Jeffries, Thomas (2015) *Optimising Blood Donation Session Scheduling in South East England.* Doctor of Philosophy (PhD) thesis, University of Kent,.

OPTIMISING BLOOD DONATION SESSION SCHEDULING IN SOUTH EAST ENGLAND

A DISSERTATION SUBMITTED TO
THE UNIVERSITY OF KENT
IN FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY IN MANAGEMENT

Thomas Jeffries
May 2015

Abstract

It is essential that all countries operate a form of blood banking service, where blood is collected at donation sessions, stored and then distributed to local healthcare providers. It is imperative that these services are efficiently managed to ensure a safe supply of blood and that costs and wastages are kept minimal. Previous works in the area of blood management have focussed primarily on the perishable inventory problem and on routing blood deliveries to hospitals; there has been relatively little work focusing on scheduling blood donation sessions.

The primary aim of this research is to provide a tool that allows the National Blood Service (the English and Welsh blood service) to schedule donation sessions so that collection targets are met in such a way that costs are minimised (the Blood Scheduling Problem). As secondary aims, the research identifies the key types of data that blood services should be collecting for this type of problem. Finally, various what-if scenarios are considered, specifically improving donor attendance through paying donors and the proposed changes to the inter-donation times for male and female donors.

The Blood Scheduling Problem is formulated as a Mixed Integer Linear Programming (MILP) problem and solved using a variable bound heuristic. Data from the South East of England is used to create a collection schedule, with all further analysis also being carried out on this data set. It was possible to make improvements to the number of units under collected in the current schedule, moreover the number of venues and panels operated could be reduced. Furthermore, it was found that paying donors to donate was uneconomical. Finally, changing the inter-donation times could lead to a reduction in the number of shortfalls, even when demand was increased by as much as 20%.

Though the model is specific to England and Wales, it can easily be adapted to other countries' blood services. It is hoped that this model will provide blood services with a model to help them better schedule donation sessions and allow them to identify the data necessary to better understand their performance.

Acknowledgement

This thesis would not have been possible without the unique contribution of many different people, a few of whom I would like to thank now.

First and foremost I would like to thank my primary supervisor Dr Jesse O'Hanley, for his continued support throughout my studies, for his willingness to share his knowledge and understanding, for his encouragement and motivation even through the difficult times, for his generosity with his time and for his patience. I could not have asked for a better supervisor.

I would also like to thank the rest of my thesis committee, Dr Maria Paola Scaparra, Professor Said Salhi and Dr Yue Wu for their encouragement and advice and for their feedback, which helped improve this thesis.

I owe a special thank you to the National Blood Service (NBS), not only for the amazing work that they do everyday, but also for providing the data used throughout this thesis. In particular, I would like to thank Crispen for taking the time to meet with me to discuss the data requirements for this thesis and for explaining how the blood service operates and some of the challenges they are currently facing.

Finally I would like to thank my family - my Mum, Joe, Alex and Matthew for their love, emotional and financial support and my fiancée Vicky for all her love, motivation, support and for always listening.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This thesis seeks to develop a model that allows blood services to evaluate how effective their collection schedules are and, if necessary, propose more efficient ones. Additionally, this model would allow for different operating policies to be evaluated by the blood service. This is achieved through the formulation of a mixed integer linear program, which in turn is solved through the use of a bounding heuristic. The model is applied to a dataset that covers the South-East of England, allowing the current collection schedule to be analysed and suggestions for improvements to be made. Furthermore, the impact of changing the inter donation interval and the potential of paying donors are considered.

This first chapter highlights why the area of managing blood inventory is of interest by pointing out some of the difficulties faced by blood services. The chapter also provides a definition for the blood collection problem, that will be examined throughout this thesis. The chapter then moves on to highlight some of the previous works that have been carried out in the area before detailing what this thesis' primary objectives are and how these objectives will assist blood services in creating more efficient and effective blood collection schedules. The chapter concludes by providing a roadmap for the remainder of this thesis.

1.1 Motivation

Blood plays a vital role in the human body. Its primary purpose is to deliver oxygen and nutrients to organs and carry waste, such as carbon dioxide, away. The body naturally regulates the levels of each component in the blood stream, maintaining some, like red blood cells within certain tolerances, while releasing additional white blood cells or platelets only when needed. If one of these components were to drop below a certain threshold, the body would naturally be less efficient and if the drop is severe, or prolonged, then a blood transfusion may be necessary to correct any imbalances.

Blood transfusions are prescribed to quickly address any imbalance in blood levels that can arise for a variety of reasons. Common causes include blood loss, through trauma, childbirth or surgery. Conversely it could be due to an illness like cancer, sickle-cell anaemia, chronic anaemia or thrombocytopenia (low platelet count). As blood cannot be manufactured it must be collected from other healthy members of the population, called blood donors. Donors can safely donate whole blood just under every four months and usually donate just under a pint of whole blood (470ml) in a session.

The National Blood Service (NBS) was established by the National Health Service (NHS) in 1946. Its remit is to ensure a reliable and safe supply of fresh blood to English and Welsh hospitals and it is responsible for collecting, testing, storing and distributing blood and blood by-products. The NBS sells blood products to hospitals at cost; for red blood cells this is currently £124 per unit. The service collects blood units from around 1.4 million registered donors, who are capable of donating two-million units of fresh blood and 200,000 units of platelets per annum. It is important to note that all of the NBS donors are volunteer donors and receive no financial renumeration for their time or donations. The donation sessions that the NBS run are held within the community and are often held at village halls, churches and schools. The blood that is collected at these sessions is transported to processing centres, where the blood is split into its various components, tested and stored before it needs to be shipped to a hospital's blood bank.

At first glance, it might appear that the process of collecting and distributing blood is relatively straight forward. However, there are several factors that drastically complicate the process. To start with, blood components have a relatively short shelf-life. As blood is a living tissue, it naturally perishes. Red cells have to be used within 45 days of donation; while white cells have to be used within 24 hours and platelets within 5 days; plasma, if it is frozen, can be stored for up to a year. Furthermore, the components require different storage conditions: red cells need to be refrigerated, plasma needs to be frozen and platelets need to be stored at room temperature with constant agitation.

