rather than red blood cells. Platelets do not require blood type match to be distributed. On the other hand, substitution between age-based classes is reportedly done to accommodate the belief that fresh blood is preferred either always or for certain patient classes. The benefits of fresher blood have not been conclusively established and one of the largest prospective clinical trials undertaken to resolve this question showed that fresher units did not improve patient outcomes after cardiac interventions (Steiner et al., 2015). The model in this chapter does not implement age-based demand preferences, but it could be easily adjusted to accommodate such preferences.

Finally, several of the models described above included batch issuance. That is not realistic. The model developed in this chapter considers each unit of demand sequentially. It assumes a single unit of demand per person but could easily be adapted to allow multiple units per person by making transition probabilities in the MDP dependent on the current patient blood type.

5.3. Approximate Dynamic Programming Background

There are nearly as many approximate dynamic approaches as there are policies. Bertsekas lays out a useful scheme for describing the anatomy of a dynamic programming algorithm in *Dynamic Programming and Optimal Control* (2012). He describes many approaches as having a policy evaluation phase and a policy improvement phase. In the policy evaluation phase, the cost function of a policy is estimated. The policy improvement phase uses that information to attempt to find an improved policy. Each of these tasks can be addressed in a variety of ways. As with exact solution methods, approximate methods are generally described as either policy iteration (PI), value iteration (VI) or hybrid approaches, based whether they produce improved policies or improved value estimates at each iteration. In approximate methods, the distinction is made between approximation in policy space versus approximation in value space. Some methods contain both. Powell (W. Powell, 2011) similarly draws a distinction between the model for approximating a value function (lookup table, parametric models, and nonparametric models) and the process of learning the value function approximation, that is, a method for computing the value estimate of a fixed policy.

Solution methods abound but can be direct or iterative and model free or not. Altogether, the chosen approach should depend on the specifics of the problem; whether finite or infinite horizon, deterministic or stochastic, the sizes of the state and action spaces and whether they are discrete or continuous, the feasibility of calculating transition probabilities and the nature of the objective function.

The problem in this chapter is formulated as an infinite horizon problem with a huge state space and limited action space. Despite the limited number of actions, the transitions are complicated and transition probabilities cannot be reasonably calculated. These characteristics necessitate an approach that does not require enumerating and recording the state space and does not require transition probabilities to be known. For example, methods requiring matrix inversion are unfeasible. Consequently, a model-free approach that can be solved iteratively is needed.

All of the methods used in this chapter rely on simulation to make cost-to-go estimates. Simulation is frequently used for this purpose as it allows complex system dynamics to be captured and allows model-free analysis. ADP with simulation has been applied to traditional OR topics such as capacity allocation (Schütz & Kolisch, 2012), transshipment (Meissner & Senicheva, 2018) and surgical scheduling (Astaraky & Patrick, 2015).

5.3.1. Rollout Background.

Rollout is part of the lookahead family of polices. This method is a natural analog to the process of sequential decision making. Say that the system is in state S_0 (Figure 23) and a decision is needed about which action to take. For each action available at S_0, the states reachable at the next step (S_1) are found. From each of those states some base policy is applied over some horizon in order to generate a "cost-to-go" approximation. Combined with information about the rewards received over the $S_0 \rightarrow S_1$ transition, the value of being in state S_0 and making each decision (in other words, the Q-factors) can be estimated. The choice with the highest Q-factor is selected. The system progresses and the process is repeated at the next state.

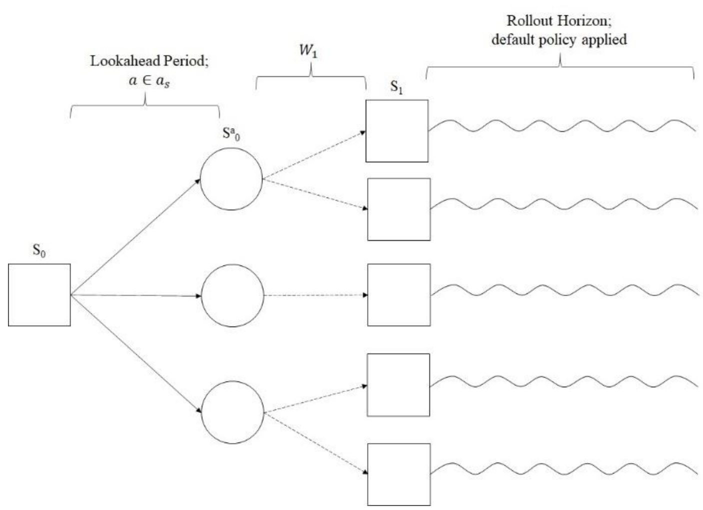

Figure 23. Diagram of Rollout Process

There are many variations of rollout through look-ahead trees:

- In multiperiod rollout the depth of the tree is increased.
- The tree may be truncated with or without addition of a terminal cost approximation.

- Sparse trees may be used by selecting branches of the tree instead of the entire tree.

Rollout is not guaranteed to give an optimal solution but it often works well in practice as famously demonstrated by Tesauro and Galperin's Backgammon AI (1996). Fortified rollout is a rollout variation guaranteed to do no worse than the base policy/heuristic (D. P. Bertsekas, 2019a).

Rollout is most useful when the problem size is small as the trees on which it is based can get impractically large, especially with multistep rollout. This curse of dimensionality has been addressed by finding ways to reduce the depth of the tree through truncation (Bhattacharya, Badyal, Wheeler, Gil, & Bertsekas, 2020) or to eliminate branches through the construction of a sparse trees (Somani, Ye, Hsu, & Lee, 2013). These are generally classed as Monte Carlo Tree Search (MCTS). Rollout is also integrated as part of more complex methods such as in Bertazzi (2015) and Moin (2018) which used a combination of MIP and rollout to address inventory routing. Adulyasak et al. (2015) used rollout to obtain feasible solutions as part of a multistage Bender's decomposition-based approach to the production routing problem. These alterations have allowed rollout's application to large problems.

In this chapter an online implementation of truncated single period rollout is used. In online rollout there is no lookup table or permanent representation of the policy. The system is implemented so that decisions are made in real time.

5.3.2. Approximate Policy Iteration with Q-factor Approximation

Background.

A single iteration of rollout will be optimal for finite-horizon deterministic problems. For more complex problems it may not be optimal, but rollout can be iterated to produce better solutions. This is policy iteration.

Policy iteration is the second of the two primary families of solution methods. See Bertsekas (2012) and Powell (W. Powell, 2011) for the essentials. Policy iteration can be viewed as a "perpetual rollout algorithm" (D. P. Bertsekas, 2019b). Rollout is used to generate a new policy which becomes the "base" heuristic for the next iteration. This approach requires a permanent representation of the rollout policy.

In general, an "approximation architecture" describes the structure of a model chosen to approximately represent either the state value (approximation in value space) or the policy itself (approximation in policy space). There are many different possible approximation architectures. Feature-based architectures are perhaps most common. For approximating costs they use the form:

$$\tilde{J}(x,r) = \hat{J}(\phi(x),r) \tag{5.1}$$

where x is the state for which to cost-to-go is to be approximated, ϕ is a feature vector, r is a vector of parameters and $\hat{J}$ is some function, denoting value. This architecture would include feature-based linear regression. The function ϕ transforms the state into a vector of features which are combined with the parameter vector r produce an approximation of the actual value function $\tilde{J}(x,r)$. These equations can be made stage dependent, for example:

$$,\tilde{J}(x_k, r_k) = \hat{J}(\phi_k(x_k), r_k) \tag{5.2}$$

In recent years neural networks and other machine learning techniques have gained popularity as approximation architectures. However, even as these methods may represent complex functions, many such methods lack interpretability.

As any data scientist knows, feature creation and selection can be a complex process as much art as science. It is a valuable process because well-chosen features can capture non-linear complexity of the value (policy) function, allowing for a simpler architecture. Handcrafted features permit the inclusion of expert opinion in these models while machine learning techniques afford a less hands-on approach to feature selection.

Once features are selected the model must be trained. Use of least-squares regression to fit linear regression models is very common.

The first policy iteration model used in this chapter approximates Q-factors. When approximating Q-factors it is not required to calculate expectations which would be very difficult in this case. Q-factors also simplify the process of deriving a policy from a value function. Now the recommended action is found by simply comparing the predicted Q-factors for the available actions (see eqn. 6.53 in Bertsekas, 2011):

$$\bar{\mu}(i) = \arg \min_{u \in U(i)} \bar{Q}_\mu(i, u, \bar{r}) \tag{5.3}$$

In this problem there are a limited number of actions to choose from but many states and many possible trajectories, making Q-factors an appealing option.

Initially, linear regression with a set of handcrafted features was used to approximate the Q-factors. A goal of this dissertation is to find practical, implementable improvements to real-world problems and linear regression offers the possibility of a model that is easily understood and implemented by blood bank operators. However, the linear regression models had low accuracy and the resulting algorithm did not perform well. Though a simple policy was the ideal, it was thought that it might be possible to find a good policy with a machine learning (ML) architecture and then further approximate it to a human-understandable form. In the worst case an ML model could be implemented in a way that does not require direct user interaction with the ML model.

Neural networks are a popular ML approximation architecture, having applications to many problems. Their use enables new functions to be approximated and enables new solution methods. However, it has not been used extensively in inventory management. See Bertsekas (2018) for a recent overview of the use of neural networks in approximate policy iteration. When tested in this application, training took an unacceptably long time and underperformed the other method considered.

Extreme gradient boosted regression tree (XGBRegressor) models from the XGBoost Python package performed better than neural networks. XGBoost models are widely used in advanced applications and many successful Kaggle competition winners have used XGBoost models (T. Chen & Guestrin, 2016). After implementing XGBRegressor models as the Q-factor approximation architecture it remained clear that more improved solution was needed.

There is a large body of literature about the best way to fit ML models. However, solving for the best training problem is a major issue outside the scope of this dissertation.

5.3.3. Policy Iteration Based on Policy Improvement by Rollout in Policy Space Approximation Background.

The second API approach changes two interrelated aspects of the first approach. The previous model uses approximation in value space; it approximates Q-factors and defines a policy which must compare the estimated Q-factors of all available options to make a recommendation at each decision point. The approach in this section uses approximation in policy space as well as approximation in value space. It trains a model on decisions made by comparing estimated Q-factors. No Q-factor models are fit. There is a single XGBClassifier model which gives a recommendation directly based on state features. This requires that the algorithm collect (state, recommended action) pairs instead of (state, action, approximate Q-factor) triplets.

5.4.Approximate Dynamic Programming Model Formulation & Solution

The choice of models in this chapter reflects the nature of the problem addressed. The state space is exceptionally large and calculating transition probabilities would be essentially impossible to do with any degree of accuracy. For these reasons it is important that the methods chosen not rely on a lookup table representation and not require calculation of transition probabilities. Both the rollout and approximate policy iteration approaches of this chapter utilize simulation to avoid calculating transition probabilities. Before describing the implementations in detail, this section begins by describing changes to the mathematical model of the previous chapter required to incorporate temporal dynamics.

5.4.1. Changes to Model Formulation to Incorporate Temporal Dynamics

States:

In the emergency model of the previous chapter the Markov chain underlying the MDP evolved through states with inventory count always decreasing. There was no replenishment or aging. Each of the models in this chapter includes aging, outdating, and replenishment.

In the previous chapter the state was represented by $(I, \mathfrak{B})$; a tuple of the inventory position, I, and current patient blood type, $\mathfrak{B}$. $\mathfrak{B}$ was limited to the four ABO types. Both aspects are expanded in this chapter. Previously, $\mathfrak{B}$ strictly described the blood type of the current patient. In this model, time is a factor and the passage of time is captured by expanding $\mathfrak{B}$ to include "Age" as a possibility. Set $\mathfrak{B}$ is now more correctly described as event type rather than patient blood type. Now the non-inventory component of the state description is an event; events are either "Age" events ($\mathfrak{B}_{Age}$) or "Use" ($\mathfrak{B}_{Use}$) events. All eight ABO/Rh types are included in $\mathfrak{B}_{Use}$.

$$\mathfrak{B} = \mathfrak{B}_{Use} \cup \mathfrak{B}_{Age} = \{A+, A-, B+, B-, AB+, AB-, O+, O-, "Age"\} \tag{5.4}$$

Additionally, the inventory variable now tracks the age of each unit.

$$I = (I_{be}) \; \forall \; b \in \mathfrak{B}_{Use}, e = 1 \dots 42 \tag{5.5}$$

Finally, because replenishment is allowed and follows a weekly schedule, the day of the week is included. See Figure 24 for an illustration of the MDP state evolution.

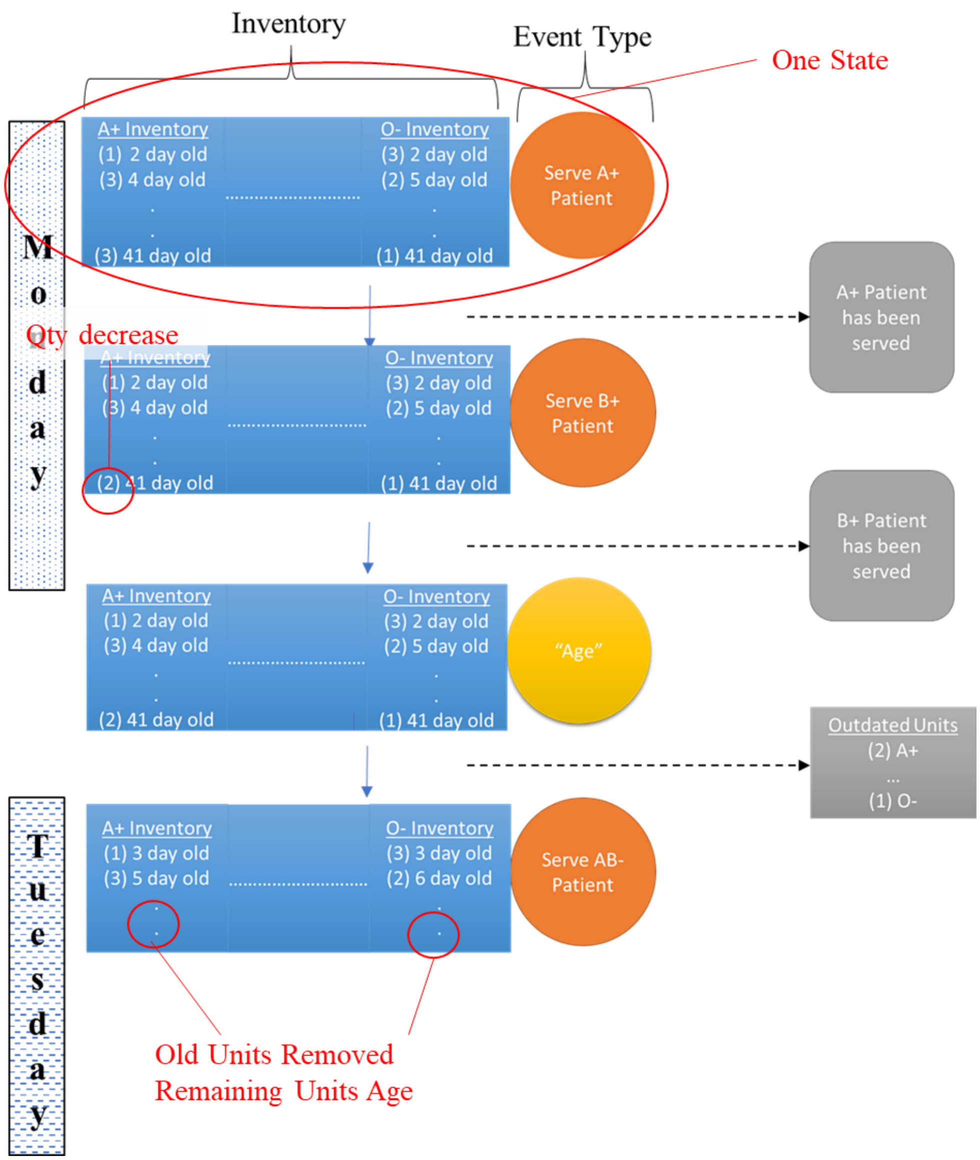

Figure 24. Illustration of the MDP Model for the Improved Allocation of Units Under Normal Conditions Problem

Actions:

The available actions are defined as in the previous chapter

$$a_{(s)} = \{0\} \cup \{\, b: b \in \mathfrak{B}_{Use}, I_b > 0\} \tag{5.6}$$

where $I_b = \sum_{e=1}^{42} I_{be}$ is the current inventory of type b of any age and 0 is the "do-nothing" action.

Transitions:

The model still contains two types of transitions, inventory transitions and transitions between patient blood types. Generation of the blood type of the next patient is as in the previous chapter's model. It is independent between patients and is dependent on the population under study.

Inventory transitions are more complex than before because of the inclusion of aging and replenishment. When the event type is "Use", a patient is demanding blood and transitions are similar to the previous chapter. It is still assumed that each patient demands a single unit. The allocation decision determines the unit type to be issued. In modeling a "Use" transition the oldest unit of the chosen type is removed from inventory to serve the patient and the next event type is generated to complete the state transition:

$$I_{b,e}^{k} \leftarrow I_{b,e}^{k-1} - 1 \tag{5.7}$$

where $e = \max(\{e: I_{b,e} > 0\})$ and b is the allocated blood type.

When the event type is "Age" it means that the end of the day has been reached and all units in inventory must age. When an "Age" event happens at period k, the day of the week advances and inventory is updated as follows:

$$I_{b,1}^{k} \leftarrow N_{b,1}^{k} \; \forall b \in \mathfrak{B}$$

$$I_{b,e}^{k} \leftarrow I_{b,e-1}^{k-1} + N_{b,x=e}^{k} \quad e = 2 \ldots 42, \qquad b \in \mathfrak{B}$$

$$I_{b,e=42}{}^{k-1} \rightarrow \text{units outdating at start of period } k$$

where $N_{b,x=e}^{k}$ is the number of replenishment units of age e, and type b arriving at the start of period k.

In the simulation models used in this chapter the age of the replenishment units,

X, is distributed according to mixture distributions that vary by blood type. It is possible for replenishment units to be older than on-hand units. It should be noted that addressing the non-uniform age of replenishment units is a contribution of this work. For more information on the mixture distributions see Appendix A.

The daily demand, the replenishment schedule with order-up-to levels, and the age distribution of new replenishment units are the simulation model parameters used to describe different hospital scenarios. Values can be found in Appendix D.

Rewards:

Allowable blood type substitutions are shown in Figure 25. Serving a patient with a compatible unit results in a positive reward. All successfully met demands have the same reward. Unmet demand or providing an incompatible unit result in negative reward. This model also includes a penalty cost for each outdated unit. Rewards are denoted $r(s, a)$. They are only state dependent in so far as determining if the unit provided is compatible with the patient's blood types. Current patient blood type is part of the state.