The second complicating factor concerns the presence of different blood groups and their distribution throughout the population, specifically this impacts how blood components are allocated to patients. There are eight primary blood groups A+, A-, B+, B-, O+, O-, AB+ and AB-, however not all of these blood groups are compatible. Table 1.1 summarises which donor and recipient blood groups are compatible. Note that O- can donate to every blood group and is sometimes called the “universal donor”, while AB+ can accept all blood types and is referred to as the “universal recipient”.

		Donors							
		A+	A-	O+	O-	B+	B-	AB+	AB-
Recipient	A+	✓	✓	✓	✓				
	A-		✓		✓				
	O+			✓	✓				
	O-				✓				
	B+			✓	✓	✓	✓		
	B-				✓		✓		
	AB+	✓	✓	✓	✓	✓	✓	✓	✓
	AB-		✓		✓		✓		✓

Table 1.1: Blood compatibility chart

Moreover, the eight major blood groups are unevenly distributed across populations and ethnicities. Table 1.2 shows the blood group distribution across the UK population. There is a significant difference between the most frequent blood type, group O blood (44% of the population) and the least common, group AB blood (4% of the population). There are also differences within the same blood group. For example, 37% of the population are O+ while only 7% are O-. These variances mean that certain blood groups are frequently in high demand, which can frequently lead to shortages.

Blood group	% of UK population with this group
O+	37%
O-	7%
Total Blood Type O	**44%**
A+	35%
A-	7%
Total Blood Type A	**42%**
B+	8%
B-	2%
Total Blood Type B	**10%**
AB+	3%
AB-	1%
Total Blood Type B	**4%**
Total Rh(D) Positive	83%
Total Rh(D) Negative	17%

Table 1.2: Distribution of blood groups in the UK population

The third factor that needs to be taken into account is the possible transmission of diseases through the transfusion of blood products. Diseases such as HIV and Hepatitis C are known to be readily transmitted through blood transfusions, however for other diseases, like vCJD, while there is no concrete proof there is still a possibility of it being transmitted. For this reason, all blood needs to be tested against various diseases before entering the supply chain, a process that takes times and incurs significant costs.

Finally, the last factor is an operational factor that requires a minimum wait period between donations, in order for donor's blood levels to recuperate. This is known as the inter-donation time and can vary from 8 to 16 weeks depending on the country and the donor's gender. This resting time has the effect of significantly reducing the number of times a particular donor can donate in a given time frame.

Taking into consideration these factors, the blood collection problem can be explicitly stated as follows: The objective is to visit collection sessions, in periods where donors are available, so that sufficient blood for each blood group is collected, that is disease free, while ensuring that the demand is met across all periods, in a way that is cost efficient. Owing to the difficulties highlighted in this chapter, the task of ensuring a fresh and efficient supply of blood to NHS hospitals is an interesting and complex area of study.

1.2 Research Contributions

Numerous studies have been carried out looking at problems relating to blood collection and delivery, including inventory management, the scheduling of deliveries, forecasting demand and donor arrivals, as well as considering the location of major blood distribution centres. This topic was particularly popular in the late 1970's and early 1980's when many key articles were published, including the review by Prastacos (1984), which summarises the main blood inventory management articles of that period. However, since the 1980's, and until more recently, research in the area had declined dramatically, even leading Pierskalla (2005) to state that "few studies of note in the application of OR/MS [to blood donation and supply] have appeared in the last two decades". More recently, there has been a renewed interest in the topic, with several more recent literature reviews appearing (Beliën and Forcé (2012) and Osorio et al. (2015)).

However, despite this renewed interest in the area there are still gaps in the literature. One of these gaps concerns how efficient blood collection schedules are generated. Previous works in this area have primarily focused on minimising the distance travelled between donation sessions by blood collection teams, turning the problem into a Travelling Salesman type problem. While this setup is utilised across America, it is in contrast to how the NBS operate, where a team operates a session for the majority of the day, making distance travelled far less important.

Therefore, the focus of this thesis will be on where a blood service should locate donation sessions, when they should be visited, and how donors are allocated to these sessions throughout a fixed planning horizon. This will be accomplished by developing a sophisticated mathematical model that better captures the constraints and objectives of how the NBS collects blood in England and Wales. The schedules produced by this model will be compared to the NBS's current schedule for the South-East of England.

Furthermore, this thesis will use data from the South-East of England to explore two key policy questions. The first question, is whether it would be of benefit to remunerate donors for their donations, while the second question explores the potential impact of changing the amount of time a donor needs to rest between donations - called the inter-donation interval - this is an important policy as it impacts how frequently donors can donate and therefore how large the regular donor pool needs to be.

To summarise this thesis seeks to:

- Identify key types of data that NBS blood centres should be collecting, specifically for the use of scheduling and other analysis.

- Develop a model and solution that can aide NBS staff to locate donation centres and allocate donors to them so that current demand is matched.
- Use this model to aid NBS staff in understanding their current collection schedule, and how this can be improved.
- Explore and evaluate various "what if" scenarios to allow NBS staff to evaluate possible policy impacts on their operations.

The primary contribution of this research is the formulation of a model that can schedule when and where a blood collection should be run and the allocation of donors to these sessions across the planning horizon. This approach differs from previous work mentioned above, which focuses on minimising the distance a collection team travels. Furthermore, a heuristic is proposed to efficiently solve the model. The model and solution methodology are applied to a detailed case study of the South-East of England to evaluate the efficiency of the current blood collection schedule. The impact of reimbursing donors is considered and changes to the inter-donation interval are also explored. The impact of the latter on the collection schedule has not yet been considered in the current OR literature.