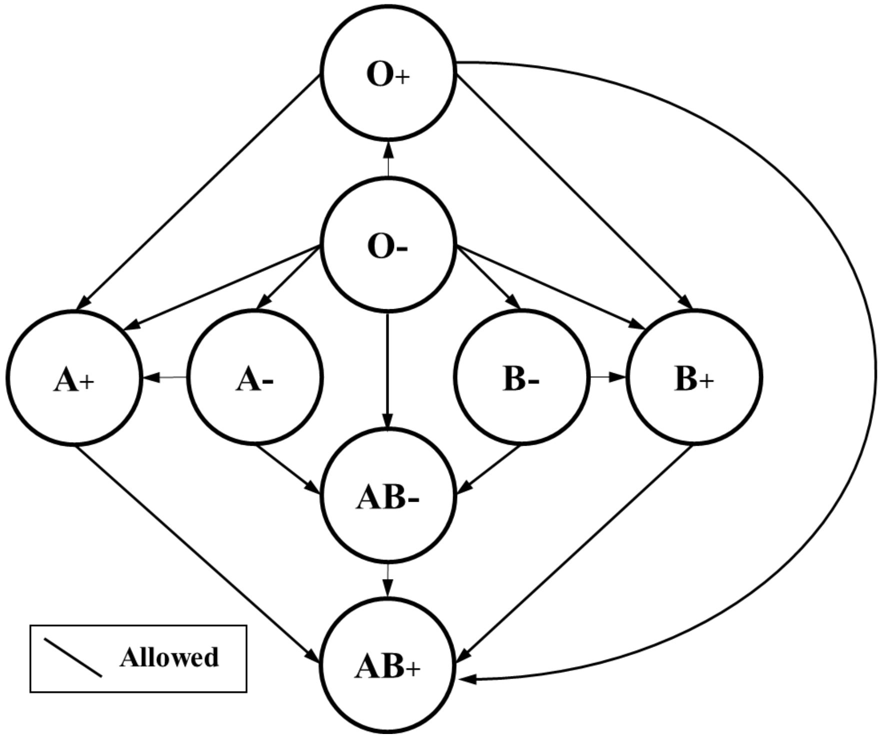

Figure 25. Allowable Blood Type Substitutions in 8-type Model

5.4.2. The Rollout model.

As described in the previous section, from the current state rollout estimates the "cost-to-go" of each available action. These are Q-factor approximations. The action with the best approximate Q-factor is the recommended action.

The model in this section is an online rollout implementation, meaning that decisions are made as they are needed and there is no iteration. This rollout implementation (Figure 26) uses simulations of multiple sample trajectories of a given length to estimate the value of taking each available action. Each of the implementations in this chapter including this rollout implementation utilize common disturbance (random noise) sequences (sample trajectories) when performing value estimation to reduce the variance of the estimates.

Step 0.

Choose base policy μ_{base}

Generate a set of sample trajectories $W = \{\omega_{nk}\}_{n\in[1,\ldots,N],k\in[1,\ldots K]}$

Set S_0 = current state, patient demanding a unit

Step 1. Find set of available action a_{S_0}

Step 2. For each $a \in a_{(s_0)}$

Step 2a. Find $S_1 = f(S_0, a)$

Step 2b.

For k = 1…K

For n = 1… N

Simulate K trajectories of length N states and controls using base policy μ_{base}

Step 2c. Calculate the sample q-factor:

$$\beta^{S_0,a} = 1/K \sum_{k=1}^{K} \Bigg(g\big(S_1, a, \omega_{1,k}\big) + \sum_{n=1}^{N} \alpha^n g\big(S_{1+n}^k, \mu_{base}\big(S_{1+n}^k\big), \omega_{n,k}\big) \Bigg)$$

Step 3. Choose action with highest estimated q-factor:

$$\mathrm{a}' = \arg\max_{a\in a_S} \beta^{S_0,a}$$

Figure 26: Pseudocode for Single Period Rollout Implementation

The rollout approach was evaluated by generating several test sequences of patient arrivals and age events and simulating with the allocation decisions made via rollout (see Figure 27). The simulation included replenishment of units according to the replenishment policy parameters.

The test sequences and the sample trajectories fix the sequence of patients per simulated day including their number and their blood types. The ages of the replenishment units cannot be fixed; this remains a stochastic element.

```
Generate test sequences
For each test_sequence in test_sequences:
          For each step in test_sequence:
     If event_type = Age
          Simulate and record age event
             Else (event_type = Use)
          Options = set of compatible options
          If |Options| =1:
             Simulate and record allocation of that unit type
          Else:
             Use rollout to make allocation decision
```

Figure 27. Pseudocode for Testing Rollout Implementation

The rollout approach resulted in improvement in the Cobre scenario, but it did not result in any positive improvement in the Mayo scenario. Further efforts were necessary; they are described in section 5.4.3.

5.4.3. Approximate Policy Iteration with Q-factor Approximation Model.

The success of rollout for Cobre but not for Mayo highlights an interesting feature of the scenarios. The normal operating conditions modeled in this chapter have high inventory set points so there should be very low unmet demand under all policies. The policies must then differentiate themselves by reducing the number of outdating units.

The API with Q-factor approximation implementation is described in Figure 28. The base policy is improved upon by using rollout to make decisions when simulating the trajectories used to estimate the Q-factors. The base policy is initially the default policy but with each iteration the base policy is updated. In this API with Q-factor approximation the policy was based on comparing estimated Q-factors.

The first approximation architecture tried was a linear regression approximation of the Q-factor using handcrafted features. Linear regression was not a good fit in this case, both in terms of regression R^2 and the improvements shown by the model. This is unsurprising given the non-linear nature of the policies revealed in the previous chapter. Because the nature of the value function was anticipated to be complex, machine learning methods were considered. Neural networks and XGBRegressor models were trained on sets of sample (state, Q-factor estimate) pairs and tuned to optimize fit. XGBRegressor models were chosen over neural networks for having equal or superior fit and lower training time.

Step 0: Initialization

Step 0a: Set the initial policy to default policy, μ^0

Step 0b: Set the iteration counter z = 1.

Step 1: Do for z = 1, . . ., Z, #perform many iterations

Step 2. Generate set of q sample trajectories, each of very short length ℓ.

$$\mathcal{Y} = \{y_{sl}\}_{s=1\ldots q, l=1\ldots\ell}$$

Step 3: Do for $s = 1 \ldots q$:

Step 4: Set initial state i^s, inventory levels &current day of week.

Step 5: For $l = 1 \ldots \ell$:

Step 5: Set current patient blood type, $b(i^{sl}) = \hat{y}_{sl}$

Step 6: Perform rollout to determine action μ' at state i^{sl}

Single period Rollout with shared sample trajectories. See Figure 26

Step 7: For $a \in a_{(i^{sl})}$

Step 8: For trajectory $w \in \mathcal{W}$:

Step 9: Simulate trajectory of states and controls beginning with action a and applying policy μ^{z-1} as base policy

Step 10: Calculate sample Q-factor, $\beta^{s,a}$

$$\beta^{s,a} = g(i^s, a) + \sum_{j=1}^{N} \alpha^j g\left(i^j, \mu^{z-1}\left(i^j\right), w_j\right)$$

Step 12 a. Return $(i^{sl}, a, \beta^{s,a})$

Step 12 b. Simulate optimal action a' to progress state

$$a' = \underset{a \in a_{(i^{sl})}}{\operatorname{argmax}} \beta^{s,a}$$

Step 13: Update Q-factor approximations

$$\tilde{Q}_a^z(i, a, \bar{r}) = F\left(\{(i^{sl}, a, \beta^{s,a}\}^z{}_{s=1\ldots q,\ l=1\ldots\ell,}\right) \ \forall\ a$$

where F is some function for fitting the approximation architecture model; in this case, XGBoost regression model. One regression model fit per action.

Step 7: Update policy, $\tilde{\mu}$

$$\tilde{\mu}(i) \in \ arg \min_{u \in U(i)} \tilde{Q}_\mu(i, u, \bar{r}), \quad i = 1, \ldots, n,$$

Figure 28. API with Q-factor Approximation Implementation.
(Flow may be easier to understand when Step 2 is omitted or if ℓ is considered to be 1. Step 2 allows gathering of ℓ related data points per starting point. In practice $\ell = 2$ was used).

After selecting XGBRegressor, reasonable ranges for tunable XGB parameters were established. The Q-factor models were tuned over these ranges in every API iteration. This adds considerable training time. In addition to the parameters of the approximation architecture, the API algorithm itself has the following additional parameters that needed to be set:

- Length of trajectory: length of simulated trajectory when collecting (i^s, u^s, r) samples
- Number of trajectory restarts from same starting state
- Sample size: how many different starting states (to simulate trajectory and collect sample data from) per iteration
- Number of iterations

5.4.4. Policy Iteration based on Policy Improvement by Rollout and Policy Approximation in Policy Space model.

Application of Q-factors is known to be difficult in cases of large state spaces (W. Powell, 2011) so it should not have been a surprise that the API with Q-factor estimation approach did not produce extraordinary results.

In the API using Q-factor approximation approach there are multiple approximation models to be trained, one for each possible action. When the policy is used to make a decision each Q-factor approximation model makes a prediction, the predicted Q-factors are compared and the action with the highest predicted Q-factor is selected. Each approximation model was trained independently but none of the fits were extremely high.

After reviewing results from initial API with Q-factor approximation attempts, this research examined if a process where the decision was given directly by the approximation equation might give better results. With the Q-factor approximation approach, even though the same set of test sequences are used to get the (state, action, Q-factor) samples for each action there is a lot of variability introduced when the Q-factor approximation equations are trained. In the second API approach examined here (referred to as the direct approach), the comparison between options is made on the Q-factor estimates collected from the same test sequence before the model is fit (see illustration of Figure 29). The samples in this method are (state, recommendation) tuples. This is similar to the concept behind advantage updating or the way that a paired t-test removes variability compared to an unpaired t-test.

Making the recommendation prior to sample collection allows for a model that gives better recommendations even if the model "fit" is the same. Using an approximation architecture that produces an action directly shifts the algorithm (see Figure 30) into the family of approximations in policy space.

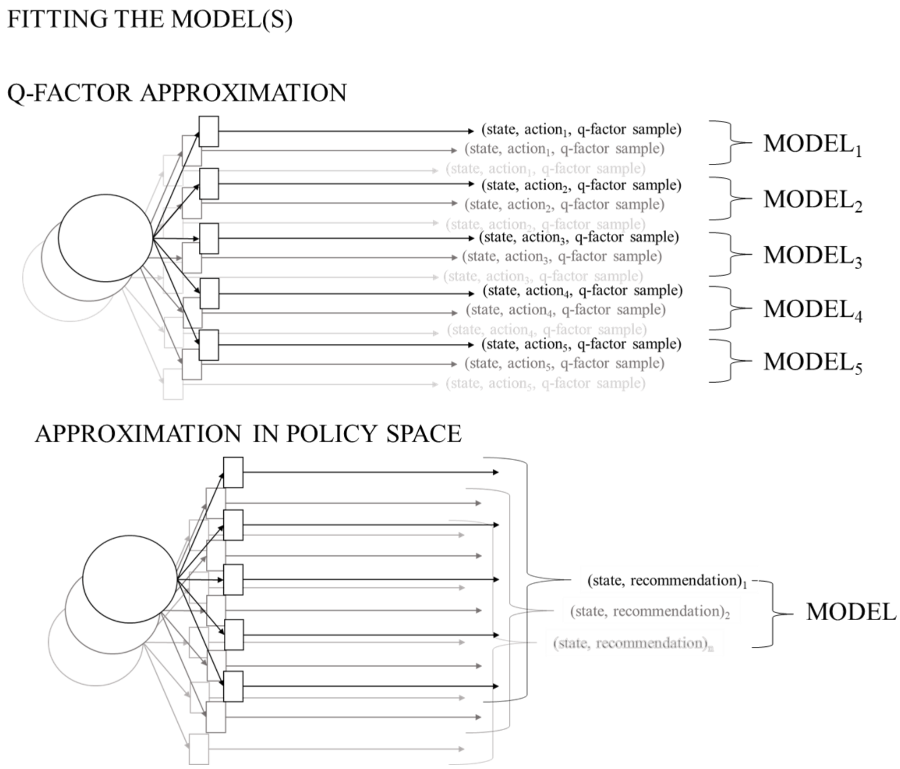

Figure 29. Collection of Samples in Two API models; Q-factor Approximation and Approximation in Policy Space

Step 0: Initialization

Step 0a: Set the initial policy to default policy, μ^0

Step 0b: Set the iteration counter z = 1.

Step 1: Do for z = 1, . . ., Z, #perform many iterations

Step 2. Generate set of q sample trajectories, each of very short length ℓ.

$$\mathcal{Y} = \{y_{sl}\}_{s=1\ldots q, l=1\ldots\ell}$$

Step 3: Do for $s = 1, \ldots, q$:

Step 4: Set initial state i^s, inventory levels &current day of week.

Step 5: For $l = 1, \ldots, \ell$:

Step 5: Set current patient blood type, $b(i^{sl}) = \hat{y}_{sl}$

Step 6: Perform rollout to determine action μ' at state i^{sl}

Single period Rollout with shared sample trajectories. See Figure 26

Step 7: For $a \in a_{(i^{sl})}$

Step 8: For trajectory $w \in \mathcal{W}$:

Step 9: Simulate trajectory of states and controls beginning with action a and applying policy μ^{z-1} as base policy

Step 10: Calculate sample Q-factor, $\beta^{s,a}$

$$\beta^{s,a} = g(i^s, a) + \sum_{j=1}^{N} \alpha^j g\left(i^j, \mu^{z-1}\left(i^j\right), w_j\right)$$

Step 11: find maximum sample Q-factor:

$$\mu'^{(i^{sl})} = arg \max_{a \in a_{(i^{sl})}} \beta^{s,a}$$

Step 12 a. Return $(i^{sl}, \mu'(i^{sl}))$

Step 12 b. Simulate optimal action to progress state

Step 13: Update policy:

$$\mu^{z+1} = F\left(\{(i^{sl}, \mu'(i^{sl})\}^z{}_{s=1\ldots q,\ \ l=1\ldots\ell,}\right)$$

where F is some function for training the approximation architecture model, in this case a boosted classifier. Use this policy for the next iteration.

$(i, \mu'(i))$: state, recommended action in state pair

*current patient blood type selected to upweight blood types with lower compatibility as substitutes so that sufficient samples are collected

Figure 30. Pseudocode for API with Policy Estimation Implementation

5.4.5. Procedures for finding good candidate policies.

This section describes how good models were extracted from the API solutions for further testing. Several different configurations of parameters for each of the API approaches were used to find candidate good policies for each of the hospital scenarios. Five different configurations were used with Cobre and three were used with Mayo. The configurations resulting in the best policies per scenario (that is, configuration “DirectPI_Cobre2” for Cobre and configuration “PIPI_direct_Mayo0.2” for Mayo) are described here next. Parameter details for these and the other configurations can be found in the Appendices C and D.

Each time before initializing the API algorithm, a set of test sequences were defined. These test sequences were simulated each time a new policy approximation was fit (once per iteration). The undiscounted reward, discounted reward, unmet demand, and number of outdating units were recorded. The API MDP models use a discounted reward objective function.

Plots of the undiscounted total reward and discounted total reward by iteration are shown in Figure 31 and 30 for Cobre and Figure 33 for Mayo. Typically, these plots are used to judge if the policy has converged since convergence in general is not guaranteed with API and policy performance will often fluctuate/chatter.

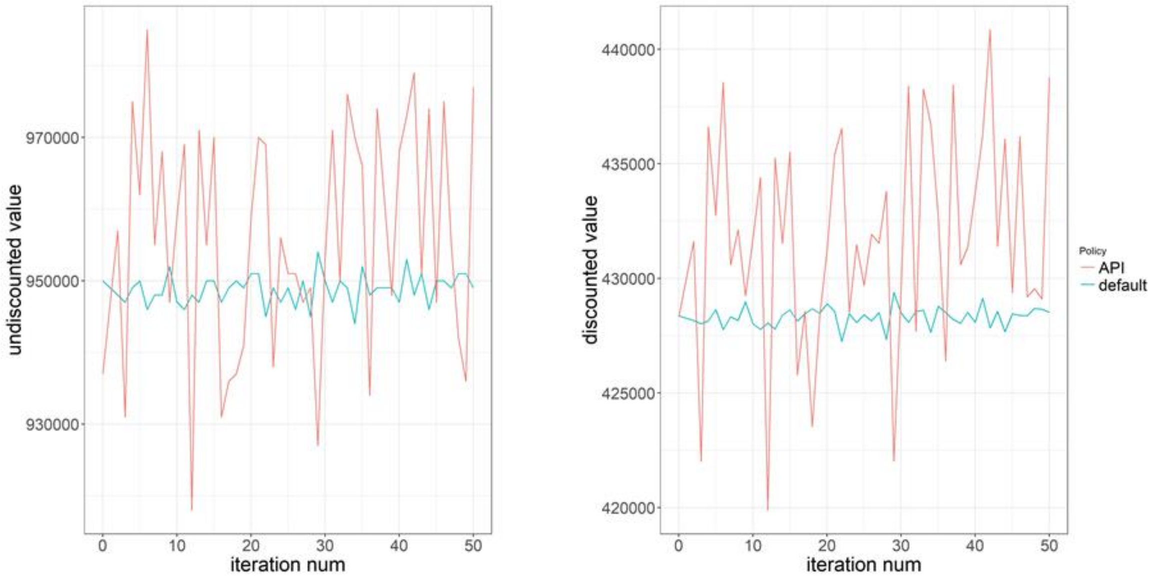

Figure 31. Value by Iteration: *DirectPI_Cobre2*

The value by iteration plots in Figure 31 do not show that the policy converged though they suggest that it may be in the process of converging. Some early iterations performed better on the test sequences than the final iterations.

Because it appeared that DirectPI_Cobre2 was possibly converging towards the end of the first 50 iterations another 100 iterations were completed to see if further convergence occurred. In this second set of runs with the same configuration the API policy is better than the default policy in terms of raw reward and discounted reward in nearly every iteration (see Figure 32). Other than the earliest and last iterations the value is consistent across iterations suggesting that the algorithm has converged.

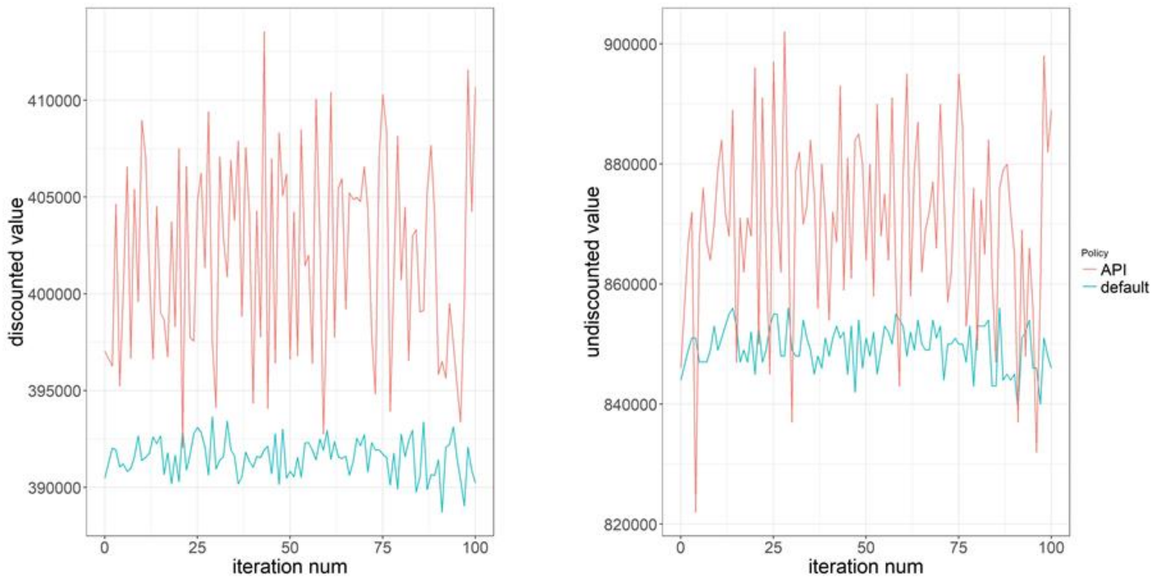

Figure 32. Value by Iteration: DirectPI_Cobre2 Extended Iterations

The policy representations (XGBClassifier models) were saved after each iteration of the API models. It was therefore possible to choose a set of policy models from intermediate iterations and between different API parameter configurations and compare them to each other on the same simulated demand streams.