The remainder of this thesis is structured as follows. Chapter 2 provides an overview of previous works, looking at general works in the blood management field, as well as focusing on locating and scheduling blood collection sessions. Our blood collection model is introduced and formulated in Chapter 3, this

Chapter also presents a method for solving this problem on a large scale data set. Chapter 4 describes the South East data set provided by the NBS which is used for the case study element of this thesis. Chapters 5 and 6 present the results, more specifically Chapter 5 presents the computational results, while Chapter 6 examines the possibilities of reimbursing donors for their donations and the impact any changes to the inter-donation interval could have on collection schedules. Finally, Chapter 7 concludes this thesis with a summary of the work, as well as some limitations and recommendations for future research.

Chapter 2

Literature Review

Chapter 2 provides an overview of the academic literature in the blood management field. Each of the papers in this chapter are categorised into either relating directly to the blood collection process or as being more general to the blood management process. The first section of this chapter provides the objectives of the literature review and introduces the blood supply chain before providing an overview of the state of the field and finally the structure for the remainder of the chapter is detailed.

2.1 Introduction

The objectives of this literature review are twofold. First, it aims to provide a comprehensive overview of the literature relating to the organisation of the blood supply chain and, secondly, it seeks to identify potential interesting areas for future research. Papers for this review were found by reviewing the references used in the literature reviews discussed below, as well as, searches performed on the Web of Knowledge and Business Source Premier databases and in the Google Scholar search engine. The search terms primarily focused on blood inventory management and the blood supply chain, with other searches being carried out to identify articles relating to blood collections and donor session scheduling. Further papers were found by exploring the papers' citations using Google Scholar's citation feature and the Web of Knowledge's forward and backward citation mapper. Papers prior to the 1960's were not included, with the exception of one. This is owing to the fact that there was very little research focusing on blood inventory management going on at the time. Some medical papers are included for completeness, however, only when the subject matter could potentially impact the blood supply chain from an operational perspective.

A simplified overview of the blood supply chain is provided in Figure 2.1. This Figure distinguishes between the two major organisations involved in the supply chain, namely the blood service (or bank) and the hospitals that they serve. The hospitals will operate a small internal blood bank from which doctors can order blood products for their patients. The hospital will also keep the blood service up to date on usage rates and indicate how much blood they expect to need in the future. Based on this information the blood service will generate a collection schedule which will be carried out in the collection phase. Once the blood has been collected the blood service will then process the blood, which

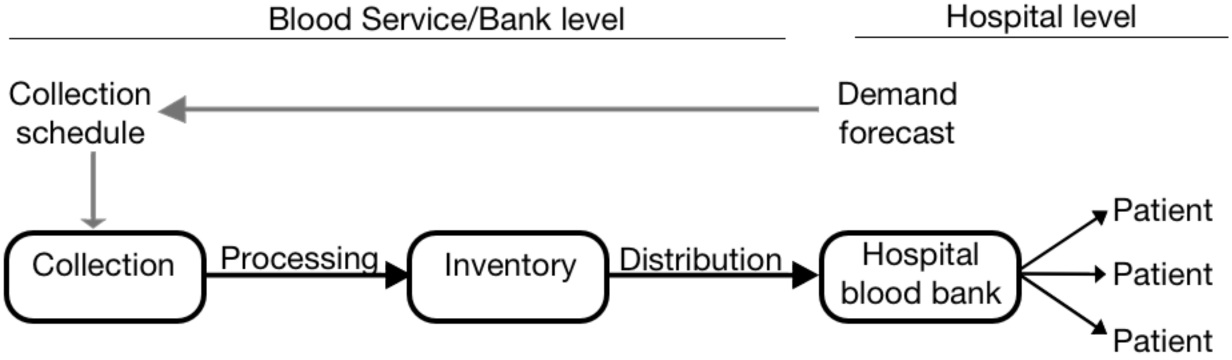

Figure 2.1: Overview of the blood supply chain

includes testing for diseases and separating the blood into its various components as indicated by the hospitals demand forecast. The blood is then placed in inventory before being distributed to the various hospitals that the blood service serve. Naturally at each of the stages of this process there are multiple decisions that need to be made and each one can have a significant impact on the other areas. Therefore, it is important to have an understanding of how each of these elements interact, which is why this literature review includes papers from each of these sections and not only the collection phase.

Research into areas relating to blood banking was particularly popular in America during the 70's and 80's, this popularity was fuelled (and funded) by government policy and incentives. Blood banking problems appealed to operation researchers primarily due to the perishable nature of blood products and as such, many contributions to the blood inventory management field were made. However, during the late 80's and 90's interest in the area had faded, Pierskalla (2005) puts forward several possible reasons for this decline: the primary reason was a lack of available funding, this naturally led researcher's to conduct research in other areas. The complexity of the remaining problems was also a contributing factor and finally, the research focus had shifted away from improving efficiency to improving the safety of blood products and transfusion practice.

More recently, there has been a renewed interest amongst operation researchers in the blood banking field. An overall interest in healthcare modelling, the predicted raise in future needs for blood products and the ongoing need for cost cutting in public spending have all contributed to this increase. Figure 2.1 (taken from Beliën and Forcé (2012)) highlights this renewed interest by comparing the number of publications that have appeared in each 5 year segment, from 1966 up until 2010, with clear peaks around the 80's, 00's and 10's.