The iterations from each model with the lowest outdating, highest undiscounted reward, and/or highest discounted reward from the Cobre API configurations are listed in Table 11. Description of the configurations can be found in Appendix C. Simulations were used to compare these policies with each other using the same test sequences per simulation run. Those results are shown later in this chapter.

Table 11

Direct PI models selected for head to head comparison in further simulations: Cobre

	Approach	Configuration	Iteration	Policy Name
1	Fixed Priority	NA	NA	default
2	Direct PI	PIDirect Cobre1	4	PIPI_direct_Cobre_0.14
3	Direct PI	PIDirect Cobre1	20	PIPI_direct_Cobre_0.120
4	Direct PI	PIDirect Cobre1	26	PIPI_direct_Cobre_0.126
5	Direct PI	PIDirect Cobre1	36	PIPI_direct_Cobre_0.136
6	Direct PI	PIDirect Cobre2	40	PIPI_direct_Cobre_0.140
7	Direct PI	PIDirect Cobre2	42	PIPI_direct_Cobre_0.142
8	Direct PI	PIDirect Cobre2	50	PIPI_direct_Cobre_0.150
9	Direct PI	PIDirect Cobre 2 extension	50	PIPI_direct_2extension50
10	Direct PI	PIDirect Cobre 2 extension	100	PIPI_direct_2extension100
11	Direct PI	DirectPI_Cobre2_reduceda	1	PIPI_direct_Cobre2_reduced_param1
12	Direct PI	DirectPI_Cobre2_reduceda	15	PIPI_direct_Cobre2_reduced_param15
13	Direct PI	DirectPI_Cobre2_reduceda	37	PIPI_direct_Cobre2_reduced_param37
14	Direct PI	DirectPI_Cobre2_reduceda	40	PIPI_direct_Cobre2_reduced_param40

The same process was repeated for the Mayo scenario. The plots of value by iteration (see Figure 33) are much less promising. There were a few iterations where the iteration’s undiscounted value exceeded the default policy’s undiscounted value but none where the iteration’s discounted value exceeded the default policy’s discounted value. At best the iteration’s total discounted reward matched that of the default policy. The structure of the MDP model is important here. Because many units are demanded per day in the Mayo scenario, the last unit of the day will be heavily discounted compared to the first and the rewards from (not) outdating will be even more heavily discounted. In the scenarios with a lot of safety-stock, the policies differentiate themselves by decreasing outdating, but (negative) rewards for outdating will be heavily discounted so it makes sense that the API iteration policy is not able to distinguish itself in terms of discounted reward.

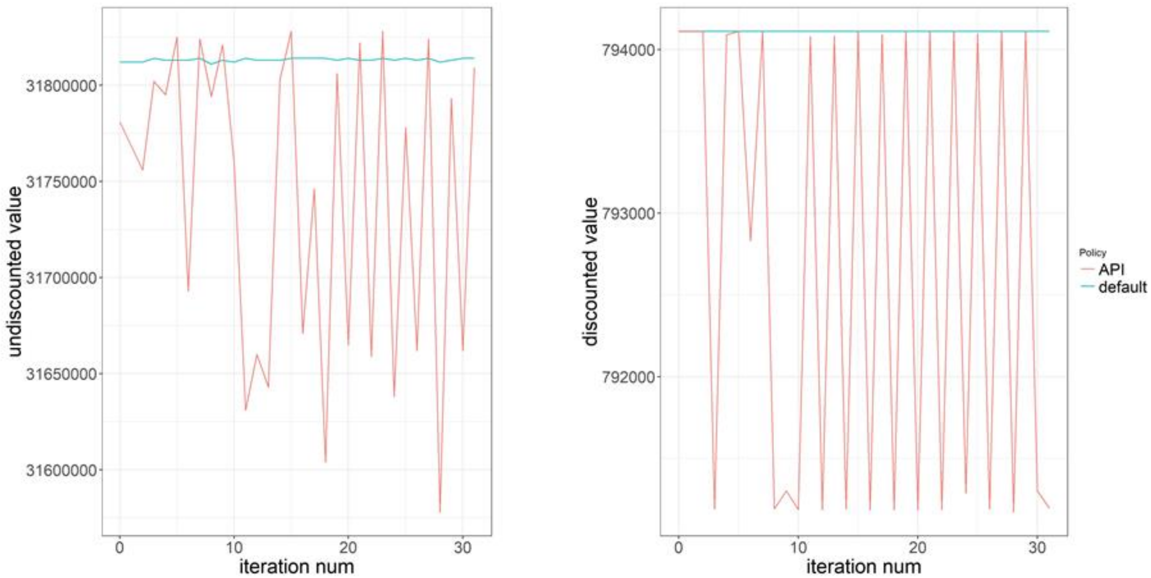

Figure 33. Value by Iteration API Mayo

Models from the iterations of the Mayo API test scenarios that had fewer outdating units, less unmet demand or higher rewards were selected for further study. A list of those policies is found in Table 12 .

Table 12

Direct PI models selected for head to head comparison in further simulation studies: Mayo

	Approach	Configuration	Iteration	Name
1	Fixed priority	NA	NA	default
2	Direct PI	PIPI_direct_Mayo0.1b	2	PIPI_direct_Mayo0.1b2
3	Direct PI	PIPI_direct_Mayo0.1b	3	PIPI_direct_Mayo0.1b3
4	Direct PI	PIPI_direct_Mayo0.2	2	PIPI_direct_Mayo0.22
5	Direct PI	PIPI_direct_Mayo0.2	5	PIPI_direct_Mayo0.25
6	Direct PI	PIPI_direct_Mayo0.2	7	PIPI_direct_Mayo0.27
7	Direct PI	PIPI_direct_Mayo0.2	9	PIPI_direct_Mayo0.29
8	Direct PI	PIPI_direct_Mayo0.2	15	PIPI_direct_Mayo0.215
9	Direct PI	PIPI_direct_Mayo0.2	21	PIPI_direct_Mayo0.221
10	Direct PI	PIPI_direct_Mayo0.2	23	PIPI_direct_Mayo0.223
11	Direct PI	PIPI_direct_Mayo0.2	27	PIPI_direct_Mayo0.227
12	Direct PI	PIPI_direct_Mayo3	4	PIPI_direct_Mayo_noreplenish4
13	Direct PI	PIPI_direct_Mayo3	6	PIPI_direct_Mayo32

5.5.Computational Results and Analysis of Approximate MDPs

The objective function of the ADP models that generated the policies compared in this section was approximate expected discounted reward value. Having a discounted measure is not only necessary for an algorithmic standpoint; it also acknowledges that individuals are able to intervene in the inventory supply chain when potential future shortages are anticipated. In practice emergency orders are common. The difficulties with the discounted reward measure with the current MDP formulation have already been discussed above

This section presents several metrics for evaluating each policy; the amount of unmet demand, the number of units outdating, and three value measures. The first value measure is the undiscounted total reward. The second is the discounted total reward with $\lambda = 0.95$. This is the same objective used in the MDP solution that results in units allocated on the same day having different discounted values. The third metric is the *daily discounted reward*. In these MDP models, fulfilling each unit of demand is a discrete event step and each subsequent step is discounted more deeply than the previous step (see Figure 24). A policy resulting in unmet demand earlier will have a lower discounted reward than a policy resulting in the same amount of unmet demand which occurs later, even if the unmet demand occurs in the same simulated day.

The *daily discounted reward* totals all undiscounted rewards for one simulated day and applies the same discount to all rewards collected on that day. The discount applied on one simulated day will be less than the discount applied on subsequent simulated days. In low demand scenarios the discounted and daily discounted values will be much closer together than in a high demand scenario.

In practice, blood bank managers may want to choose a policy based on one of these value metrics or choose a policy based on the pareto frontier of "outdating" vs "unmet demand".

Unless noted otherwise, each analysis in this results section features five starting points each with 20 test sequences (resulting in 100 different test sequences), with four repetitions for each test sequence. Each simulation run lasted 12 simulated weeks (84 days). The first 42 days served as warmup. Results from the second 42-day period are reported in the analyses. The mean values across the four repetitions for each reward metric were found for each starting point and test sequence combination. The difference between the default policy mean value for each starting point and test sequence combination and the mean for the same starting point and test sequence combination for each test policy were found and two-tailed t-tests (sample size 100 (Cobre) or 120 (Mayo), with type 1 error rate of 5%) were performed on these differences.

5.5.1. Comparison of alternate fixed priority policies.

The policy used as the default policy throughout this dissertation can be described as a fixed priority policy. For this policy the allocation decision is based entirely on priority and does not consider the age of units in the inventory or the relative amount of inventory by type. This policy was selected for use as the default because it is the type of policy used by each of the blood bank managers contacted during this research[2] and is standard in the field (Armstrong et al., 2008). The specific default policy used is the allocation policy in use by Mayo Clinic at the time the research was begun (Dumkrieger et al., 2014).

The default policy is one of several possible fixed priority policies. To confirm that the default policy is not arbitrarily bad and therefore an unfair way to establish improvement over a baseline, several other fixed priority policies were simulated. The results are presented below.

The alternate fixed priority policies are described in Figure 32. In policies *FP1* and *FP3* ABO type takes priority over Rh+/-. In policy *FP1*, which is the default policy, blood type A takes priority over B and in *FP2* B takes priority over A. In fixed priority policies *FP2* and *FP4* Rh+ units are prioritized over Rh- units for A and B but all A and B types have priority over O. In *FP2* blood type A is prioritized over B and in *FP3* the

---

[2] In one case the manager had decades of experience and would occasionally make alternative allocation decision based on the relative inventories and age of units on hand. That manager achieved very good results unfortunately the strategy of hiring an expert to oversee the allocation of each unit may be impractical on a global scale.

opposite is true. In policies *FP5* and *FP6* Rh+ units are prioritized over Rh- for ABO, meaning that O+ has priority over A- or B-. In *FP5* blood type A has priority over B and in *FP6* blood type B has priority over A. In all cases all compatible AB units have priority over units of any other compatible type.

Fixed Priority Policy

FP1. Afirst_alt (default)

Patient BT	Priority
A+	[A+,A-,O+,O-]
A-	[A-,O-]
B+	[B+,B-,O+,O-]
B-	[B-,O-]
AB+	[AB+,AB-,A+,A-,B+,B-,O+,O-]
AB-	[AB-,A-,B-,O-]
O+	[O+,O-]
O-	[O-]

FP3. Bfirst_alt

Patient BT	Priority
A+	[A+,A-,O+,O-]
A-	[A-,O-]
B+	[B+,B-,O+,O-]
B-	[B-,O-]
AB+	[AB+,AB-,B+,B-,A+,A-,O+,O-]
AB-	[AB-,B-,A-,O-]
O+	[O+,O-]
O-	[O-]

FP2. ABfirst_pos

Patient BT	Priority
A+	[A+,A-,O+,O-]
A-	[A-,O-]
B+	[B+,B-,O+,O-]
B-	[B-,O-]
AB+	[AB+,AB-,A+,B+,A-,B-,O+,O-]
AB-	[AB-,A-,B-,O-]
O+	[O+,O-]
O-	[O-]

FP4. BAfirst_pos

Patient BT	Priority
A+	[A+,A-,O+,O-]
A-	[A-,O-]
B+	[B+,B-,O+,O-]
B-	[B-,O-]
AB+	[AB+,AB-,B+,A+,A-,B-,O+,O-]
AB-	[AB-,B-,A-,O-]
O+	[O+,O-]
O-	[O-]

FP5. ABOfirst_pos

Patient BT	Priority
A+	[A+,O+,A-,O-]
A-	[A-,O-]
B+	[B+,O+,B-,O-]
B-	[B-,O-]
AB+	[AB+,AB-,A+,B+,O+,A-,B-,O-]
AB-	[AB-,A-,B-,O-]
O+	[O+,O-]
O-	[O-]

FP6. BAOfirst_pos

Patient BT	Priority
A+	[A+,O+,A-,O-]
A-	[A-,O-]
B+	[B+,O+,B-,O-]
B-	[B-,O-]
AB+	[AB+,AB-,B+,A+,O+,B-,A-,O-]
AB-	[AB-,B-,A-,O-]
O+	[O+,O-]
O-	[O-]

Figure 34. Fixed Priority Policies

The alternate fixed priority policies were simulated for both Cobre and Mayo scenarios. For the Cobre scenario five different starting points, each with 20 different test sequences, were simulated. For the Mayo scenario 6 different starting points were used.

Results for the low demand scenario, Cobre, are shown in Table 13. Results for the high demand scenario, Mayo, are shown in Table 14. For the Cobre scenario none of the alternate fixed priority policies performed statistically better than the default policy. *FP3*, *FP5*, and *FP6* performed better than default but the difference was not significant. For the Mayo scenario *FP3* is consistently better than the default policy but the difference was not statistically significant. Thus, it can be concluded that the default policy is a reasonable one to use as a baseline.

Blood bank managers may still wish to consider these policies based on their own situations. For example, if blood type B is more available in their region or less expensive, they may prefer a policy which prioritizes type B blood. Because the performance of all of these policies are dependent on the characteristics of the hospital and patient population any change should be proceeded by simulation of that particular facilities situation.

Table 13

Simulation Results: Cobre Alternate Fixed Priority Policy Simulation Comparison.

	Unmet Demand		Outdated Units			Undiscounted Reward		Discounted Reward		Daily Discount	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.09	0.015%	299.59	33.7%	0.0%	0.00	NA	0	NA	0.00	NA
FP1	0.09	0.015%	303.05	34.0%	-1.2%	-400.00	0.146	-11	0.917	-91.64	0.522
FP3	0.09	0.015%	301.47	33.8%	-0.6%	-217.50	0.413	66	0.533	5.43	0.970
FP4	0.11	0.018%	302.90	34.0%	-1.1%	-400.00	0.141	-21	0.845	-112.68	0.458
FP5	0.11	0.018%	300.45	33.8%	-0.3%	-117.50	0.662	107	0.351	68.78	0.658
FP6	0.13	0.022%	303.96	34.0%	-1.5%	-540.00	0.063	148	0.124	39.83	0.775

*period is 42 days

Table 14

Simulation Results: Mayo Alternate Fixed Priority Policy Simulation Comparison. 6 starting points with 20 test sequences Each sequence repeated four times.

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.00	0.000%	790	1.9%	0.0%	0	NA	0	NA	0.00	NA
FP1	0.00	0.000%	792	1.9%	-0.3%	-242	0.771	-15	0.460	-453	0.267
FP3	**0.00**	**0.000%**	**774**	**1.9%**	**2.0%**	**1,858**	**0.088**	**8**	**0.797**	**305**	**0.579**
FP4	0.00	0.000%	772	1.9%	2.3%	2,090	0.038	-107	0.071	544	0.291
FP5	0.00	0.000%	794	1.9%	-0.5%	-467	0.535	1	0.978	-231	0.540
FP6	0.00	0.000%	764	1.9%	3.3%	3,000	0.003	-128	0.035	676	0.189

*period is 42 days

5.5.2. Rollout results.

The same test sequences and starting points used in the alternate fixed policy comparison for Cobre were simulated using rollout. Rollout performed similarly to the default policy for Cobre (Table 15). There was slightly more unmet demand with the rollout simulation (0.044%) vs the default simulation (0.015%) and a small (6.6%) decrease in outdated units. Despite this relatively small improvement the difference in each of the value metrics is significant (improvement: undiscounted reward 2210, $p < 0.001$; discounted reward 431, $p = 0.009$; daily discounted reward 1005, $p = < 0.001$). The test sequences and starting points are used again in the main Cobre API simulations.

Rollout did not improve upon the default policy at all for Mayo (therefore computational results are not shown). The performance of the two policies were the same across all metrics for that scenario. The default policy and the rollout implementation did not make the same decisions at each step of the test sequences, but they still ended up in the same place when it came time to outdate. The frequent replenishment in the Mayo scenario may have also contributed.

Table 15

Simulation Results: Cobre Rollout Comparison

Policy Name	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reductio n from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.09	0.015%	307	34.3%	NA	0	NA	0	NA	0	NA
rollout	0.26	0.044%	286	32.7%	6.6%	2210	7.89E-09	431	9.07E-03	1005	1.98E-06

5.5.1. Results of Policy Iteration based on Policy Improvement by Rollout and Policy Approximation in Policy Space.

5.5.1.1. Cobre API results.

The 14 candidate policies selected from across the API configurations (see Table 11) were simulated on the same set of tests sequences and starting points that were used in the Cobre rollout and Cobre fixed-priority policy comparisons. Results are presented in Table 16.

The best Cobre policy is *PIPI_direct_2extension50*. When simulated (with the same 100 test sequences as used in the rollout comparisons) it shows positive results for improvement in discounted and daily discounted reward compared to the default policy, but these differences are not significant. Noticing that few of the results are significant in any direction it was considered that the power of the test may be too low. Two-hundred additional simulation runs, with four repetitions each, were run for *PIPI_direct_2extension50* and default policies only (Table 17, Figure 36). With the increased test power the improvement in undiscounted reward and the improvement in daily discounted reward are significant (undiscounted: p=0.019, discounted p = 0.017), while the improvement in discounted reward continues to fail to reach significance (p = 0.084).

The rollout policy and *PIPI_direct_2extension50* were simulated on the same test sequences, enabling comparison. The difference in performance between the two policies was calculate and t-tests were performed. The rollout policy outperformed the API in each of the value metrics (undiscounted mean improvement: 1760, p <0.001; discounted mean improvement: 1000, $p < 0.001$; daily discounted mean improvement: 619, $p = 0.003$). The improvement in reward metrics by the rollout policy are shown in Figure 37.