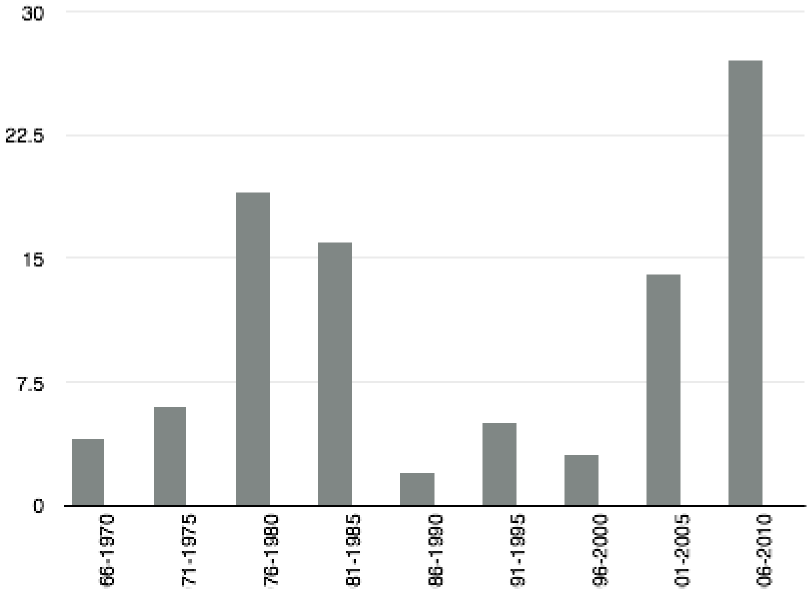

Figure 2.2: Number of articles in 5 year segments (Beliën and Forcé (2012))

Over the years, several reviews relating to blood bank management have appeared in the literature, the first major contribution was the review undertaken by Prastacos (1984), who provided an overview of the inventory studies undertaken in the 70's and 80's and is de facto review for this era. Pierskalla (2005), in this book chapter, provides an excellent overview of the blood banking problem, covering aspects relating to the regionalisation of blood centres in great detail, albeit this is not a full literature review. More recently, Beliën and Forcé (2012) published a more up to date review, however their paper primarily classifies the current literature to aide researchers in identifying relevant articles, at which it does an excellent job. Another recent review is that by Osorio et al. (2015) that provides an interesting review of the quantitive approaches used to solve decision making models within the blood supply chain.

This chapter is split up into three additional sections, section two of this Chapter focuses on the literature relating to the inventory, distribution, processing and forecasting activities carried out by the blood service and at hospital blood banks additionally the impact of the age of blood on the supply chain is considered in this section. Section three focuses on the blood collection problem specifically and covers planning for collection, improving the donation process and locating blood facilities. Furthermore, this section summarises the most frequent solution methods used when solving blood collection problems and discusses the classification of the blood collection problem in terms of the classical operation research problems. Finally, section four concludes by summarising the main themes found to be running throughout the literature, while also presenting how this research fills some of the gaps in the current literature.

2.2 Blood inventory management

This section considers how blood inventory levels are managed at a hospital blood bank as well as its supplier - the regional blood bank. Inventory management was the area that first attracted researchers to investigate the blood supply chain, due to factors such as: the perishable nature of blood, the need for constant restocking, the potential fatal consequences of a shortage and its wide application to other industries (for example foodstuffs and pharmaceuticals). Owing to this, the vast majority of research into the blood supply chain has focused on managing blood inventory levels and the various policies a blood bank manager can implement to improve overall efficiency.

A blood service is responsible for ensuring a fresh supply of blood to the region it serves. All donated blood is sent to the region's blood bank, where it is separated into its various components, tested for viruses and stored until needed. Each hospital in the region requests blood from the regional blood bank. However, to ensure there are no delays in transit and because of the large volume of blood needed on a daily basis, each hospital also operates an internal blood bank. This internal blood bank is where doctors place orders for patients needing transfusions. Fresh blood components are routinely delivered to the hospital at set intervals, however, if a shortage arises then the hospital manager can order an emergency delivery, although it is possible that the blood service would charge the hospital a non-routine delivery charge.

The obvious consequences of a hospital adopting an inefficient inventory policy are undue wastage and shortages, which besides wasting financial resources can also put patients' lives at risk. If a regional blood centre operates inefficiently, then there is the added risk that the region's hospitals will not be able to carry out routine operations and, more importantly, life-saving procedures. Therefore, it makes sense that inventory shortages (being unable to fulfil a doctor's request for blood), wastage or outdates (the number of units that expire before being used) are important benchmarks to measure how efficient a particular inventory policy is. Moreover, it also makes sense to measure the current age of each blood unit held in inventory, as this can provide insight into the effectiveness of the hospital's ordering practice. A very old blood inventory is much more likely to incur a higher outdate rate.

The management of blood inventory is clearly an important subject. The major developments in this area are first explored, starting by briefly explaining the theoretical work regarding optimal ordering levels for blood inventories, before moving on to look at more empirical research investigating ideal inventory levels at a hospital blood bank. Papers relating to blood cross matching and redistribution policies and the age of blood are then considered. Finally, current hospital best practice and factors affecting benchmark targets are considered, as well as, the need to accurately forecast demand levels and some of the works in the area.

2.2.1 Perishable inventories

Perishable stocks pose a unique problem to supply chain managers, as they must ensure sufficient stock to meet demand while taking into account the fact that their inventory has only a limited shelf-life. Over the years this problem has received attention from operation researchers investigating the perishable nature of a wide range of different items, from fashion to foodstuffs. Several researchers have also considered blood inventory systems however, these theoretical models often contain several simplifications. Firstly, most models assume that all inventory arriving at the hospital's blood banks is new or of the same age. Secondly, analytical models normally assume that a blood bank operates a FIFO (First-In-First-Out) distribution policy. Third, the models do not incorporate, for reasons of complexity, regional blood bank functions such as unused stock recall and redistribution before a unit expires, which have been shown to reduce wastage rates. Consequently, as Prastacos (1984) points out, model results may have limited application in actual blood banking practise. For an in-depth review of perishable inventories, the reader is referred to Nahmias (1982), Raafat (1991), Goyal and Giri (2001) and Nahmias (2011). Some of the key theoretical developments relating to blood inventories are highlighted next before moving on to more empirical works.

Theoretical works relating to blood inventories

Millard (1960) has been credited as the first to apply mathematical inventory modelling techniques to the blood bank inventory problem. In his paper, Millard

(1960) suggests that the criteria for evaluating blood bank performance should be the probability of a shortage occurring and the probability of a unit expiring, though the actual number of units can also be used. This work formed the basis for other researchers to develop more sophisticated blood inventory models.