Table 16

API Simulation Results: Cobre Standard policy comparison

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period *	p-value
default	0.11	0.018%	303	34.0%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_Cobre_0.14	0.98	0.165%	315	34.9%	-4.1%	-2,118	1E-06	-103	4E-01	-640	7E-04
PIPI_direct_Cobre_0.120	0.33	0.055%	306	34.2%	-1.1%	-555	5E-02	81	4E-01	**-101**	5E-01
PIPI_direct_Cobre_0.136	0.35	0.059%	345	36.9%	-13.9%	-5,053	3E-22	-592	3E-05	-1,861	5E-15
PIPI_direct_Cobre_0.140	0.56	0.096%	327	35.7%	-7.9%	-3,133	1E-12	-358	5E-03	-1,141	2E-08
PIPI_direct_Cobre_0.26	2.34	0.397%	409	41.1%	-35.2%	-14,095	3E-49	-2,367	2E-23	-5,820	1E-39
PIPI_direct_Cobre_0.242	1.91	0.324%	381	39.3%	-25.7%	-10,448	7E-34	-1,552	6E-20	-4,198	3E-31
PIPI_direct_Cobre_0.250	1.04	0.176%	381	39.3%	-25.7%	-9,728	2E-34	-1,263	2E-12	-3,879	3E-28
PIPI_direct_2extension50	**0.24**	**0.040%**	**302**	**33.9%**	**0.2%**	**-25**	**9E-01**	**130**	**2E-01**	**99**	**5E-01**
PIPI_direct_2extension100	0.65	0.110%	361	38.0%	-19.0%	-7,088	3E-28	-947	1E-09	-2,802	1E-21
PIPI_direct_Cobre2_reduced_param1	1.97	0.335%	335	36.3%	-10.6%	-5,203	7E-15	-310	3E-02	-1,739	1E-09
PIPI_direct_Cobre2_reduced_param15	1.41	0.239%	334	36.2%	-10.4%	-4,668	4E-14	-399	4E-03	-1,580	8E-11
PIPI_direct_Cobre2_reduced_param37	0.67	0.114%	320	35.3%	-5.8%	-2,485	3E-09	-184	2E-01	-933	9E-06
PIPI_direct_Cobre2_reduced_param40	1.32	0.224%	322	35.4%	-6.4%	-3,208	8E-10	-29	8E-01	-998	2E-05

Table 17

API Simulation Results: Additional Basic comparisons runs for Cobre

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.03	0.005%	306	34.0%	0.0%	0.0	NA	0	NA	0.00	NA
PIPI_direct_2extension50	**0.29**	**0.049%**	**300**	**33.6%**	**2.0%**	**486.3**	**1.9E-02**	**144**	**8.4E-02**	**274.15**	**1.7E-02**

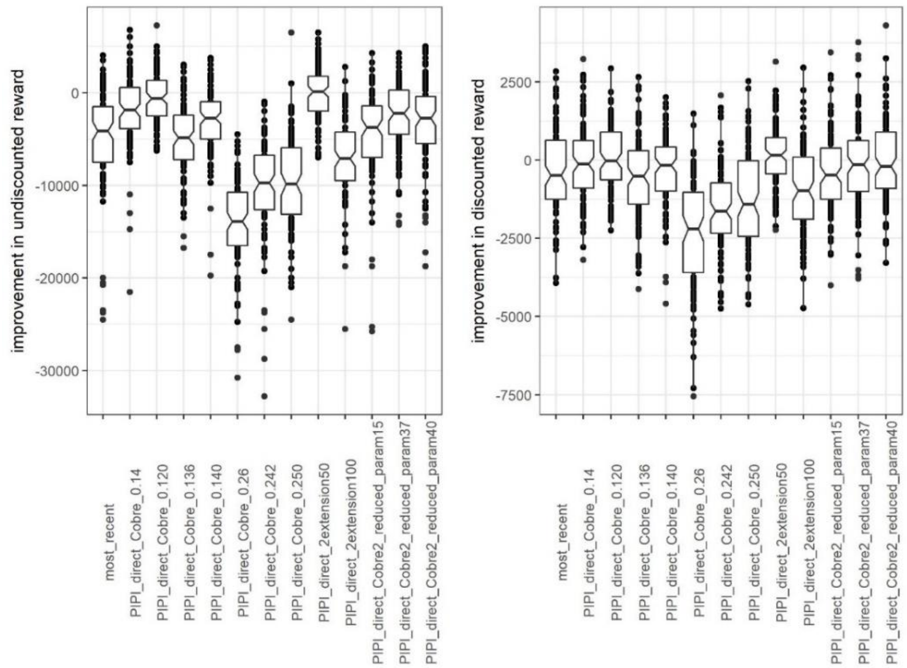

Figure 35. Cobre Standard Comparison Reward Boxplots

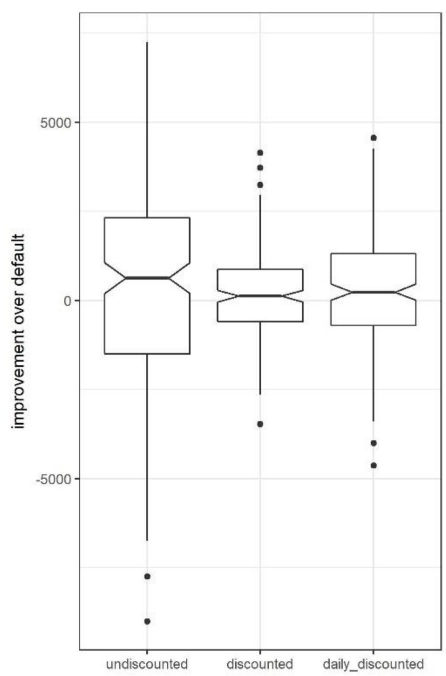

Figure 36. Cobre *PIPI_direct_2extension50* Standard Extra Runs Boxplot

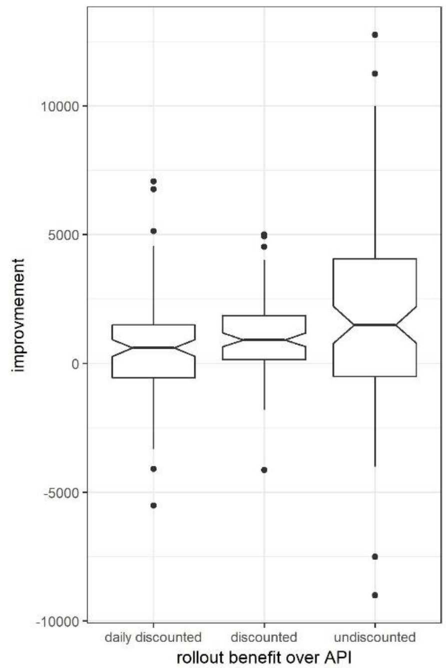

Figure 37. Cobre Improvement of Rollout Over API

Table 18

Units Allocated by Type, by Policy: Cobre

Policy	Unmet Units	Allocated Type							
		A+	A-	B+	B-	AB+	AB-	O+	O-
default	0.02%	33%	7%	8%	2%	3%	1%	37%	8%
PIPI_direct_Cobre_0.14	0.17%	31%	6%	9%	1%	2%	0%	31%	20%
PIPI_direct_Cobre_0.120	0.06%	35%	5%	6%	3%	1%	1%	36%	14%
PIPI_direct_Cobre_0.136	0.06%	14%	5%	4%	5%	2%	1%	58%	12%
PIPI_direct_Cobre_0.140	0.10%	24%	6%	1%	9%	1%	0%	42%	17%
PIPI_direct_Cobre_0.26	0.40%	17%	16%	2%	9%	2%	0%	25%	29%
PIPI_direct_Cobre_0.242	0.32%	10%	18%	3%	8%	0%	1%	41%	19%
PIPI_direct_Cobre_0.250	0.18%	4%	11%	2%	6%	2%	2%	53%	20%
PIPI_direct_2extension50	0.04%	33%	7%	5%	2%	2%	1%	37%	12%
PIPI_direct_2extension100	0.11%	9%	12%	0%	6%	3%	1%	48%	20%
PIPI_direct_Cobre2_reduced_param1	0.33%	36%	0%	0%	0%	0%	0%	42%	22%
PIPI_direct_Cobre2_reduced_param15	0.24%	30%	8%	0%	0%	0%	0%	42%	20%
PIPI_direct_Cobre2_reduced_param37	0.11%	36%	3%	0%	0%	0%	0%	45%	16%
PIPI_direct_Cobre2_reduced_param40	0.22%	36%	0%	0%	0%	0%	0%	46%	18%

Table 19

To:From matrix: Cobre Default

Patient Blood Type	A+	A-	B+	B-	AB+	AB-	O+	O-	% of Allocated
	Allocated Units								
A+	33.1	-	-	-	-	-	-	-	33.1
A-	-	7.1	-	-	-	-	-	0.0	7.1
B+	-	-	8.4	0.1	-	-	0.0	-	8.5
B-	-	-	-	1.8	-	-	-	0.1	1.8
AB+	0.0	-	-	-	3.1	0.0	-	-	3.1
AB-	-	0.0	-	-	-	1.2	-	-	1.2
O+	-	-	-	-	-	-	37.0	-	37.0
O-	-	-	-	-	-	-	-	8.1	8.1
Total Allocated	33.1	7.1	8.4	1.8	3.1	1.2	37.0	8.2	100
Outdated, % of outdating units	17.7	7.2	5.9	7.5	15.2	8.4	25.8	12.1	100
Total Units Ordered, % of total	27.9	7.1	7.6	3.8	7.2	3.7	33.2	9.6	100
Outdated, % of Total Ordered	21.6	34.2	26.7	67.9	71.8	78.0	26.4	43.1	33.96

Table 20

To:From matrix: Cobre PIPI_direct_2extension50

Patient Blood Type	A+	A-	B+	B-	AB+	AB-	O+	O-	% of Allocated
	Allocated Units								
A+	33.1	0.0	-	-	-	-	-	-	33.2
A-	-	7.1	-	-	-	-	-	0.0	7.1
B+	-	-	4.5	0.8	-	-	0.0	3.2	8.5
B-	-	-	-	1.3	-	-	-	0.6	1.8
AB+	0.0	0.0	0.0	-	2.4	0.6	-	-	3.1
AB-	-	0.1	-	0.3	-	0.8	-	-	1.2
O+	-	-	-	-	-	-	37.0	-	37.0
O-	-	-	-	-	-	-	-	8.1	8.1
Total Allocated	33.2	7.2	4.5	2.4	2.4	1.4	37.0	11.9	100
Outdated, % of outdating units	17.8	7.2	9.8	6.6	16.3	7.8	26.1	8.4	100
Total Units Ordered, % of total	28.0	7.2	6.3	3.8	7.1	3.6	33.3	10.7	100
Outdated, % of Total Ordered	21.6	33.9	52.7	58.8	77.6	73.8	26.6	26.7	33.9

Another factor that blood bank managers or regional blood bank managers are interested in is the percentage of demand being met with units of the same blood type. This information can be found in Table 19 for the default policy or Table 20 for

PIPI_direct_2extension50. Under the default policy 99.8% of demand is allocated a direct match and 99.9% are allocated an ABO match. Only 0.3% of O goes to non-O patients. Under *PIPI_direct_2extension50* 94.3% of demand is allocated a direct match, 95.7% is allocated an ABO match and 7.7% of O goes to non-O patients. 37.0% of allocated units are O+ under both policies but a smaller percentage of allocated units are O- under the default policy (8.1%) vs *PIPI_direct_2extension50* (11.9%). While *PIPI_direct_2extension50* orders and allocates more O- units it outdates a smaller percentage of the O- units it orders than the default policy. These values would depend on the replenishment policy as well. For the API policies, a replenishment policy that orders fewer O units, or if younger O units were provided, it would be expected that fewer O units would be allocated.

5.5.1.2. **Mayo API results.**

PIPI_direct_Mayo0.221 and *PIPI_direct_Mayo0.223* are the most promising policies for the Mayo scenario. As seen in Table 21 and Figure 38 both policies result in slightly more unmet demand but also a large reduction in outdated units compared to the default policy. *PIPI_direct_Mayo0.221* is significantly better than the default policy in terms of undiscounted reward ($p<0.001$), and daily discounted reward ($p<0.001$) but not discounted reward ($p=0.097$). *PIPI_direct_Mayo0.223* is not significantly better in terms of discounted reward ($p=0.314$) but is better in terms of undiscounted reward and daily discounted reward ($p<0.001$). *PIPI_direct_Mayo0.223* shows more improvement in outdating than *PIPI_directMayo0.221* which may make it attractive to a blood bank manager.

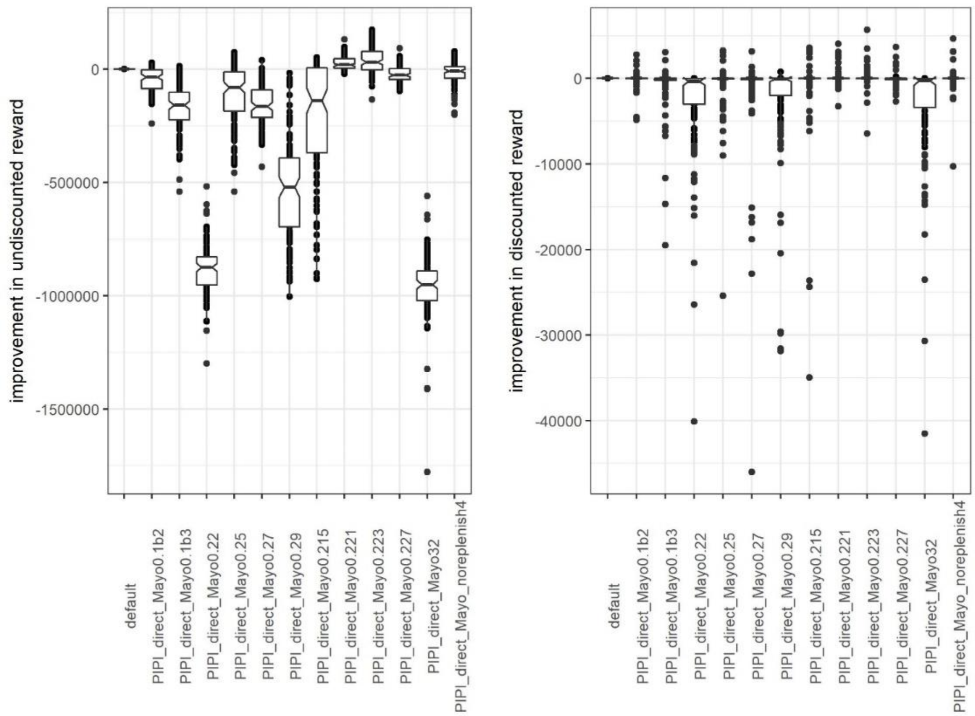

Figure 38. Mayo API Standard Comparison Boxplots. Figures show the improvement of each new policy over the default policy in terms of undiscounted reward and discounted reward.

Table 21

API Simulation Results: Mayo Standard Policy Comparison. 6 starting points with 20 test sequences of 42 days each. 42 day warmup. Each sequence repeated four times. The same test sequences are used in Table 14.

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.00	0.000%	789	1.9%	0.0%	0.00	NA	0	NA	0	NA
PIPI_direct_Mayo_noreplenish4	2.85	0.007%	926	2.3%	-17.4%	-18,173	1.4E-04	-47	6.7E-01	-7,976	5.0E-04
PIPI_direct_Mayo0.1b2	0.00	0.000%	1,189	2.9%	-50.8%	-46,227	6.4E-17	-119	1.4E-01	-20,909	2.5E-14
PIPI_direct_Mayo0.1b3	31.43	0.078%	1,978	4.7%	-150.8%	-162,625	9.2E-34	-767	2.6E-03	-69,541	3.1E-30
PIPI_direct_Mayo0.22	43.71	0.109%	8,133	16.8%	-931.5%	-882,769	1.5E-108	-2,930	1.1E-07	-366,490	6.7E-86
PIPI_direct_Mayo0.25	4.84	0.012%	1,763	4.2%	-123.6%	-116,342	1.6E-17	-583	2.4E-02	-51,473	2.7E-15
PIPI_direct_Mayo0.27	0.40	0.001%	2,176	5.1%	-175.9%	-160,373	2.4E-40	-1,347	8.7E-03	-68,161	3.3E-31
PIPI_direct_Mayo0.29	19.30	0.048%	5,216	11.5%	-561.5%	-526,492	4.7E-51	-2,549	1.8E-05	-223,922	9.5E-45
PIPI_direct_Mayo0.215	2.04	0.005%	2,721	6.3%	-245.1%	-224,602	2.1E-16	-882	3.6E-02	-89,678	8.0E-14
PIPI_direct_Mayo0.221	**0.22**	**0.001%**	**579**	**1.4%**	**26.5%**	**23,942**	**8.8E-15**	**112**	**9.7E-02**	**10,434**	**3.2E-12**
PIPI_direct_Mayo0.223	**2.22**	**0.006%**	**444**	**1.1%**	**43.7%**	**37,956**	**3.3E-11**	**92**	**3.1E-01**	**16,028**	**6.2E-09**
PIPI_direct_Mayo0.227	0.36	0.001%	961	2.3%	-21.8%	-20,156	5.3E-09	-80	1.8E-01	-8,659	4.2E-07
PIPI_direct_Mayo32	65.31	0.162%	8,664	17.8%	-998.7%	-961,419	2.4E-100	-3,100	1.9E-07	-390,668	2.3E-95

Table 22

Units Allocated by Type, by Policy: Mayo

		Allocated Type							
	Unmet Units	A+	A-	B+	B-	AB+	AB-	O+	O-
Default	0%	39%	8%	10%	1%	0%	0%	35%	6%
PIPI_direct_Mayo0.1b2	0%	36%	7%	6%	9%	0%	0%	35%	7%
PIPI_direct_Mayo0.1b3	0%	36%	7%	9%	6%	0%	0%	27%	14%
PIPI_direct_Mayo0.22	0%	1%	9%	9%	5%	0%	0%	44%	30%
PIPI_direct_Mayo0.25	0%	34%	9%	8%	3%	0%	0%	30%	17%
PIPI_direct_Mayo0.27	0%	36%	9%	6%	3%	0%	0%	24%	21%
PIPI_direct_Mayo0.29	0%	36%	5%	6%	2%	0%	0%	19%	32%
PIPI_direct_Mayo0.215	0%	24%	10%	10%	4%	0%	0%	34%	19%
PIPI_direct_Mayo0.221	0%	36%	8%	7%	1%	0%	0%	34%	14%
PIPI_direct_Mayo0.223	0%	28%	8%	10%	3%	0%	0%	40%	11%
PIPI_direct_Mayo0.227	0%	37%	7%	7%	0%	0%	0%	36%	13%
PIPI_direct_Mayo_noreplenish4	0%	27%	13%	7%	3%	0%	0%	39%	12%
PIPI_direct_Mayo32	0%	3%	13%	3%	5%	0%	0%	59%	17%

Table 23

To:From Matrix: Mayo default

	Allocated Units								% of
Patient Blood Type	A+	A-	B+	B-	AB+	AB-	O+	O-	Allocated
A+	36.0	-	-	-	-	-	-	-	36.0
A-	-	7.0	-	-	-	-	-	-	7.0
B+	-	-	10.0	-	-	-	-	-	10.0
B-	-	-	-	1.0	-	-	-	-	1.0
AB+	3.4	-	-	-	0.4	0.2	-	-	4.0
AB-	-	0.9	-	-	-	0.1	-	-	1.0
O+	-	-	-	-	-	-	35.0	-	35.0
O-	-	-	-	-	-	-	-	6.0	6.0
Total Allocated	39.4	7.8	10.0	1.0	0.4	0.4	35.0	6.0	100
Outdated, % of outdating units	0.3	12.9	0.6	14.9	0.0	0.0	29.4	41.9	100
Total Units Ordered	38.7	7.9	9.8	1.3	0.4	0.4	34.9	6.7	100
Outdated, as % of Total Ordered	0.0	3.1	0.1	22.7	0.0	0.0	1.6	12.0	1.9

Table 24

To:From Matrix: Mayo PIPI_direct_mayo0.221

Patient Blood Type	A+	A-	B+	B-	AB+	AB-	O+	O-	% Allocated
	Allocated Units								
A+	35.3	0.5	-	-	-	-	0.3	0.0	36.0
A-	-	6.9	-	-	-	-	-	0.0	7.0
B+	-	-	4.8	0.1	-	-	0.0	5.1	10.0
B-	-	-	-	0.2	-	-	-	0.8	1.0
AB+	1.1	0.6	1.9	0.0	-	0.1	0.2	0.0	4.0
AB-	-	0.3		0.5		0.1		0.0	1.0
O+	-						33.6	1.4	35.0
O-	-							6.0	6.0
Total Allocated	36.4	8.3	6.7	0.8	-	0.2	34.0	13.5	100
Outdated, % of outdating units	0.3	12.9	0.6	14.9	0.0	0.0	29.4	41.9	100
Total Units Ordered	35.7	8.4	6.6	1.1	0.0	0.2	34.0	14.1	100
Outdated, as % of Total Ordered	0.0	2.9	0.2	27.1	100	0.1	1.7	5.7	1.9

With the default policy 96% of demand is met with direct match and 96% is met with an ABO match. 0% of O blood goes to non-O patients. The new policy has direct matching for only 87% of demand and ABO match for 89% of demand. 13.8% of allocated O units go to non-O patients. The default policy has high direct match percentage and a lower percentage of O units allocated to non-O patients, but it also has more O units outdating. 12.0% of ordered O- units outdate under the default policy, but only 5.7% of ordered O- units outdate under the new policy.