Nahmias and Pierskalla (1973) investigate the optimal ordering policy for a product with a two period shelf-life, with the assumption that delivery is instantaneous (zero lead time), importantly they incorporate a cost penalty for expiring items. Later on, Nahmias (1975b) and Fries (1975) also consider the optimal amount to order however, they generalise Nahmias and Pierskalla's work to include a product with a finite shelf-life of m periods. Using dynamic programming both authors demonstrated that their models were able to analyse optimal ordering policies. However, both authors noted that this approach was only feasible when the product life was very short, less than 2 days, due to difficulties calculating large multidimensional dynamic programs. Nahmias and Pierskalla (1976) extend Nahmias (1975a) work to consider an inventory that incorporates both a perishable and non-perishable product to fulfil demand. In the case of blood banking, the perishable product would be fresh blood and the non-perishable product would be frozen blood - which has a much longer shelf-life.

Building on these works, Cohen (1976) and Chazan and Gal (1977) considered the critical number ordering policy: a policy where stock ordered is to replenish any used or discarded stock during that period, ensuring that the stock of blood remains constant. Both authors used Markov chain analysis as their principal method, with Chazan and Gal (1977) looking specifically at how large an inventory can be before an unacceptable number of units start expiring. They calculated an upper and lower bound for expected outdating per period as well as demonstrated that an approximate age distribution of the units in stock can be calculated when assuming daily demand is a Poisson distribution. This is important as an older stock signifies that the blood bank is ordering too much blood at a time and that the stock is more likely to be wasted.

Nahmias and Pierskalla (1976) attempted to create an algorithm for approximating the critical number policy. This was attempted due to the complexities in calculating an optimal order policy when the product lifetime is greater than three periods. Nahmias' algorithm performed well when tested against a simulated optimal policy for a product with both a two and three period lifetime, using an Erlang or exponential demand distribution (see Nahmias (1975b) for more details). The approximations were on average 3% to 5% below the optimal ordering policy, while cost was higher on average by 0.5%.

Brodheim et al. (1975) adopted a different approach to Nahmias and Pierskalla (1976). Instead of trying to determine an order up to policy (a fixed number of blood units that should always be in stock), they assume a fixed amount of units are delivered at the start of each day. This has the distinct advantage of allowing the total amount of inventory for each different age to be easily calculated. Previous works had suffered due to the large inventory vector needed to keep track of this. However, all of these models have several drawbacks. Firstly, they can only deal with a very short shelf-life, around two to three periods. This is mainly due to the complexity of the calculations involved. Secondly, the models only capture certain elements of the blood inventory process.

Empirical works

Owing to the drawbacks noted in the previous paragraph and the limited real world application of the theoretical models to blood banking, researchers, turned to more empirical studies to determine the optimal inventory levels to be held at the blood bank.

Jennings (1973) conducted a simulation study to better understand the benefits that an improved inventory system would have on an individual hospital as well as on collaborating hospital blood banks. Data was gathered from a Boston hospital and was used to develop a shortage-outdating operating curve for a single hospital. This curve showed for a selected inventory size the expected shortages and outdating percentage rates. This model was then extended to include multiple cloned hospitals, in order to explore the potential to further limit shortages and outdating across a shared inventory. Jennings (1973) demonstrated that when two hospitals collaborate, under a common inventory policy, then outdating and shortages can be reduced by around 45%. The savings are even greater when more hospitals are added to the system, with 5 and 20 hospitals achieving, 64% and 72% reductions, respectively. However, the practical cost of implementing this system would be astronomical due to the need to ship one unit at a time from one hospital to another to satisfy that particular shortage. There is also the problem of needing and keeping track of detailed inventory data. In light of these costs, a threshold transfer policy was tested. This policy only transfers units from one hospital to another once the inventory drops below a certain threshold. The policy showed that significant reductions were still possible (54% instead of 64% for 5 hospitals), but with much lower information costs.

Using simulation and similar to Jennings (1973), Vrat and Khan (1976) sought to minimise the shortfall and expiration rates of blood at an Indian blood bank. They were able to develop optimal inventory policies for rhesus positive blood, however they struggled to determine optimal levels for rhesus negative blood. This is due to the fact that rhesus negative blood has a much rarer occurrence in the Indian population (compared to more western populations). For this reason, the blood bank must either allow a certain level of wastage to occur in rhesus negative blood, or as Vrat and Khan (1976) suggest, the centre should keep a list of rarer donors that can be called upon when needed. The results of this paper highlight an important point that blood bank managers must take into account while trying to improve efficiency: blood groups are not evenly distributed throughout populations, meaning that optimising inventory levels will only go so far. Rarer blood groups must still be stocked, incurring additional costs, either through a higher outdate rate or through maintaining special collection efforts for the rare groups. This is especially relevant to ethnically-diverse countries such as the United Kingdom.

Brodheim et al. (1976) present a method for determining the inventory levels, for each blood group, to be held at a hospital, based on the average daily demand for blood and a permitted shortage rate (ranging from 1% to 20%). This method allowed blood bank managers to easily see the effects that changing inventory levels had on shortage rates, as such Brodheim et al. (1976) report that the method had been used as an inventory guide at several hospitals.

Extending the work by Brodheim et al. (1976), Cohen and Pierskalla (1979) article developed a decision rule for determining the optimal amount of inventory to be held at a blood bank. Using data collected from a Chicago based hospital and regional blood centre, Cohen and Pierskalla (1979) proceeded to identify the key variables that influence blood bank shortage and outdate rates. A simulation model was run varying each of the variables one at a time, resulting in 96 individual values for the optimal inventory level (S^*) (Cohen and Pierskalla (1979)). Regression was applied to these results to determine the input variables that most significantly influenced the optimal level of stock for a specific blood type. From this analysis an optimal decision rule (2.1) was determined.

$$S^* = 6.03(d_m)^{.7604}(p)^{.1216}(D)^{-.0677} \tag{2.1}$$

Where:

d_m : the mean daily demand for the particular blood type.

p: the ratio of the amount of blood reserved for transfusion and the amount of blood actually transfused

D: the amount of time (in days) that blood is reserved for a patient before being released back into common inventory.