5.5.1.3. **Sensitivity Analysis**

Because these policies are tailored to the specific profile of the simulated hospital it is important to check that they are robust to some misspecification in the relevant parameters. This section evaluates how the best direct API policies found for Cobre and Mayo perform when the demand is altered or when the ages of the replenishment units change.

Cobre Scenario

When the demand is approximately 10% higher than the demand used to find the policy *PIPI_direct_2extension50* performs approximately the same as the default policy (Table 25). When the demand is approximately 10% lower than used to find the policy *PIPI_direct_2extension50* performs better than the default across all three metrics (undiscounted reward $p = 0.009$; discounted reward $p = 0.26$; daily discount $p = 0.032$) (see Table *26*). The difference between the policies in terms of discounted reward is not significant but it affirms that the new policy will not perform worse than the default policy even if the actual demand is 10% lower than modeled. This new policy may perform even better at lower demands but increased demand would likely cause it to perform less well than default. There is a limitation to these models in that the replenishment policy was the same in all Cobre simulations. If there was a significant change in demand then the replenishment policy would likely change to reflect the demand shift. However, if it were a matter of misspecification then this analysis provides an accurate estimation of the effect.

Table 25

API Simulation Results: Cobre Increased Demand

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.01	0.002%	263	28.5%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_2extension50	0.29	0.044%	260	28.3%	1.3%	158	4.8E-01	68	3.8E-01	127	2.5E-01

Table 26

API Simulation Results: Cobre Decreased Demand

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.03	0.006%	358	41.6%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_2extension50	0.13	0.026%	352	41.2%	1.6%	584	8.7E-03	98	2.6E-01	247	3.2E-02

Table 27

API Simulation Results: Cobre Older Age. Each distribution of the mixture distribution mean age increased by 1.

Policy Name	Unmet Demand: mean annual	Unmet Demand: % Demand Unmet	Outdated Units: mean annual	Outdated Units: % of Total Units	Outdated Units: % Reduction from default	Undiscounted Reward Improvement Over default: mean per period*	Undiscounted Reward Improvement Over default: p-value	Discounted Reward Improvement Over Default: mean per period*	Discounted Reward Improvement Over Default: p-value	Daily Discount Improvement Over Default: mean per period*	Daily Discount Improvement Over Default: p-value
default	0.07	0.011%	365	38.3%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_2extension50	0.27	0.046%	360	37.9%	1.5%	446	5.3E-02	60	4.9E-01	128	3.1E-01

Table 28

API Simulation Results: Cobre Young Age

Policy Name	Unmet Demand: mean annual**	Unmet Demand: % Demand Unmet	Outdated Units: mean annual**	Outdated Units: % of Total Units	Outdated Units: % Reduction from default	Undiscounted Reward Improvement Over default: mean per period*	Undiscounted Reward Improvement Over default: p-value	Discounted Reward Improvement Over Default: mean per period*	Discounted Reward Improvement Over Default: p-value	Daily Discount Improvement Over Default: mean per period*	Daily Discount Improvement Over Default: p-value
default	0.04	0.007%	120	17.1%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_2extension50	0.56	0.097%	116	16.7%	2.9%	-15	9.5E-01	-69	3.6E-01	-61	6.0E-01

The age of replenishment units is an important factor in blood inventory management. Table 27 shows the results of comparison between the API policy and the default policy when the mean age per each component of the age mixture distribution is increased by one. The improvement of each metric is positive but none of the differences are significant any longer.

The low demand scenario (Cobre) is representative of the many small rural hospitals across the United States. These hospitals often receive fresher units than the larger hospitals and their aging stock will be retrieved by the regional blood provider and rotated to higher demand hospitals. However, none of the small hospitals contacted were able to provide data on the actual age of their new units or the age of units collected for rotation. For that reason, the same age distribution as the large hospital were used in all simulations except those presented as Cobre Young Age in Table *28*. In these simulations the age distribution for each blood type was distributed according to a normal distribution with mean 5 and standard deviation 2. This is an estimate of the youngest units that might reach one of the rural hospitals. The simulations show that the new policy does not perform better than the default policy when the age is very young, but it does not perform statistically worse either. This does not invalidate the results of simulations using the older age distributions. If a hospital were going to implement one of these policies then they would want to tailor it to their situation to the best of their abilities and that would include a more accurate estimation of the age of new units. It does illustrate the benefits of correctly tailored policies.

Mayo Scenario:

Simulating the Mayo scenario with increased demand shows that *PIPI_direct_Mayo0.221* continues to perform better than default, with the difference being significant across all three metrics. However, *PIPI_direct_Mayo0.223* performs less well than default across all three metrics, though none of the differences are significant (Table *29*). When simulated with decreased demand both of the promising policies (*PIPI_direct_Mayo0.221*, *PIPI_direct_Mayo0.223*) have statistically significant improvement in undiscounted reward and in daily discounted reward ($p<0.001$) compared to default. The difference in discounted reward in each case is positive but not significant (Table 30).

When faced with older age replenishment units (Table 31) both Mayo policies have significantly better undiscounted reward and daily discounted reward ($p<0.001$) than the default policy. *PIPI_direct_Mayo0.221* shows significant improvement in discounted reward ($p = 0.01$) but the improvement in discounted reward for *PIPI_direct_Mayo0.223* is not significant ($p = 0.19$).

Table 29

API Simulation Results: Mayo Increased Demand

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over Default		Discounted Reward Improvement Over Default		Daily Discount^	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.00	0.000%	481	1.1%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_Mayo0.221	3.34	0.008%	378	0.8%	21.4%	9,178	2.6E-03	214	6.4E-03	4,632	1.5E-04
PIPI_direct_Mayo0.223	23.79	0.054%	384	0.9%	20.2%	-8,003	4.1E-01	-194	3.3E-01	-3,847	3.6E-01

Table 30

API Simulation Results: Mayo Decreased Demand

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.00	0.000%	1,260	3.4%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_Mayo0.221	0.41	0.001%	959	2.6%	23.9%	34,403	9.4E-21	110	6.2E-01	13,165	3.5E-13
PIPI_direct_Mayo0.223	1.17	0.003%	658	1.8%	47.8%	68,488	1.3E-24	234	4.1E-01	26,898	3.1E-17

140

Table 31

API Simulation Results: Mayo Older Age

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual	% Demand Unmet	mean annual	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.00	0.000%	1,152	2.8%	0.0%	0	NA	0	NA	0	NA
PIPI_direct_Mayo0.221	0.22	0.001%	913	2.2%	20.8%	27,510	2.0E-13	445	9.7E-03	11,328	2.4E-10
PIPI_direct_Mayo0.223	1.13	0.003%	697	1.7%	39.5%	51,675	3.1E-15	399	1.9E-01	21,712	3.4E-12

5.5.1.4. Application of Mayo policy to Cobre Scenario.

To evaluate the benefit of tailored policies with the same approach as in the previous chapter the best Cobre and Mayo direct API policies were used in simulations of the Cobre scenario. These comparisons used 40 test sequences instead of the usual 20. Plots comparing the improvement due to the use of the Cobre policy are presented in Figure 39 below. The difference between the best Cobre policy (*PIPI_direct_2extension50*) and *PIPI_direct_Mayo0.221* was positive (indicating better performance by the Cobre policy) for all three reward metrics (undiscounted mean improvement: 573, $p = 0.006$; discounted: 31, $p = 0.698$; daily discounted: 186, $p = 0.099$) but only the undiscounted mean improvement was significant.

The best Cobre policy was significantly better than PIPI_direct_Mayo.223 for all three metrics (undiscounted mean improvement: 10132, $p < 0.001$; discounted reward improvement: 1607, $p<0.001$; daily discounted improvement: 4150, $p<0.001$). This may indicate that *PIPI_direct_Mayo0.221* and *PIPI_direct_2extension50* are more similar policies than are *PIPI_direct_Mayo0.223* and *PIPI_direct_2extension50*. Together, these results support the tailoring of policies to hospital characteristics.

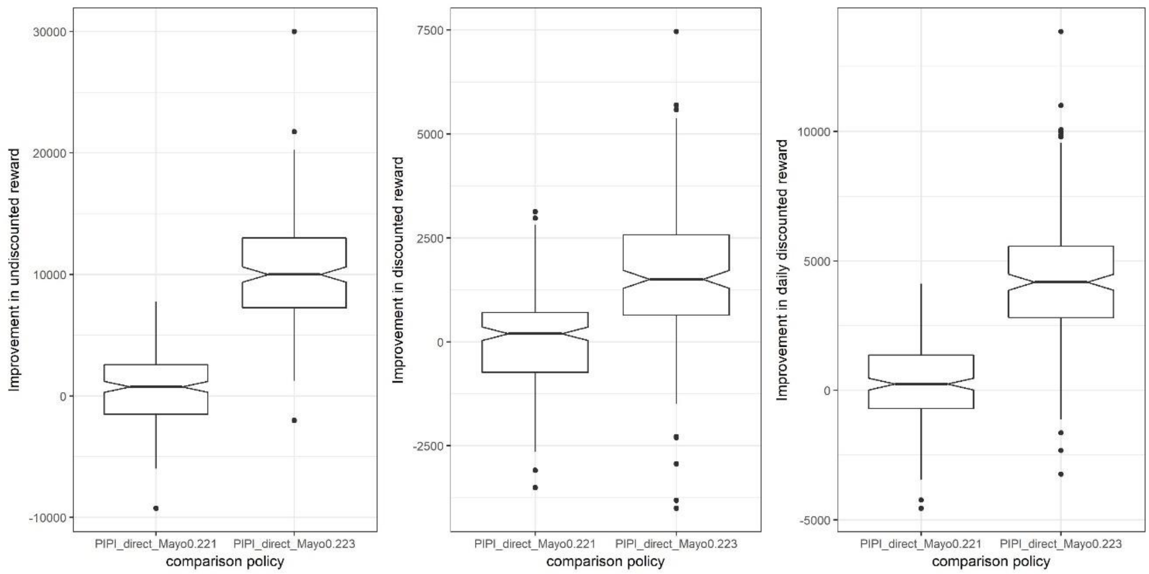

Figure 39. Cobre Policy Improvement over Mayo Policies in Cobre scenario. Positive Value Indicates Cobre policy has higher value.

5.5.1.5. Feature importance for best API model.

Looking directly at the final XGBRegressor model for the best Mayo policy (*PIPI_direct_Mayo0.221*) itself may suggest a path forward for simpler and better models. In these models a large feature vector was used to describe the state. These features were used to fit the XGBClassifier models. They included the blood type of the current patient, the current day of the week, and a multitude of information on the inventory by blood type including the number of units of each age of each blood type (e.g. number of 6 day old O+ units, number of 4 day old A- units etc.), the oldest unit, the 25th, and 75th age percentiles and the total number of units by type (see Appendix F).

Unsurprisingly, the most important feature overall was the blood type of the current patient. Surprisingly, the day of the week was not a significant feature. This would probably be different if ordering were not so frequent.

One of the handcrafted features included was the total number of compatible units available for a patient of each blood type (not the current patient type). For example the total number of compatible units available to a patient with AB+ blood would be the total number of units in inventory; the total number of compatible units available to a patient with A- blood would be the number of A- units plus the number of O- units. This feature was in the top 10 features for 6 of the 8 blood types. The ranking of these features were: A+: 4; A-: 9; B+: 5; B-: 1; AB+: 2; AB-: 3. For O+ the most important feature was the total number of (O+) units in inventory, while total number of units was 3rd most important for O-. The number of 10- and 26-day old units were the most important features for O-. For A+ and B+ the most important feature was the age of the oldest unit.

5.6. Conclusions

Rollout performed better than the API for Cobre but provided no improvement in the high demand case. It may be possible that a longer lookahead would enable gain with rollout for a high demand scenario, but it will come at an impractically steep time penalty. Rollout could be easily implemented for small hospitals and will be more robust to misspecification than a more tailored algorithm.

The high outdate rate seen in the Cobre simulations may be misleading as hospitals of this size typically have aged units collected by the blood provided and distributed for use at hospitals with higher demand. In this way units which were counted as outdated in the simulation would still find use in reality. Information on the specifics of how those rotation decisions are made was not available so that aspect was not able to be included. The simulations show what would happen if rotation was not a factor.

These analyses also provide information that could be used to set different rotation policies, perhaps forgoing rotation altogether or allowing the units to be rotated at an older age.

Assuming the rotation policy is based on some internal age cutoff, the age dependent policies could be applied with that internal cutoff as the target maximum shelf life. This would reduce the number of units reaching the age where they need to be rotated. This is beneficial to the system because rotated units lose useful shelf life during each rotation process. Additionally, the units which were rotated would be younger when provided to the larger hospitals.

For a high demand hospital, the direct API results are very promising. In the Mayo scenario the API policies showed the potential for outdating to be reduced without changing any other policy. The effect on the larger system is hard to judge because of the interconnectedness.

The new policies showed good robustness to misspecification but simulation with facility-specific data is recommended for facilities looking to implement this type of tailored policy.

The new policies used more O blood and more often gave it to non-O patients but also wasted less of the O type blood.

The models presented here are realistic representations, based on actual data, that forgo many of the simplifications found in other inventory models. The use of distributed age of replenishment units and the use of machine learning models is innovative, particularly the use of XGBoost and the use of a categorical model rather than a regression model.

The simulations did not include a crossmatching policy or, as described before, rotation of units from the low-demand hospital. The analysis would be improved by the simultaneous adjustment of replenishment policies. With the involvement of a regional blood center additional benefits could be gained from system-wide coordination of replenishment, allocation and rotation policies. The information about variable importance gained from the XGBoost model can be used to create more user-friendly control policies.

CHAPTER 6

SUMMARY AND CONCLUSIONS

This dissertation addressed the hospital level blood inventory management problem with an eye towards developing more effective policies for issuing blood units to patients. Many of the interacting facets which complicate the allocation of units to patients are captured in the developed policies.

Summary

In Chapter 3 simulation was used to examine the interaction of forecast accuracy with age of new units, replenishment schedule, lead time, daily demand and safety stock. A forecast derived from a hospital's procedure schedule was modeled. It simulated a forecast based on procedure schedule. It was shown that increased age of replenishment units and less frequent replenishment can increase wastage even under a perfect forecast. Consideration of two models with significantly different daily demand showed that the interaction of safety stock with forecasting error variance leads to much more wastage in the low demand case with an unbiased forecast.

In Chapter 4 Markov decision process models were used to find optimal issuing policies for red blood cell units under emergency conditions where there is no replenishment arriving. This model did not consider perishability because of the short time horizon of the emergency. The resulting policies are stock level-dependent allocation policies that make optimal decisions of what blood type to give to patients based on current inventory. Simple Monte Carlo simulation analysis showed that the new policies reduced outdating compared to the default policy. Additionally, it showed the benefits of having location specific policies. In a discrete event simulation-based case study of a mass casualty event the new allocation policies were shown to reduce the number of unserved patients in most scenarios. In a pandemic influenza simulation-based case study the new policies were shown to reduce the number of unserved patients and reduce the number of outdated units in all of the considered cases. The model in Chapter 4 only included ABO type and expansion to include ABO/Rh type may prove beneficial.

In Chapter 5 approximate policy iteration was used to find improved issuing policies under normal operations. This model included perishability and replenishment. It also included all eight ABO/Rh blood types. The inclusion of this information led to an explosion in problem size, necessitating an approximate dynamic programming approach. Simple rollout showed improvement over the default policy for the low demand case but not in the high demand case. Policies found using model-free policy iteration using rollout to improve the value approximation and a XGBClassifier model to approximate the policy did show improvement over the default policy for both the high and low demand cases.

The models in Chapter 5 used complex age distributions for replenishment units which can improve the accuracy of model results. Both Chapter 4 and Chapter 5 make recommendations for issuing a single unit at a time. Many existing models assume batch demand or batch allocation.

Chapter 4 and Chapter 5 both utilized substitution to improve the inventory performance. In Chapter 4 optimal choices about blood unit type to issue increases the number of patients that can be served consecutively. In Chapter 5 the policies are age and stock level dependent. Substitution may be recommended even if a direct match is present in order to prevent units outdating.

The common themes of these chapters are the focus on tailoring analysis to the specific characteristics and context of a facility as well as maintaining fidelity to the real world through use of historical data and simulation. For example, Chapters 3 and 5 both consider the size of the facility specifically while Chapters 4 and 5 both analyze multiple simulated patient blood type distributions. Each chapter uses simulation to capture the complex interactions present in the blood supply chain and each chapter utilizes historical data whenever possible, to avoid relying on inaccurate but simple distributional assumptions.

Dissertation Research Contributions

The specific contributions of this dissertation can be described as:

- Quantifying the benefits that may accrue through demand forecasting of blood needs at the hospital level and showing that the benefits are highly dependent on hospital characteristics.
- Creating relative stock level-dependent policies for the issuing of red blood cell units to patients under emergency conditions considering ABO blood typing and utilizing proactive blood-type substitution; and then showing that these policies outperform the fixed priority default policy. These results show that customized policies outperform default policies and stock-dependent policies which are not tailored for the specific location.

- Creating detailed simulations that evaluate the performance of stock level-dependent policies with proactive blood-type based substitution in response to a mass casualty event and under a pandemic influenza scenario and show that the stock level-dependent policies reduce outdating under each scenario considered as well as reduce the unmet demand in all pandemic scenarios and most mass casualty event scenarios considered.
- Showing via simulation that the MDP-derived policies improve outdating and unmet demand under normal operating conditions as well as under emergency conditions. A method is introduced for extending the policies developed on a reduced state space to a larger state space by using an approximating hyperplane.
- Developing the first age- and stock level-dependent allocation policy with proactive substitution between blood types for improved allocation of perishable red blood cell inventory. This model incorporates ABO/Rh typing and compatibility, the lifetime of the product, set replenishment policies, varying age of new replenishment units and active substitution. Previous models have assumed batch demand, ignored unit blood type and/or substitution or permitted only priority-based substitution.
- Utilizing approximate dynamic programming with machine learning and approximation in policy space to develop the proactively substituting age- and stock level-dependent policies. The resulting policies are defined in easily implementable machine learning models.

- Demonstrating that the policies based on approximate dynamic programming models outperform the default policy. These models have fewer assumptions compared to existing models and policies and work on problems of any realistic size.

Possible Future Research Directions

Future work on the forecasting of demand based on hospital procedure schedule should increase the range of hospital sizes considered. It should also look for better ways to place orders in a scenario where some of the demand is forecastable (for example elective procedures) while the other demand is random and cannot be forecast from the procedure schedule.

Similarly, it would be beneficial to use additional data sets from facilities of a variety of sizes and configurations to use the methods developed in this dissertation to determine new allocation policies as in Chapters 4 and 5. More study of small hospitals or hospitals in regions without steady supply is especially warranted.

Chapter 5 did not include a lead time in the models and in both chapters, it was assumed that each patient required a single unit each. These assumptions are reasonable but multiple units per person could be included and studied.

For hospitals to take action based on the results of this dissertation, it would be recommended that they first perform simulations with their own input data and derive their own allocation policies specific to their daily demand, age of replenishment units from their blood suppliers and blood type distribution of their patient population.

For a region-wide application, the inclusion of information on the regional blood center would be useful as well as inventory management policies within the region. It is possible that allocation and replenishment policies could be coordinated across a region along with the blood provider's rotation policies to optimize inventory usage across the region.