By using this equation, blood bank managers could determine the amount of blood to be stocked for each of the blood types in such a way that shortages and outdates are minimised, both at the individual hospital and regional levels.

Other researchers have considered the effects that different issuing policies (e.g., Last-In-First-Out (LIFO) versus First-In-First-Out (FIFO)) and length of reservation periods (the amount of time a unit is kept for a particular patient before being returned to the general stock) have on the overall inventory level, including shortages and outdates. Issuing policies are discussed in the next section with reservation periods being discussed in a later section.

2.2.2 Issuing policies

Issuing policies can have a major impact on the availability of blood, knowing which issuing policy to adopt is an important decision for the blood bank manager. The main issuing policies to choose from are LIFO (Last-In-First-Out) and FIFO (First-In-First-Out). However, due to the fact that not all blood arriving at a blood bank is of the same age, researchers have generally studied what they call a modified LIFO or FIFO policy. The modified policies seek to issue blood based on the actual age of the unit rather than the strict order which the blood arrives at the blood bank.

Pierskalla and Roach (1972) considered what the optimal issuing policy for blood was. They proved mathematically that when demand and supply are deterministic, and the objective function was to minimise the number of units backlogged, then a FIFO issuing policy is best. They later relaxed the deterministic constraints to allow random demand and supply which also yielded FIFO as being the best issuing policy when compared to LIFO. However, certain medical procedures require much fresher blood, consequently a strict FIFO policy is not always practical in real-world situations. Therefore, it would be of interest to further explore what effects the requests for fresher blood have on the overall inventory levels, both at a hospital and a regional level.

Other notable papers looking at issuing policies include: Pegels and Jelmert (1970), who used absorbed Markov chains to investigate issuing policies, focusing on average inventory levels and the age of blood at transfusion. Rabinowitz (1973) proposed assigning patients that have a higher likelihood of being transfused older blood and also explored the impact of transfusing, compatible, rhesus negative blood to rhesus positive patients. Both issuing policies reduced wastage however, today's blood banks try to match donors with their own blood type, so transfusing patients with differing rhesus factors would not routinely be permitted. And, finally, Deuermeyer et al. (1976) who used simulation to demonstrate that a small department within a hospital that routinely over estimates the number of units needed could benefit from adopting a LIFO issuing policy instead. Deuermeyer et al. (1976) demonstrated this could result in a reduction of wasted units throughout the hospital.

A policy decision related to issuing and often investigated at the same time, is the cross-match and reservation policy, this is because if all issued units are transfused to the patient then the cross-match policy reduces to an issuing policy. Different cross-match policies are explored in the following section.

2.2.3 Cross-matching policies

Dumas and Rabinowitz (1977) considered this problem and found that on average one third of all requested blood was not used. To reduce wastage they tested a double cross-match policy, where if a doctor orders a transfusion for two patients, say one unit each, then one older unit is cross-matched for both of them as well as one younger unit. This results in the older unit having a higher probability of being used. The reasons for also cross-matching a younger unit are twofold. First, enough blood must be cross-matched in case both units are needed. Second, if the younger unit is not used, then it can be returned into unused stock with sufficient shelf-life left for future cross-matching. Based on data Dumas and Rabinowitz (1977) collected from a local hospital, the implementation of a double cross-match policy improved the probability of the first unit being used to 5/9, representing a 67% improvement.

Double cross-matching does have several major implementation issues. The primary issue is that patients need to be compatible with each other. As there are eight major blood groups this is not always the case, making matching difficult and exacerbated by the fact that orders are not placed at fixed intervals. Double cross-matching also increases the workload of the hospital blood bank; if every unit is double cross-matched then, logically, the work load also doubles. For this reason, Dumas and Rabinowitz (1977) propose that only blood of a certain age should be cross-matched, thus reducing the work load but also increasing the probability that the older blood is used before it expires.

The cross-match release period, which refers to the number of days a unit of blood is reserved for after the initial request, also has an impact on wastage rates. Pierskalla (2005) modelled eight LIFO and eight FIFO inventory systems, each with a different release period, ranging from zero days to one week. The FIFO systems consistently performed better than any LIFO system across all release periods. When the cross-match release period is near zero then the number of wasted units is greatly reduced. Therefore, it was suggested that a hospital should adopt a FIFO issuing policy while at the same time trying to minimise the cross-match release period.

Jagannathan and Sen (1991) proposed a model that would help blood bank managers determine the expected number of outdates and shortages in a stock that included cross matching. The advantage of this work is the ability to specify differing inventory parameters relating to cross-matching, which previous works had lacked. The model can also be used as a decision support system, allowing blood bank managers to determine the cross-match release period that minimises outdates. However, the main drawback of this model is that it does not incorporate double cross-matching, which is widely used in hospital blood banks.

With both medical and technological advances, over the past few decades, the way blood laboratories test for patient compatibility has evolved. Georgsen and Kristensen (1998) and Chapman et al. (2000) discuss the use of electronic (computer) cross-matching as opposed to traditional serological cross-matching. Both authors found that computer cross-matching had a positive effect on outdating and laboratory workload, with Chapman et al. (2000) also observing lower staff stress levels. More recently, Reesink et al. (2013) conducted a survey on the use of computer cross-matching in the United Kingdom and found that despite these benefits only 54% of hospitals currently use the technique.