REFERENCES

AABB. (2008). *Disaster Operations Handbook — Hospital Supplement Coordinating the Nation's Blood Supply*. Retrieved from www. aabb.org/Content/Programs_and_Services/Disaster_Response/

AABB Interorganizational Task Force on Pandemic Influenza and the Blood Supply. (2009). *Pandemic Influenza Issues Outline*. Retrieved from http://www.aabb.org/programs/disasterresponse/Pages/taskforcepibs.aspx

Abbasi, B., & Hosseinifard, S. Z. (2014). On the Issuing Policies for Perishable Items such as Red Blood Cells and Platelets in Blood Service. *Decision Sciences*, *45*(5), 995–1020. https://doi.org/10.1111/deci.12092

Abdulwahab, U., & Wahab, M. I. M. (2014). Approximate dynamic programming modeling for a typical blood platelet bank. *Computers and Industrial Engineering*, *78*, 259–270. https://doi.org/10.1016/j.cie.2014.07.017

Adulyasak, Y., Cordeau, J. F., & Jans, R. (2015). Benders decomposition for production routing under demand uncertainty. *Operations Research*, *63*(4), 851–867. https://doi.org/10.1287/opre.2015.1401

Agnew, T. (2012, October 31). Red Cross : We need blood. *Suffold News Herald*, pp. 8–11. Retrieved from https://www.suffolknewsherald.com/2012/10/31/red-cross-we-need-blood/

Alfonso, E., Xie, X., Augusto, V., & Garraud, O. (2012). Modeling and simulation of blood collection systems. *Health Care Management Science*, *15*(1), 63–78. https://doi.org/10.1007/s10729-011-9181-8

American Red Cross. (2016). Blood Facts and Statistics Facts about blood needs Facts about American Red Cross Blood Services. Retrieved December 7, 2016, from www.redcrossbloo.org/print/leran-about-blood/blood-facts-and-statistics

American Red Cross. (2020). *American Red Cross faces severe blood shortage as coronavirus outbreak threatens availability of nation's supply*. Retrieved from https://www.redcross.org/about-us/news-and-events/press-release/2020/american-red-cross-faces-severe-blood-shortage-as-coronavirus-outbreak-threatens-availability-of-nations-supply.html

Arellanos-Soto, D., B.-d. l. Cruz, V., Mendoza-Tavera, N., Ramos-Jiménez, J., Cázares-Taméz, R., Ortega-Soto, a., & Rivas-Estilla, a. M. (2015). Constant risk of dengue virus infection by blood transfusion in an endemic area in Mexico. *Transfusion Medicine*, (April), n/a-n/a. https://doi.org/10.1111/tme.12198

Armstrong, B., Wilkinson, R., & Smart, E. (2008). Compatibility testing. *ISBT Science Series*, *3*, 197–215.

Asamoah-Akuoko, L., Hassall, O. W., Bates, I., & Ullum, H. (2017). Blood donors' perceptions, motivators and deterrents in Sub-Saharan Africa – a scoping review of evidence. *British Journal of Haematology*, *177*(6), 864–877. https://doi.org/10.1111/bjh.14588

Astaraky, D., & Patrick, J. (2015). A simulation based approximate dynamic programming approach to multi-class, multi-resource surgical scheduling. *European Journal of Operational Research*, *245*(1), 309–319. https://doi.org/10.1016/j.ejor.2015.02.032

Baesler, F., Nemeth, M., Martínez, C., & Bastías, A. (2014). Analysis of inventory strategies for blood components in a regional blood center using process simulation. *Transfusion*, *54*(2), 323–330. https://doi.org/10.1111/trf.12287

Bakker, M., Riezebos, J., & Teunter, R. H. (2012). Review of inventory systems with deterioration since 2001. *European Journal of Operational Research*, *221*(2), 275–284. https://doi.org/10.1016/j.ejor.2012.03.004

Bekiempis, V. (2020, March 21). "This could kill our patients": US faces blood shortage as donation sites shutter. *The Gua*. Retrieved from http://www.theguardian.com/us-news/2020/mar/21/us-blood-donations-shortage-coronavirus

Beliën, J., & Forcé, H. (2012). Supply chain management of blood products: A literature review. *European Journal of Operational Research*, *217*(1), 1–16. https://doi.org/10.1016/j.ejor.2011.05.026

Bertazzi, L., Bosco, A., & Laganà, D. (2015). Managing stochastic demand in an Inventory Routing Problem with transportation procurement. *Omega (United Kingdom)*, *56*, 112–121. https://doi.org/10.1016/j.omega.2014.09.010

Bertsekas, D. (2011). Approximate policy iteration: A survey and some new methods. *Journal of Control Theory and Applications*, *9*(3), 310–335. https://doi.org/10.1007/s11768-011-1005-3

Bertsekas, D. P. (2012). *Dynamic Programming and Optimal Control* (4th ed.). Nasua: Athena Scientific.

Bertsekas, D. P. (2018). *Feature-Based Aggregation and Deep Reinforcement Learning: A Survey and Some New Implementations*. (April), 1–52. https://doi.org/arXiv:1804.04577v2

Bertsekas, D. P. (2019a). 2.4 Rollout. In *Reinforcement Learning and Optimal Control* (1st ed., pp. 90--1). Belmont, Massachusetts: Athena Scientific.

Bertsekas, D. P. (2019b). *Reinforcement Learning and Optimal Control* (1st ed.). Belmont, Massachusetts: Athena Scientific.

Bhattacharya, S., Badyal, S., Wheeler, T., Gil, S., & Bertsekas, D. (2020). Reinforcement Learning for POMDP: Partitioned Rollout and Policy Iteration with Application to Autonomous Sequential Repair Problems. *IEEE Robotics and Automation Letters*, *5*(3), 3967–3974. https://doi.org/10.1109/LRA.2020.2978451

Blake, J., & Hardy, M. (2013). Using simulation to evaluate a blood supply network in the Canadian maritime provinces. *Journal of Enterprise Information Management*, *26*(1), 119–134. https://doi.org/10.1108/17410391311289587

Bloch, E. M., Cohn, C., Bruhn, R., Hirschler, N., & Nguyen, K. A. (2014). A cross-sectional pilot study of blood utilization in 27 hospitals in Northern California. *American Journal of Clinical Pathology*, *142*(4), 498–505. https://doi.org/10.1309/AJCP8WFIQ0JRCSIR

Boseley, S. (2020). WHO declares coronavirus pandemic. *The Guardian*. Retrieved from https://www.theguardian.com/world/2020/mar/11/who-declares-coronavirus-pandemic

Britten, R. (1932). The Incidence of Epidemic Influenza, 1918-1919. In *Public Health Weekly Reports* (Vol. 47).

Broekmeulen, R. A. C. M., & van Donselaar, K. H. (2009). A heuristic to manage perishable inventory with batch ordering, positive lead-times, and time-varying demand. *Computers & Operations Research*, *36*(11), 3013–3018. https://doi.org/10.1016/j.cor.2009.01.017

Brown, J. B., Sperry, J. L., Fombona, A., Billiar, T. R., Peitzman, A. B., & Guyette, F. X. (2015). *Pre-Trauma Center Red Blood Cell Transfusion Is Associated with Improved Early Outcomes in Air Medical Trauma Patients*. 797–808. https://doi.org/10.1016/j.jamcollsurg.2015.01.006

Cattani, K. D., Jacobs, F. R., & Schoenfelder, J. (2011). Common inventory modeling assumptions that fall short: Arborescent networks, Poisson demand, and single-echelon approximations. *Journal of Operations Management*, *29*(5), 488–499. https://doi.org/10.1016/j.jom.2010.11.008

Cerveny, R. P. (1980). An Application of Warehouse Location Techniques to Bloodmobile Operations An Application of Warehouse Location Techniques to Bloodmobile Operations *. *Interfaces*, *10*(6), 88–96.

Chandra, T., & Gupta, A. (2012). Frequency of ABO & Rhesus blood groups in blood donors in Punjab. *Asian Journal of Transfusion Scienceournal of Medical Sciences*, *6*(1), 52–53. https://doi.org/10.4103/0973-6247.95057

Chapman, J. F., Hyam, C., & Hick, R. (2004). Blood inventory management. *Vox Sanguinis*, *87 Suppl 2*, 143–145. https://doi.org/10.1111/j.1741-6892.2004.00472.x

Chen, T., & Guestrin, C. (2016). XGBoost : A Scalable Tree Boosting System. *KDD*, 785–794.

Chen, X., Hao, G., Li, X., & Yiu, K. F. C. (2012). The impact of demand variability and transshipment on vendor's distribution policies under vendor managed inventory strategy. *International Journal of Production Economics*, *139*(1), 42–48. https://doi.org/10.1016/j.ijpe.2011.05.005

Cheng, C. K., Trethewey, D., & Sadek, I. (2010). Comprehensive survey of red blood cell unit life cycle at a large teaching institution in eastern Canada. *Transfusion*, *50*(1), 160–165. https://doi.org/10.1111/j.1537-2995.2009.02375.x

Civelek, I., Karaesmen, I., & Scheller-Wolf, A. (2015). Blood platelet inventory management with protection levels. *European Journal of Operational Research*, *243*(3), 826–838. https://doi.org/10.1016/j.ejor.2015.01.023

Coelho, L. C., & Laporte, G. (2014). Optimal joint replenishment, delivery and inventory management policies for perishable products. *Computers and Operations Research*, *47*, 42–52. https://doi.org/10.1016/j.cor.2014.01.013

Costantino, F., Di Gravio, G., Shaban, A., & Tronci, M. (2014). The impact of information sharing and inventory control coordination on supply chain performances. *Computers and Industrial Engineering*, *76*, 292–306. https://doi.org/10.1016/j.cie.2014.08.006

Demario, V. M., Sikorski, R. A., Efron, D. T., Serbanescu, M. A., Buchanan, R. M., Wang, E. J., … Frank, S. M. (2018). Blood utilization and mortality in victims of gun violence. *Transfusion*, *58*(October), 2326–2334. https://doi.org/10.1111/trf.14925

Deniz, B., Karaesmen, I., & Scheller-Wolf, A. (2010). Managing Perishables with Substitution: Inventory Issuance and Replenishment Heuristics. *Manufacturing & Service Operations Management*, *12*(2), 319–329. https://doi.org/10.1287/msom.1090.0276

Dillon, M., Oliveira, F., & Abbasi, B. (2017). A two-stage stochastic programming model for inventory management in the blood supply chain. *International Journal of Production Economics*, *187*(May 2016), 27–41. https://doi.org/10.1016/j.ijpe.2017.02.006

Duan, Q., & Liao, T. W. (2013). A new age-based replenishment policy for supply chain inventory optimization of highly perishable products. *International Journal of Production Economics*, *145*(2), 658–671. https://doi.org/10.1016/j.ijpe.2013.05.020

Duan, Q., & Liao, T. W. (2014). Optimization of blood supply chain with shortened shelf lives and ABO compatibility. *International Journal of Production Economics*, *153*, 113–129. https://doi.org/10.1016/j.ijpe.2014.02.012

Dumkrieger, G., Huschka, T. R. ;, & Stubbs, J. R. (2014). Modeling the Effect of Shorter Shelf life on Red Blood Cells. *Proceedings of the 2014 Winter Simulation Conference*, 1293–1304.

Fildes, R., & Kingsman, B. (2011). Incorporating demand uncertainty and forecast error in supply chain planning models⋆. *Journal of the Operational Research Society*, *62*(3), 483–500. https://doi.org/10.1057/jors.2010.40

Flavelle, C. (2020, June 2). Red Cross Warns of a “Staggering” Drop in Blood Supplies. *The New York Times*.

Flegel, W. a, Natanson, C., & Klein, H. G. (2014). Does prolonged storage of red blood cells cause harm? *British Journal of Haematology*, *165*(1), 3–16. https://doi.org/10.1111/bjh.12747

Flommersfeld, S., Mand, C., Kuhne, C. A., Bein, G., Ruchholtz, S., & Sachs, U. J. (2018). Unmatched Type O RhD + Red Blood Cells in Multiple Injured Patients. *Transfusion Medicine and Hemotherapy*, *45*, 158–161. https://doi.org/10.1159/000485388

Fortsch, S. M., & Khapalova, E. A. (2016). Reducing uncertainty in demand for blood. *Operations Research for Health Care*, *9*, 16–28. https://doi.org/10.1016/j.orhc.2016.02.002

Garcia, B. (2016, June 8). United Blood Services appeals to donors amidst blood shortage. *The Monitor*, pp. 11–13. Retrieved from http://www.themonitor.com/news/local/united?blood?services?appeals?to?donors?amidst?blood? shortage/article_7b07070e?2ddf?11e6?a340?77dac6647482.html

Germain, R., Claycomb, C., & Dröge, C. (2008). Supply chain variability, organizational structure, and performance: The moderating effect of demand unpredictability. *Journal of Operations Management*, *26*(5), 557–570. https://doi.org/10.1016/j.jom.2007.10.002

Glasgow, S., Davenport, R., Perkins, Z., Tai, N., & Brohi, K. (2013). A comprehensive review of blood product use in civilian mass casualty events. *The Journal of Trauma and Acute Care Surgery*, *75*(3), 468–474. https://doi.org/10.1097/TA.0b013e318298efb9

Goh, C.-H., Greenberg, B. S., & Matsuo, H. (1993). Two-Stage Perishable Inventory Models. *Management Science*, *39*(5), 633–649. https://doi.org/10.1287/mnsc.39.5.633

Goodson, J. C., Ohlmann, J. W., & Thomas, B. W. (2013). Rollout policies for dynamic solutions to the multivehicle routing problem with stochastic demand and duration limits. *Operations Research*, *61*(1), 138–154. https://doi.org/10.1287/opre.1120.1127

Goyal, S. K., & Giri, B. C. (2001). Recent trends in modeling of deteriorating inventory. *European Journal of Operational Research*, *134*(1), 1–16. https://doi.org/10.1016/S0377-2217(00)00248-4

Guan, L., Tian, X., Gombar, S., Zemek, A. J., Krishnan, G., Scott, R., … Pham, T. D. (2017). *Big data modeling to predict platelet usage and minimize wastage in a tertiary care system*. *114*(43). https://doi.org/10.1073/pnas.1714097114

Gürler, Ü., & Özkaya, B. Y. (2008). Analysis of the (s, S) policy for perishables with a random shelf life. *IIE Transactions (Institute of Industrial Engineers)*, *40*(8), 759–781. https://doi.org/10.1080/07408170701730792

Haijema, Rene. (2011). Optimal Issuing of Perishables with a Short Fixed Shelf Life. In J. W. Bose, C. Jahn, H. Hu, X. Shi, R. Stahlbock, & S. Voss (Eds.), *Internation Conference on Computational Logistics* (pp. 160–169). Verlag Berlin Heidelberg: Springer.

Haijema, Rene. (2014). Optimal ordering, issuance and disposal policies for inventory management of perishable products. *International Journal of Production Economics*, *157*(1), 158–169. https://doi.org/10.1016/j.ijpe.2014.06.014

Haijema, René. (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. https://doi.org/10.1016/j.ijpe.2011.05.021

Haijema, René, van der Wal, J., & van Dijk, N. M. (2007). Blood platelet production: Optimization by dynamic programming and simulation. *Computers & Operations Research*, *34*(3), 760–779. https://doi.org/10.1016/j.cor.2005.03.023

Hawkins, S. A., & Gillett, J. (2015). Reconstructing Aboriginal Identity and the Meaning of the Body in Blood Donation Decisions. *Journal of Religion and Health*, *54*(4), 1286–1301. https://doi.org/10.1007/s10943-014-9879-2

Hayn, D., Kreiner, K., Kastner, P., Breznik, N., Hofmann, A., Gombotz, H., & Schreier, G. (2016). Data driven methods for predicting blood transfusion needs in elective surgery. *Studies in Health Technology and Informatics*, *223*, 9–16. https://doi.org/10.3233/978-1-61499-645-3-9

Hemmelmayr, V., Doerner, K. F., Hartl, R. F., & Savelsbergh, M. W. P. (2009). Delivery strategies for blood products supplies. *OR Spectrum*, *31*(4), 707–725. https://doi.org/10.1007/s00291-008-0134-7

Hemmelmayr, V., Doerner, K. F., Hartl, R. F., & Savelsbergh, M. W. P. (2010). Vendor managed inventory for environments with stochastic product usage. *European Journal of Operational Research*, *202*(3), 686–695. https://doi.org/10.1016/j.ejor.2009.06.003

Igarashi, T., Fujita, H., Asaka, H., Takada, Y., Ametani, R., Naya, I., … Kasai, A. (2018). Patient rescue and blood utilization in the Ogasawara blood rotation system. *Transfusion*, *58*(March), 788–794. https://doi.org/10.1111/trf.14444

Inghilleri, G. (2010). Prediction of transfusion requirements in surgical patients: A review. *Transfusion Alternatives in Transfusion Medicine*, *11*(1), 10–19. https://doi.org/10.1111/j.1778-428X.2009.01118.x

Ishii, H., & Nose, T. (1996). Perishable inventory control with two types of customers and different selling prices under the warehouse capacity constraint. *International Journal of Production Economics*, *5273*(95).

Jagannathan, R., & Sen, T. (1991). Storing Crossmatched Blood : A Perishable Inventory Model with Prior Allocation. *Management Science*, *37*(3), 251–266.

Janssen, L., Claus, T., & Sauer, J. (2016). Literature review of deteriorating inventory models by key topics from 2012 to 2015. *International Journal of Production Economics*, *182*, 86–112. https://doi.org/10.1016/j.ijpe.2016.08.019

Janssen, L., Sauer, J., Claus, T., & Nehls, U. (2018). Development and simulation analysis of a new perishable inventory model with a closing days constraint under non-stationary stochastic demand. *Computers and Industrial Engineering*, *118*(November 2016), 9–22. https://doi.org/10.1016/j.cie.2018.02.016

Jester, B., Uyeki, T., & Jernigan, D. (2018). Readiness for Responding to a Severe Pandemic 100 Years After 1918. *American Journal of Epidemiology*, *187*(12), 2596–2602. https://doi.org/10.1093/aje/kwy165

Johansen, S. G., & Thorstenson, A. (2014). Emergency orders in the periodic-review inventory system with fixed ordering costs and compound Poisson demand. *International Journal of Production Economics*, *157*(1), 147–157. https://doi.org/10.1016/j.ijpe.2014.07.013

Kamp, C., Heiden, M., Henseler, O., & Seitz, R. (2010). Management of blood supplies during an influenza pandemic. *Transfusion*, *50*(1), 231–239. https://doi.org/10.1111/j.1537-2995.2009.02498.x

Katsaliaki, K., & Brailsford, S. C. (2006). Using simulation to improve the blood supply chain. *Journal of the Operational Research Society*, *58*(2), 219–227. https://doi.org/10.1057/palgrave.jors.2602195

Kendall, K. E., & Lee, S. M. (1980). Formulating Blood Rotation Policies with Multiple Objectives. *Management Science*, *26*(11), 1145–1157. https://doi.org/10.1287/mnsc.26.11.1145

Kok, A. Gurhan; Fisher, M. L. (2007). Demand Estimation and Assortment Optimization Under Substitution : Methodology and Application. *Operations Research*, *55*(6), 1001–1021.

Kopach, R. (2004). *Finding the optimal switchover level between two levels of demand in a Blood Bank*.

Kopach, R., Balcıoğlu, B., & Carter, M. (2008). Tutorial on constructing a red blood cell inventory management system with two demand rates. *European Journal of Operational Research*, *185*(3), 1051–1059. https://doi.org/10.1016/j.ejor.2006.01.051

Kuehnert, M. J. (2016). Screening of blood donations for Zika virus infection — Puerto Rico, April 3-June 11, 2016. *MMWR. Morbidity and Mortality Weekly Report*, *65*(24), 627–628. Retrieved from https://www.cdc.gov/mmwr/volumes/65/wr/pdfs/mm6524e2.pdf

Lardieri, A. (2020, July 7). Arizona Coronavirus Death Toll Mounts as Hospitals Near Capacity. *U.S. News & World Report*.