2.2.4 Blood Redistribution

Not all hospital blood banks are homogeneous, some, such as large urban hospitals, will have a much higher demand and usage rate compared to smaller, more rural hospitals. This means that while large hospitals could achieve a lower wastage rate through proper inventory management smaller hospitals generally cannot. This is due to the fact that the smaller blood banks are much more sensitive to the unpredictable demand for blood, meaning that they must stock a larger quantity of blood units to avoid any shortages. However, this over stocking leads to a higher number of units sitting in stock until they expire and become unusable. Researchers and regional blood bank managers recognised this problem and considered the impact of rotating the blood stock on overall wastage and availability rates, along with any added transportation costs.

Blood rotation seeks to rotate nearly expired blood from a low usage hospital to a higher usage one. Thus, increasing the probability that the unit will be transfused before expiring. As part of their Programmed Blood Distribution System (PBDS) in Long Island, New York Brodheim and Prastacos (1979) implemented a once a period blood rotation policy, where each hospital blood bank received a shipment of fresh blood (one to two days old), which was eligible for rotation and a fixed number of older units, depending on the hospital's usage levels, which were not eligible for rotation. Any eligible rotation units that were not required at the end of the period were returned to the regional blood centre for redistribution as older blood. It was reported that this single rotation system contributed to the PBDS decreasing wastage by 80% (Brodheim and Prastacos, 1979).

Kendall and Lee (1980) claimed that Brodheim and Prastacos (1979) rotation model would only be of benefit to regions that had similar goals, a similar geographical area and a similar population and hospital mix as Long Island. As such, Kendall and Lee (1980) used goal programming to create a more general blood rotation model that could be applied to various regions. By adopting this approach, each region could adjust the various constraints (blood availability rate, inventory levels, fresh blood availability, blood outdate rate, average stock age and cost) in accordance with the regions overall policies and objectives.

Denesiuk et al. (2006) implemented and tested a blood rotation rule in several Northern Canadian territories. Denesiuk et al. (2006) demonstrated that despite relatively long shipping time (7 - 9 hours) between hospitals, older units of blood could still be salvaged by sending them to high usage hospitals. Despite the additional transhipment and packaging costs the paper demonstrated that the redistribution system still represented value for money. Denesiuk et al. (2006) also note that more research is needed to improve the reliability of packaging during transportation (blood must be kept within strict temperature limits). This research would be of great benefit to developing countries, especially those in more tropical climates, allowing them to implement a blood rotation system to manage blood inventory more effectively.

2.2.5 The age of blood

Though the current rated shelf-life for blood throughout developed nations varies between 35 and 42 days, it is possible to store blood for much longer. However, more recently, clinicians have started to call for the shelf-life of blood to be drastically reduced, arguing that for some patient categories, blood no older than seven days should be used. This appeal is being fuelled by some research that suggests the use of older blood has increased the mortality rate amongst transfused patients. This subsection will first look at the impact of raising the shelf-life of blood from 21 to 35 days had on inventory, before moving on to summarise the current debate around shortening the shelf life of blood, as well as, exploring the impacts this could have on current inventory practice.

Pegels et al. (1977) suggested that extending the shelf-life of blood could reduce the number of units needing to be collected each day, however if collection rates stayed the same then the average inventory for the region would increase (Pegels et al. (1977) observed a 40% increase in their model). However, the wastage rate only drops marginally, as the region is now over collecting (assuming no changes to the demand or usage rate have occurred). On the other hand, if collection levels were reduced to match inventory levels then under a 21 day shelf-life Pegels et al. (1977) model suggests that outdating would drop by approximately 50%. Clearly, extending the shelf-life of blood could lead to drastic improvements, nevertheless the impact of relaxing other inventory measures need to be considered.

Cohen et al. (1983) use a combination of simulation and regression to explore the impact extending the shelf-life would have on outdates and shortages. The main factors that influence the number of outdates and shortages were shown to be the inventory levels, the demand levels, the mean age of blood at delivery, the cross-match release period and the cross-match to transfusion ratio. It was found that if all contributing factors remained the same under the extended shelf-life then, as Pegels et al. (1977) found, shortages and outdates fell drastically (43% and 80% respectively). However, if blood bank managers allowed certain factors to be relaxed, such as increasing the cross-match release time or accepting the delivery of older blood then the benefits of the extended shelf-life would be wiped out. As an example, Cohen et al. (1983) found that if a blood bank manager relaxed the transfusion to cross-match ratio from 0.5 to 0.25 (with all other factors remaining equal) shortages would increase by 50% and outdates by 450%. Similar results were observed by varying the other factors that influence outdates and shortages. Cohen et al. (1983) do note that the increase in shelf-life can still be beneficial if the hospital's internal blood bank policies are good, however external factors (such as the age of blood on delivery and number of units delivered) deteriorate.

Kendall (1984) also sought to estimate the effects of adding adenine to blood, however, his focus was more on examining what the impact would be on the age of blood at transfusion. Using mathematical equations they created various tables which blood bank managers could use for updating operating targets under the extended shelf-life. Kendall (1984) argues that blood bank managers must set updated goals to reap the benefits of the extended shelf-life, if they do not then the outdate rate will remain the same (collaborating Pegels et al. (1977) findings). Kendall (1984) then applies his analysis to a mid-western blood bank in the US and reported that the region was able to reduce outdates from 10% to zero and that collections would need to be reduced by approximately 10%. Kendall (1984) concludes by remarking that the extended shelf-life of blood should eliminate the seasonality of blood donations. However, sadly, this has not been observed in the United Kingdom.

The literature clearly suggests that increasing the shelf-life of blood will aide blood bank managers in improving their outdate and shortage rates, though they will need to adjust current operating policies to get the full benefits. However, over the past few decades there has been a growing concern amongst clinicians that using older blood could be detrimental to patient outcomes, most notably in cardiac and paediatric patients. These concerns have surfaced both in the medical literature, as well as, in popular US media. The main driving factors behind these concerns are twofold. Firstly, there is a lack of understanding in how tissues utilise and receive oxygen from blood. Secondly, it has been well established that blood goes through numerous chemical and physiological changes while in storage. However, the impact these changes have on patients is still widely debated (Steiner and Stowell, 2009). Note should be taken, to the conclusion made by Wang et al. (2012) in their recent meta-analysis of previous studies. Wang et al. (2012) conclude that surgical patients that have been transfused with older blood have a significantly higher mortality rate, regardless of the type of surgery they underwent.