Laughland, O. (2020, July 7). Fauci : US is ‘ still knee-deep in first wave ’ of pandemic as it passes Register for free and continue reading. *The Gua*, p. 130.

Lestari, F., Anwar, U., Nugraha, N., & Azwar, B. (2017). Forecasting demand in blood supply chain (case study on blood transfusion unit). *Lecture Notes in Engineering and Computer Science*, *2230*, 5–8.

Mahajan, S., & van Ryzin, G. (2001). Inventory Competition Under Dynamic Consumer Choice. *Operations Research*, *49*(5), 646–657. https://doi.org/10.1287/opre.49.5.646.10603

Marks, P. (2020). *FDA Statement Coronavirus (COVID-19) Update : Blood Donations*.

Mcquilten, Z. K., Mercer, G., Phillips, L., Luangwilai, T., Brown, R., Ozolins, I., … Wood, E. M. (2014). A dynamic mathematical model of red blood cell clinical demand to assess the impact of prolonged blood shortages and transfusion restriction policies. *Transfusion*, *54*(October), 2705–2715. https://doi.org/10.1111/trf.12525

Meissner, J., & Senicheva, O. V. (2018). Approximate dynamic programming for lateral transshipment problems in multi-location inventory systems. *European Journal of Operational Research*, *265*(1), 49–64. https://doi.org/10.1016/j.ejor.2017.06.049

Moin, N. H., & Halim, H. Z. A. (2018). Solving inventory routing problem with stochastic demand. *AIP Conference Proceedings*, *1974*(June). https://doi.org/10.1063/1.5041635

Mustafee, N., Taylor, S. J. E., Katsaliaki, K., & Brailsford, S. (2009). Facilitating the Analysis of a UK National Blood Service Supply Chain Using Distributed Simulation. *Simulation*, *85*(2), 113–128. https://doi.org/10.1177/0037549708100530

Nahmias, S. (1977). On Ordering Perishable Inventory when Both Demand and Lifetime are Random. *Management Science*, *24*(1), 82.

Nahmias, S. (1982). Perishable inventory theory: a review. *Operations Research*, *30*(4), 680–708. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10298625

Nahmias, S. (2011). *Perishable Inventory Systems*. Springer.

Nahmias, S., & Pierskalla, W. P. (1973). Optimal ordering policies for a product that perishes in two periods subject to stochastic demand. *Naval Research Logistics Quarterly*, *20*(2), 207–229. https://doi.org/10.1002/nav.3800200202

Netessine, S., & Rudi, N. (2003). Centralized and Competitive Inventory Models with Demand Substitution. *Operations Research*, *51*(2), 329–335.

Osorio, A. F., Brailsford, S. C., & Smith, H. K. (2015). A structured review of quantitative models in the blood supply chain: a taxonomic framework for decision-making. *International Journal of Production Research*, *1*(1), 1–22. https://doi.org/10.1080/00207543.2015.1005766

Parlar, M., Perry, D., & Stadje, W. (2010). FIFO Versus LIFO Issuing Policies for Stochastic Perishable Inventory Systems. *Methodology and Computing in Applied Probability*, *13*(2), 405–417. https://doi.org/10.1007/s11009-009-9162-2

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., … Duchesnay, É. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Perera, G., Hyam, C., Taylor, C., & Chapman, J. F. (2009). Hospital Blood Inventory Practice: The factors affecting stock level and wastage. *Transfusion Medicine*, *19*(2), 99–104. https://doi.org/10.1111/j.1365-3148.2009.00914.x

Pierskalla, W. P. (2004). Supply Chain Management of Blood Banks. In M. L. Brandeau & F. Sainfort (Eds.), *Operations Research and Health Care* (pp. 103–145). https://doi.org/10.1007/1-4020-8066-2_5

Pierskalla, W., & Roach, C. (1972). Optimal Issuing Policies for Perishable Inventory. *Management Science*, *18*(11), 603.

Powell Elizabeth, K., Hinckley William, R., Gottula Adam, W., Hart Kimberly, J., Lindsell Christopher, T., & Mcmullan Jason, T. (2016). Shorter times to packed red blood cell transfusion are associated with decreased risk of death in traumatically injured patients. *Journal of Trauma and Acute Care Surgery*, *81*(3), 458–462. https://doi.org/10.1097/TA.0000000000001078

Powell, E. K., Hinckley, W. R., Gottula, A. W., Hart, K. J., Lindsell, C. T., & Mcmullan, J. T. (2016). Shorter times to packed red blood cell transfusion are associated with decreased risk of death in traumatically injured patients. *Journal of Trauma and Acute Care Surgery*, *81*(3), 458–462. https://doi.org/10.1097/TA.0000000000001078

Powell, W. (2011). *Approximate Dynamic Programming*. Hoboken, New Jersey: John Wiley & Sons, Inc.

Powell, W. B. (2010). Merging AI and OR to solve high-dimensional stochastic optimization problems using approximate dynamic programming. *INFORMS Journal on Computing*, *22*(1), 2–17. https://doi.org/10.1287/ijoc.1090.0349

Prastacos, G. P. (1984). Blood inventory management: An Overview of Theory and Practice. *Management Science*, *30*(7), 777–800. https://doi.org/10.1111/j.1741-6892.2004.00472.x

Puranam, K., Novak, D. C., Lucas, M. T., & Fung, M. (2017). Managing blood inventory with multiple independent sources of supply. *European Journal of Operational Research*, *259*(2), 500–511. https://doi.org/10.1016/j.ejor.2016.11.005

Reuters News Services. (2000, September 21). U. S. faces critical blood shortage. *Deseret News*. Retrieved from http://www.deseretnews.com/article/783849/US?faces?critical?blood?shortage.html ?pg=all

Rex-Kiss, B. ., Szabó, L. ., Szabó, S., & Hartmann, É. (1973). ABO , MN , Rh Blood Groups , Hp Types and Hp Level , Gm (1) Factor Investigations on the Gypsy Population of Hungary. *Human Biology*, *45*(1), 41–61.

Rockwell Automation. (2016). *Arena Simulation Software*.

Sainathan, A. (2013). Pricing and replenishment of competing perishable product variants under dynamic demand substitution. *Production and Operations Management*, *22*(5), 1157–1181. https://doi.org/10.1111/poms.12004

Sapountzis, C. (1989). Allocating Blood to Hospitals. *The Journal of the Operational Research Society*, *40*(5), 443–449.

Sarhangian, V., Abouee-Mehrizi, H., Baron, O., & Berman, O. (2017). Threshold-Based Allocation Policies for Inventory Management of Red Blood Cells. *Manufacturing & Service Operations Management*, (April 2019). https://doi.org/10.1287/msom.2017.0650

SAS Institute Inc. (2018). *JMP Pro*. Cary, NC: SAS Institute Inc.

Schantz-Dunn, J., & Nour, N. M. (2011). The use of blood in obstetrics and gynecology in the developing world. *Reviews in Obstetrics and Gynecology*, *4*(2), 86–91. https://doi.org/10.3909/riog0160

Schmidt, P. (2002). Blood and Disaster - Supply and Demand. *New England Journal of Medicine*, *346*(8), 617–620.

Schütz, H. J., & Kolisch, R. (2012). Approximate dynamic programming for capacity allocation in the service industry. *European Journal of Operational Research*, *218*(1), 239–250. https://doi.org/10.1016/j.ejor.2011.09.007

Semenza, J. C., & Domanović, D. (2013). Blood supply under threat. *Nature Climate Change*, *3*(5), 432–435. https://doi.org/10.1038/nclimate1867

Shah, M. D., Goldstein, D. P., McCluskey, S. A., Miles, B. A., Hofer, S. O., Brown, D. H., … Gilbert, R. W. (2010). Blood Transfusion Prediction in Patients Undergoing Major Head and Neck Surgery With Free-Flap Reconstruction. *Archives of Otolaryngology - Head & Neck Surgery*, *136*(12), 10–15.

Sharrar, R. G. (1969). National influenza experience in the USA, 1968-69. *Bulletin of the World Health Organization*, *41*(3), 361–366.

Shen, Z., Dessouky, M., & Ordonez, F. (2011). Perishable inventory management system with a minimum volume constraint. *Journal of the Operational Research Society*, *62*(12), 2063–2082. https://doi.org/10.1057/jors.2010.181

Somani, A., Ye, N., Hsu, D., & Lee, W. S. (2013). DESPOT: Online POMDP planning with regularization. *Advances in Neural Information Processing Systems*, *58*, 231–266.

Stanger, S. H. W., Wilding, R., Yates, N., & Cotton, S. (2012). What drives perishable inventory management performance? Lessons learnt from the UK blood supply chain. *Supply Chain Management: An International Journal*, *17*(2), 107–123. https://doi.org/10.1108/13598541211212861

Stanworth, S. J., Davenport, R., Curry, N., Seeney, F., Eaglestone, S., Edwards, A., … Brohi, K. (2016). Mortality from trauma haemorrhage and opportunities for improvement in transfusion practice. *BJS*, *103*, 357–365. https://doi.org/10.1002/bjs.10052

Steiner, M. E., Ness, P. M., Assmann, S. F., Triulzi, D. J., Sloan, S. R., Delaney, M., … Stowell, C. P. (2015). Effects of red-cell storage duration on patients undergoing cardiac surgery. *New England Journal of Medicine*, *372*(15), 1419–1429. https://doi.org/10.1056/NEJMoa1414219

Stephan, R. B. (2008). *Letter to Emergency Management Agency from Department of Homeland Security*. Retrieved from http://www.aabb.org/programs/disasterresponse/Documents/letterdhs.pdf

Suárez, I. M. B., Fernández-Montoya, A., Fernández, A. R., López-Berrio, A., & Cillero-Peñuela, M. (2004). How regular blood donors explain their behavior. *Transfusion*, *44*(10), 1441–1446. https://doi.org/10.1111/j.1537-2995.2004.04102.x

Teo, D., Ng, L. C., & Lam, S. (2009). Is dengue a threat to the blood supply? *Transfusion Medicine*, *19*(2), 66–77. https://doi.org/10.1111/j.1365-3148.2009.00916.x

Tesauro, G.J.; Galperin, G. R. (1996). On-Line Policy Improvement Using Monte Carlo Tree Search. *NIPS*.

Titmuss, R. M. (1971). Points of View WHY GIVE TO STRANGERS ? *The Lancet*, 123–125.

Tounkara, a., Diakite, M., Noumsi, G. T., Sarro, Y. D. S., Siddiqui, S., & Parta, M. (2009). Retrospective surveillance of HIV prevalence in blood donors can help in the selection of the best social group for blood donation in Mali. *Transfusion Medicine*, *19*(5), 252–259. https://doi.org/10.1111/j.1365-3148.2009.00950.x

Tserenpuntsag, B., Nelson, K., Lamjav, O., Triner, W., Smith, P., Kacica, M., & McNutt, L. A. (2010). Prevalence of and risk factors for hepatitis B and C infection among Mongolian blood donors. *Transfusion*, *50*(1), 92–99. https://doi.org/10.1111/j.1537-2995.2009.02387.x

Velindre NHS Trust. (2008). *Integrated Medium Term Plan 2018-2021*.

Whitaker, B. I. (2011). *The 2011 National Blood Collection and Utilization Survey Report*.

Wong, C. Y., Boon-itt, S., & Wong, C. W. Y. (2011). The contingency effects of environmental uncertainty on the relationship between supply chain integration and operational performance. *Journal of Operations Management*, *29*(6), 604–615. https://doi.org/10.1016/j.jom.2011.01.003

World Health Organization. (2007). *Mass Casualty Management Systems Strategies and guidelines*.

World Health Organization. (2011). *Maintaining a Safe and Adequate Blood Supply during Pandemic Influenza Guidelines for Blood Transfusion Services*. 1–8.

Xie, S. R., Bish, D. R., Bish, E. K., Slonim, A. D., & Stramer, S. L. (2012). Safety and waste considerations in donated blood screening. *European Journal of Operational Research*, *217*(3), 619–632. https://doi.org/10.1016/j.ejor.2011.09.045

Yazdanbakhsh, K., & Nandi, V. (2016). Red blood cell non-ABO-identical transfusions are harmful: Really? *Transfusion*, *56*(3), 543–545. https://doi.org/10.1111/trf.13505

Zabarenko, D. (2014, September 19). The Nation Has a Major Blood Shortage. *ABC News*, pp. 1–2. Retrieved from http://abcnews.go.com/Health/story?id=117954

Zou, S. (2006). Potential Impact of Pandemic Influenza on Blood Safety and Availability. *Transfusion Medicine Reviews*, *20*(3), 181–189. https://doi.org/10.1016/j.tmrv.2006.03.001

APPENDIX A

MIXTURE DISTRIBUTIONS FOR AGE AT RECEIPT

A common assumption in blood inventory research is that units are new or the same age upon entering inventory. It often assumes that all units are new (Deniz et al., 2010; Sarhangian et al., 2017); or of the same age (Nahmias, 2011). However, in their study of a product with a random shelf life, which simulates a product with a fixed life at the upper echelon, **Gürler and Özkaya (2008) found "shape of the shelf life distribution has a significant impact on the costs and a precise estimation of shelf life distribution may result in substantial savings" and failure to properly account for age had significant costs.** This is supported by Broekmeulen & van Donselaar (2009) who found that taking the age distribution into account when making replenishment decisions can result in a large cost reduction in practice.

Most hospitals which transfuse red cells do not have their own internal donor centers and even those that do may require supplemental units from a regional blood center. Regional blood centers provide units of specified type, but not of a specified age, to hospitals that order them. Regional blood centers will frequently rotate blood from smaller or more isolated hospitals to larger hospitals with higher demand to reduce outdating. This rotation by blood centers means that the red cells delivered are of widely varying ages when they are delivered. Limited data is available to characterize the age of the units received by hospitals of varying sizes. Cheng et al. (2010) found that at the Queen Elizabeth II Health Sciences Center in Halifax Nova Scotia, Canada the overall mean age of units at receipt was 11.7 days with individual blood types having means ranging from 8.5 (B+) days to 29 days (AB+). Average age of O- units was 27.3 days at receipt. Furthermore data from the Mayo Clinic, Rochester demonstrates that, at least in some cases, the age is not normally distributed (Figure 40). The shape suggests that the distributions may be well modeled as mixture distribution. It is important to note that Queen Elizabeth II Health Center and Mayo Clinic have similar (very high) numbers of annual transfusions. A facility with lower demand would likely receive fresher units. Additionally, the mean ages by blood type of units received by each facility are different. The age distribution of the units will depend on the characteristics of the region and the polices of the regional blood center.

A mixture distribution may provide a representation that is more compact than an empirical distribution and more true to life than those assumed in literature. Here it is shown how a single set of historical age at delivery data was fitted by mixture distribution. Such mixture distributions were used as inputs in Chapter 5.

Mixture distributions were fit with data from six of the blood types using Python's sklearn (Pedregosa et al., 2011) package. Insufficient data was available for AB positive and AB negative blood to be able to fit a distribution. The Gaussian mixture model function of sklearn uses an expectation maximization (EM) algorithm to fit and therefore is not guaranteed to find the optimal fitting models. To address this, multiple restarts were used for each fitting. Full covariances were used but other covariance types are available. Data from each blood type was fit to a Gaussian mixture model of 1-9 components. Histograms of the A positive data compared to samples from the fit mixture model can be seen in Figure 40. Similar plots were found for the other blood types. The fit statistics average log-likelihood, AIC, and BIC were recorded and can be seen in Figure 41.

The blood types A positive and O positive have the largest numbers of samples making them more robust to the EM algorithm.

Mixture distributions with three components were chosen to model the age distributions of new units for A- and B+. For O+ a mixture distribution with four components was chosen and for A+, B- and O- mixture distributions with five components each were chosen. These were chosen to balance the tradeoff between accuracy and complexity. The final overall mean age of replenishment units at receipt were A+: 15.4; A-: 17.0; B+: 17.0; B-:19.9; AB+: 21; AB-: 21; O+: 20.7 and O-:15.

This analysis is a proof of concept. It is limited by the presence of data from only one specific (large) hospital being served by one specific regional blood center. In particular, the size of the hospital means that the units are likely to be older than units that go to hospitals with less demand. If this type of information were more widely available it may encourage researchers to incorporate it into their models.

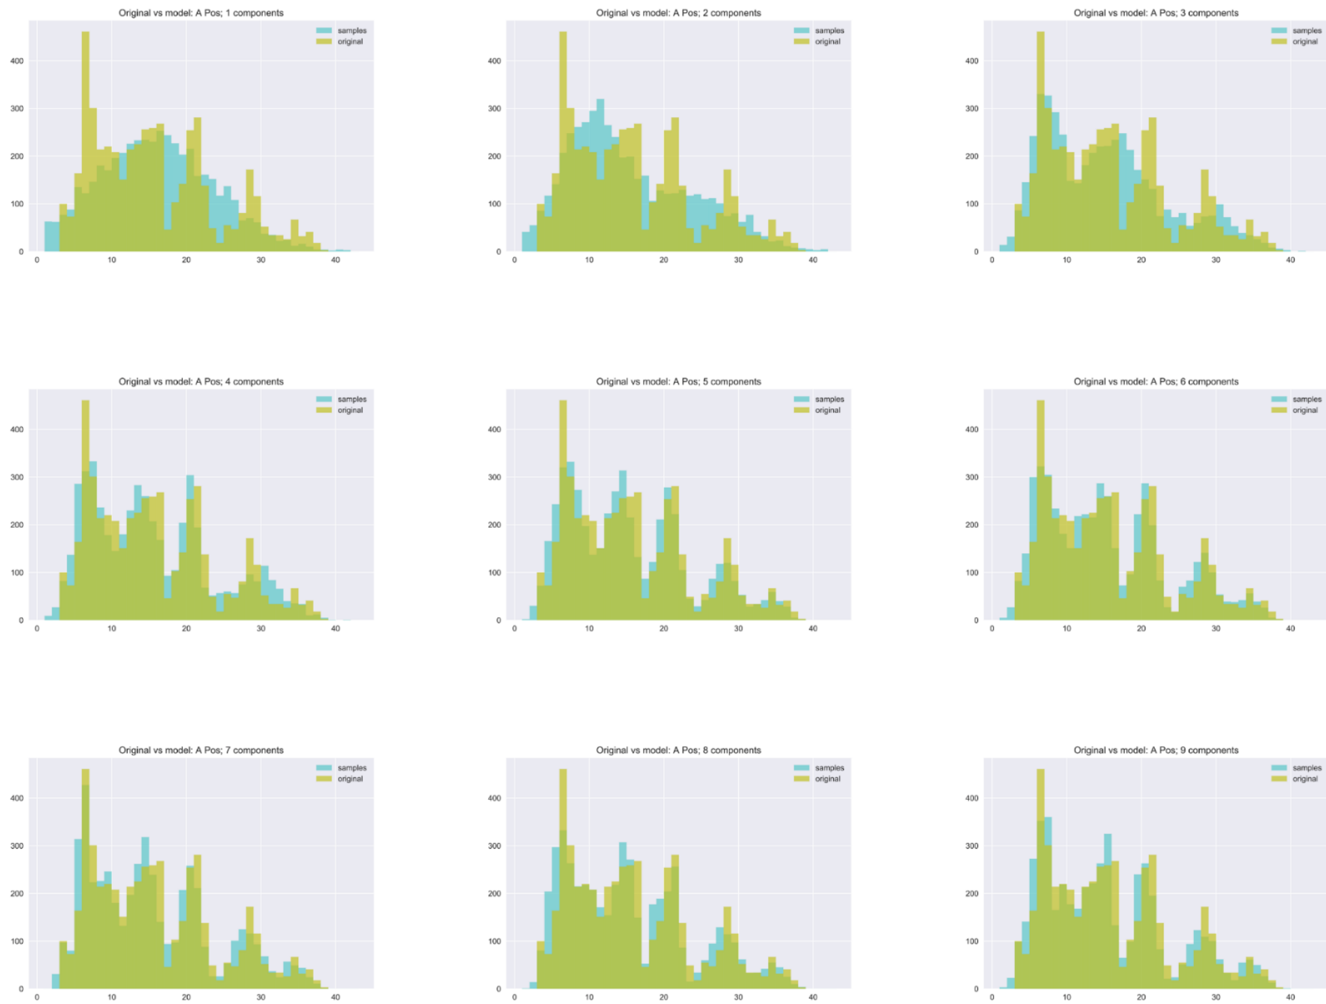

Figure 40. Original Data vs Model Fits for A Positive Age at Receipt Mixture Distribution

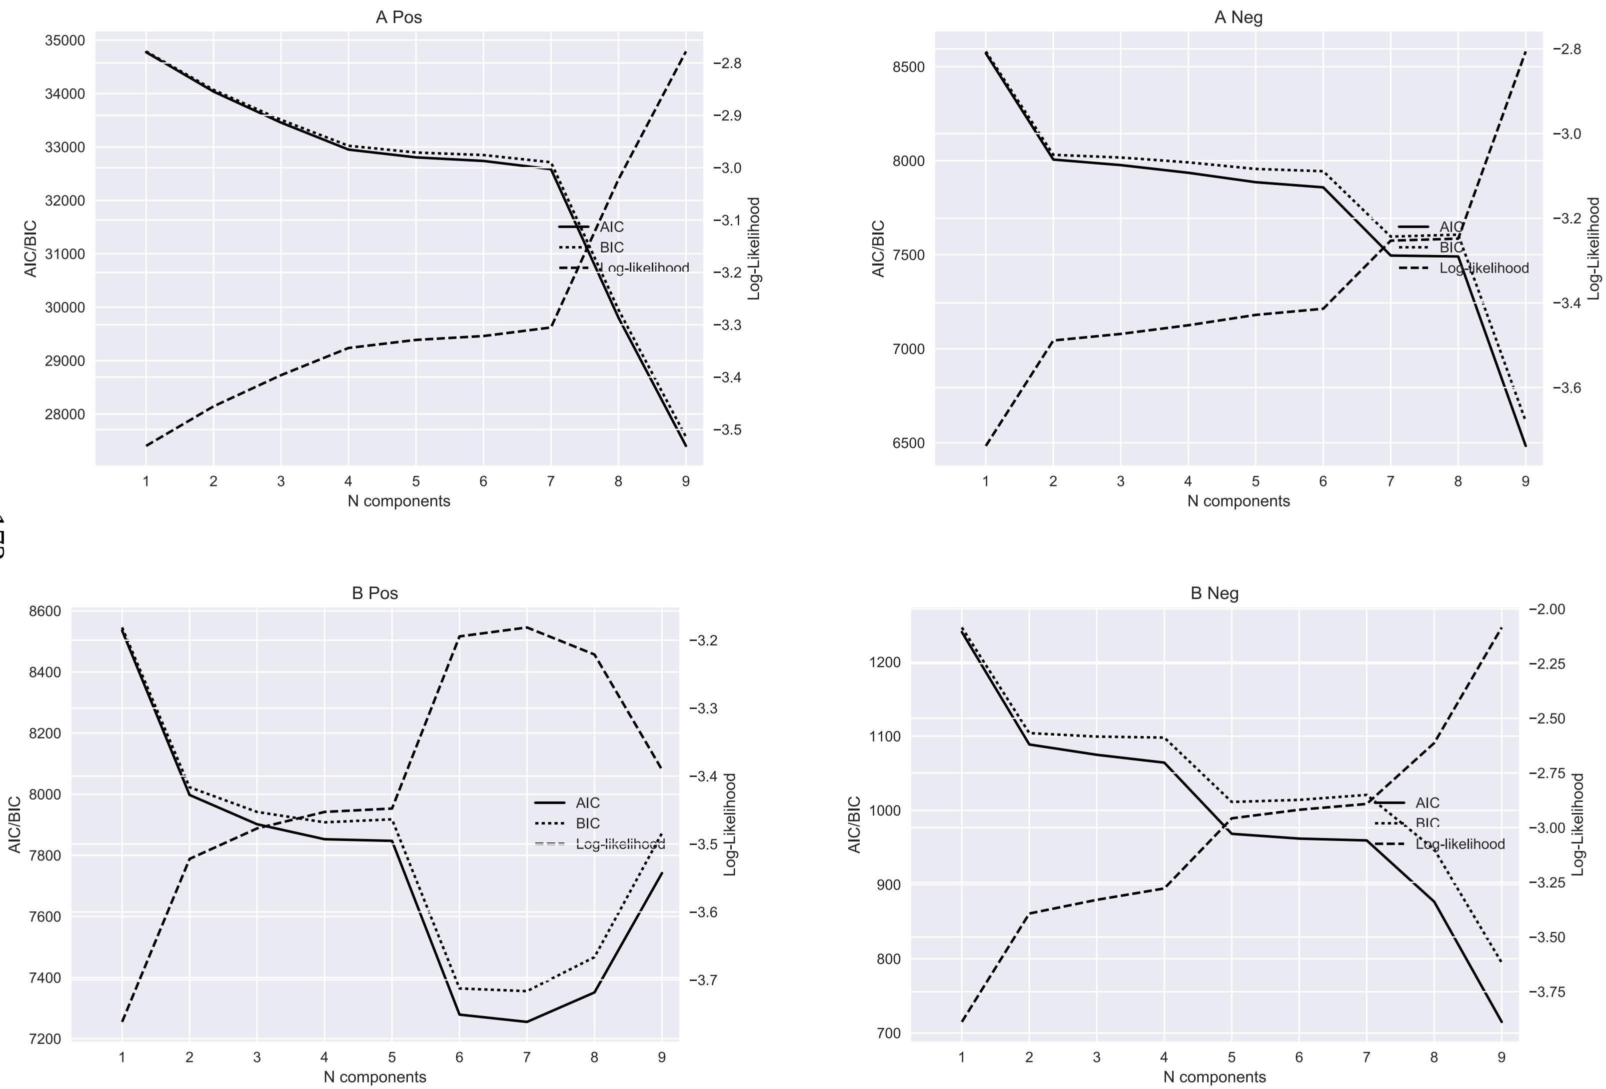
A Pos
A Neg
B Pos
B Neg
AIC/BIC
Log-Likelihood
N components
AIC
BIC
Log-likelihood

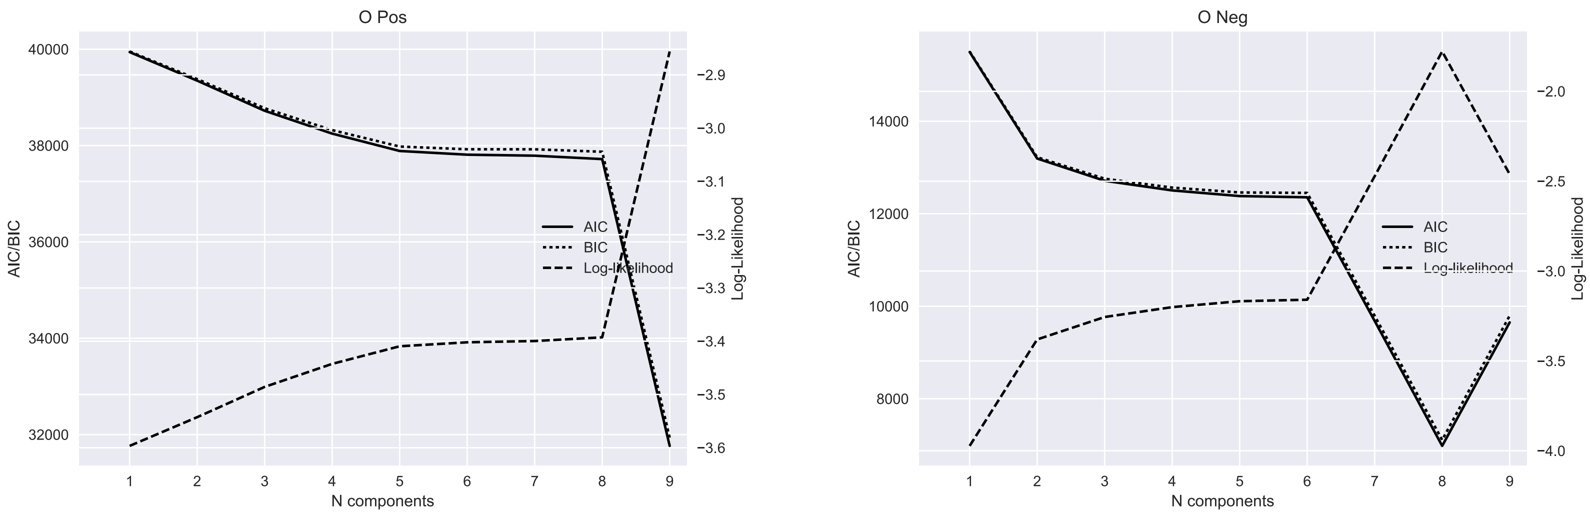

Figure 41. Fit Characteristic Curves for Each Blood Type

APPENDIX B

COBRE ALTERNATIVE FIXED PRIORITY POLICY COMPARISON

Table 32

API Simulation Results: Cobre Alternate Fixed Priority Policy Simulation Comparisons. Same additional test sequences as Table *17.*

	Unmet Demand		Outdated Units			Undiscounted Reward Improvement Over default		Discounted Reward Improvement Over Default		Daily Discount Improvement Over Default	
Policy Name	mean annual**	% Demand Unmet	mean annual**	% of Total Units	% Reduction from default	mean per period*	p-value	mean per period*	p-value	mean per period*	p-value
default	0.02	0.004%	307.86	34.2%	0.0%	0.00	NA	0	NA	0.00	NA
default1	0.04	0.007%	309.75	34.3%	-0.6%	-235.00	0.248	2	0.983	-52.91	0.631
default2	0.03	0.005%	308.08	34.2%	-0.1%	-33.75	0.865	26	0.736	-23.64	0.819
default3	0.03	0.005%	308.22	34.2%	-0.1%	-50.00	0.789	61	0.411	9.20	0.928
default4	0.04	0.007%	305.84	34.0%	0.7%	216.25	0.284	129	0.103	158.91	0.151
default5	0.05	0.009%	308.48	34.2%	-0.2%	-97.50	0.619	48	0.499	23.53	0.815

* Period is 6 weeks

APPENDIX C

DIRECT API SOLUTION SETUP CONFIGURATIONS: COBRE

Table 33

Direct API Solution Setup Configurations: Cobre

Direct Policy Iteration Setup					
	PI_direct_Cobre_0.1d	PI_direct_Cobre_0.2	PI_direct_Cobre2_reduced_param	PIDirect_Cobre2_extension	PI_direct_Cobre_0.3
bt_weight_use	Cobre_All_eight	Cobre_All_eight	Cobre_All_eight	Cobre_All_eight	Cobre_All_eight
use_mixtures	Mayo A Pos, Mayo A Neg, Mayo B Pos, Mayo B Neg, random, random, Mayo O Pos, Mayo O Neg				
prob_use	0.617162	0.617162	0.617162	0.617162	0.617162
target inventory	18,4,4,2,4,2,18,6	18,4,4,2,4,2,18,6	18,4,4,2,4,2,18,6	18,4,4,2,4,2,18,6	18,4,4,2,4,2,18,6
reorder_policy	M	M	M	M	M
resupply	TRUE	TRUE	TRUE	TRUE	TRUE
successful_service_reward	2000	2000	2000	2000	2000
no_unit_cost	-5000	-5000	-5000	-5000	-5000
mismatch_cost	-10000	-10000	-10000	-10000	-10000
outdate_cost	-1000	-1000	-1000	-1000	-1000
simLength	40	40	40	40	40
simRestarts	2	2	2	2	2
lengthOfTrajectory	2	2	2	2	2
discountFactor	0.9	0.95	0.95	0.95	0.95
sample_size	600	600	600	600	600
numberOfIterations	50	50	50	100	50
simCount	10	10	20	10	10
warmup_length	60	60	60	60	60
rollout_tail_length	40	40	40	40	40
rollout_tail_repeats	1	1	1	1	1
rollout_tail_count	2	2	2	2	2
rollout_test_sequences_count	3	3	3	3	3

rollout_test_length	40	40	40	40	40
rollout_test_repeats	5	5	5	5	5
rollout_length	0	0	0	0	0

APPENDIX D

DIRECT API SOLUTION SETUP CONFIGURATIONS: MAYO

Table 34

Direct API Solution Setup Parameter Configurations: Mayo

	PI_direct_Mayo0.1	PI_direct_Mayo0.2	PI_direct_Mayo_0.3
bt_weight_use	Mayo_All_eight	Mayo_All_eight	Mayo_All_eight
use_mixtures	Mayo A Pos, Mayo A Neg, Mayo B Pos, Mayo B Neg, random, random, Mayo O Pos, Mayo O Neg		
prob_use	0.990986	0.990986	0.990986
target inventory	650,180,170,30,1,1,650, 190	650,180,170,30,1,1,650, 190	650,180,170,30,1,1,650, 190
reorder_policy	M,W,F	M,W,F	M,W,F
resupply	TRUE	TRUE	TRUE
successful_service_reward	2000	2000	2000
no_unit_cost	-5000	-5000	-5000
mismatch_cost	-10000	-10000	-10000
outdate_cost	-1000	-1000	-1000
simLength	40	800	800
simRestarts	2	2	2
lengthOfTrajectory	2	2	2
discountFactor	0.95	0.95	0.95
sample_size	600	600	600
numberOfIterations	50	50	50
simCount	10	10	10
warmup_length	60	60	60
rollout_tail_length	40	400	1000
rollout_tail_repeats	1	1	1
rollout_tail_count	2	50	50
rollout_test_sequences_count	3	3	50
rollout_test_length	40	40	1000
rollout_test_repeats	5	5	1
rollout_length	0	0	0

APPENDIX E

API SIMULATION SETTINGS

Table 35

Settings for API policy comparison simulations. Basic setup. Alterations for additional analyses are described in the text.

	Cobre Scenario	**Mayo Scenario**
num_threads	4	4
bt_weight_use	Cobre_All_eight	Mayo_All_eight
use_mixtures	Mayo A Pos, Mayo A Neg, Mayo B Pos, Mayo B Neg, random, random, Mayo O Pos, Mayo O Neg	Mayo A Pos, Mayo A Neg, Mayo B Pos, Mayo B Neg, random, random, Mayo O Pos, Mayo O Neg
prob_use	0.617162	0.990986
Target Inventory	18,4,4,2,4,2,18,6	650,180,170,30,1,1,650,190
Reorder Schedule	M	M, W, F
resupply	TRUE	TRUE
Successful service reward	2000	2000
No Unit Cost	-5000	-5000
Mismatch Cost	-10000	-10000
Outdate Cost	-1000	-1000
warm_starting_point	TRUE	TRUE
warmup_length	~42 days	~42 days
simCount	20	20
simLength	42 days	42 days
simRestarts	4	4
discountFactor	0.95	0.95

APPENDIX F

PIPI_directMayo0.221 FINAL XGBCLASSIFIER PARAMETERS AND VARIABLE IMPORTANCE

Table 36

PIPI_directMayo0.221 final XGBClassifier paramters

Parameter	Value
'base_score'	0.5,
'booster'	'gbtree'
'colsample_bylevel'	1
'colsample_bytree'	1
'gamma'	0
'learning_rate'	0.01
'max_delta_step'	0
'max_depth'	5
'min_child_weight'	9
'n_estimators'	100
'n_jobs'	1
'nthread'	16
'objective'	'multi:softprob’
'random_state'	0
'reg_alpha'	0
'reg_lambda'	1
'scale_pos_weight'	1
'seed'	None
'silent'	0
'subsample'	1

		A+	A-	B+	B-	AB+	AB-	O+	O-
# of units of this age and type	1		0.0042						
	2			0.0047					0.0053
	3	0.0046	0.0037	0.0052				0.0024	0.0042
	4	0.0062	0.0035	0.0042				0.0026	0.0038
	5	0.0052	0.0074	0.0029	0.0039			0.0047	0.0003
	6	0.0049	0.0008	0.0057				0.0046	0.0062
	7	0.0029		0.0001	0.0062			0.0022	0.0048
	8	0.0014	0.0015	0.0051					0.0035
	9	0.0034	0.0015		0.0034			0.0058	0.0037
	10	0.0058	0.0071	0.0006				0.0042	0.0077
	11	0.0034	0.0033	0.0040	0.0035			0.0001	0.0037
	12	0.0004	0.0019	0.0042				0.0078	0.0051
	13	0.0057	0.0052	0.0034				0.0066	0.0049
	14	0.0051	0.0022		0.0037			0.0031	0.0028
	15	0.0076	0.0040	0.0037	0.0046			0.0056	0.0014
	16		0.0037	0.0065				0.0061	
	17	0.0038	0.0012	0.0025				0.0069	0.0046
	18	0.0060	0.0027					0.0057	
	19	0.0039	0.0031	0.0035				0.0054	
	20	0.0046	0.0065	0.0036				0.0045	
	21	0.0043	0.0039					0.0039	
	22		0.0033	0.0013				0.0062	
	23	0.0037	0.0075					0.0051	
	24	0.0052	0.0032	0.0043	0.0031			0.0039	
	25		0.0001	0.0028				0.0026	
	26	0.0044	0.0010	0.0044				0.0062	0.0072
	27	0.0075		0.0046	0.0053			0.0021	0.0037
	28	0.0047	0.0069	0.0049	0.0034			0.0059	0.0063
	29		0.0058	0.0041	0.0041			0.0052	0.0027
	30	0.0052	0.0028	0.0025				0.0017	0.0063
	31	0.0032	0.0043	0.0040				0.0033	0.0055
	32	0.0038	0.0016					0.0052	
	33	0.0050	0.0049	0.0040					0.0036
	34		0.0047	0.0037				0.0048	0.0018
	35	0.0037	0.0055	0.0056	0.0052			0.0054	0.0042
	36	0.0049	0.0046	0.0037				0.0057	0.0011
	37	0.0028	0.0045	0.0051	0.0069			0.0061	0.0058
	38	0.0044	0.0047					0.0054	0.0039
	39		0.0047					0.0042	0.0035
	40							0.0034	0.0019
	41							0.0026	
	42								0.0042
mean age		0.0025	0.0039	0.0040	0.0032	0.0041	0.0125	0.0055	0.0044
Q100		0.0084		0.0112	0.0068				
Q25		0.0027	0.0049	0.0060	0.0059			0.0064	0.0036
Q50				0.0081	0.0043			0.0040	0.0046
Q75			0.0033	0.0019				0.0010	0.0045
sum_compat		0.0072	0.0050	0.0058	0.0069	0.0066	0.0045	0.0053	
total units		0.0050	0.0042	0.0049	0.0038	0.0163		0.0085	0.0066

Figure 42. Variable Importance of *PIPI_directMayo0.221* Final XGBClassifier Model. Blue is more important; red is less important. Blank is not important. Current patient blood type importance 0.0477