Naturally, the question of shortening the shelf-life of blood has caused blood bank managers and researchers alike to re-investigate and re-evaluate what impacts this change could have on their current policies. Fontaine et al. (2010) develop six scenarios, each with a different shelf-life (scenarios 1-5 varying between 7-35 days while scenario 6 allows a mixture of shelf-lives) and run simulations to determine the impact on availability and outdates at a large hospital in America. Fontaine et al. (2010) observed a large reduction in the number of units of blood available to the hospital when the shelf-life is dramatically reduced (7, 14 and 21 day shelf-life results in a 50, 20 and 10% respective decrease in availability). On the other hand, reducing the shelf-life to 35 days (the current legal shelf-life in the UK) then the availability of blood reduced by only 0.8%. Outdates were only found to increase by small amounts, when the shelf-life is reduced to 21 days then outdates increase by 3.2% and under a 35 day policy an outdate rate of only 0.4% was observed.

Interestingly, as the shelf-life is shortened, Fontaine et al. (2010) observed a large increase in the number of O- and O+ units transfused to non O- and O+ patients, indicating that patients cannot be matched with their respective blood type and that O blood will become much scarcer. It would be of interest to investigate the impact this increase in demand for O blood would have on the cancellation of routine surgeries, as O blood is often used for trauma patients, any shortfalls would result in excessive cancellations or emergency blood drives just for O blood. Fontaine et al. (2010) paper has, rightly, been criticised by Pereira (2011) as they assume that the blood centre would not change their collection schedules, casting doubt over their availability rate findings. As noted earlier, it is important that blood banks adapt their policies when the shelf-life of blood is either extended or shortened.

Blake and Hardy (2013) extend and improve on Fontaine et al. (2010)'s work by investigating the impact of shorter shelf-lives across a whole region in Canada. They build and validate a simulation model to evaluate the effects that a 28, 21 and 14 day shelf-life would have on the region's blood supply chain. Blake and Hardy (2013) found that larger hospitals are more capable of absorbing the short shelf-lives compared with smaller ones, they also found that the Quebec region is capable of handling either a shelf-life of 28 or 21 days with only a minor impact on shortages and outdates. However, when the shelf-life is further reduced to 14 days then all hospitals would face inventory issues with outdates reaching an unacceptable high 6.64% (Blake and Hardy, 2013).

If the shelf-life of blood is to be reduced to 21 days or less then Blake and Hardy (2013) note that drastic changes to the organisation of the blood supply chain would be needed and stresses the importance of a centralised blood banking system with an efficient blood distribution network.

Currently, blood can be stored for up to 42 days and with further medical and technological advances the shelf-life could be further extended. Increasing the shelf-life of blood could lead to significant reductions in outdates and shortages and reduction in the number of donors needed. However, recent medical research has called into question the safety and benefits of using older blood and have called for a reduction in its legal shelf-life. Any reductions would result in considerable changes to a regions operating polices and special consideration should be taken when rescheduling blood collection to meet demand.

2.2.6 Performance benchmarking and best practice

A more recent trend in the study of blood bank inventory management is to investigate current best practice within hospital blood banks and to set benchmarks for improvement. The primary objective of these studies is to identify the most efficient blood banks (usually those with very low outdate and shortage rates) and seek to replicate their success by implementing the efficient banks policies at less efficient but similar (in terms of size and function) banks. Researchers also evaluate the impact these improvements would have on the overall supply of blood.

The Blood Stock Management Scheme (BSMS) was set up in 2001 and monitors blood stocks at hospitals and blood centres throughout England and Wales. The BSMS collects data on blood product inventory and wastage levels, which are used to calculate two performance indicators, useful when comparing hospitals, these are: an issuable stock index and wastage as a percentage of issues (WAPI). The issue stock index measure is used to estimate the number of days of unassigned stock that are available and WAPI monitors the number of units wasted over a given period, as a percentage. In the literature, both indicators are used when comparing the efficiency of blood banks within England and Wales. The main advantage of these indicators is that comparison can happen regardless of the blood banks size, location and type, with similar measures being used throughout the world (Services, 2011).

Papers seeking to determine blood banks best practise include Perera et al. (2009) who used data from the BSMS and close-ended inventory practice surveys to evaluate the performance, in regard to stock level and WAPI, of over 300 hospital blood banks in England and Wales. Perera et al. (2009) found that hospitals who use a shorter cross-match reservation period (24 hours instead of 48 hours) have a lower stock index, as do the hospitals that only cross-match blood no more than 24 hours in advance (for elective surgeries for instance). Improved stock and wastage levels were also noted in hospitals that used computer systems to determine order sizes, as opposed to humans. The most likely reason for this, is that humans are likely to panic-order when they believe stocks are low, whereas a computer can keep track of every unit in the system. Hospitals that use a blood redistribution system were also observed to have lower wastage levels (indicated by a lower WAPI rate) but, on the other hand, slightly higher stock levels, these findings are consistent with previous works described above.

Perera et al. (2009) were unable to explore the impact of applying all of the aforementioned policies, as no hospital currently implements them all. Owing to the fact that there are a wide range of operating policies in use, it is important that national best practice guidelines are drawn up.

In a similar paper to Perera et al. (2009), Stanger et al. (2012b) identified seven

UK hospitals that have below national wastage rates (national wastage rate for 2008 was 2.27%, the selected hospitals had wastage rates in the region of 0.262% - 0.98%) and conducted open-ended interviews with the respective blood bank managers. The objective of the study was to gain insights into which policies these efficient hospitals had adopted and which ones most contributed to their low wastage rates. Stanger et al. (2012b) noted that current inventory models found in the perishable inventory literature were not in use at any of the hospitals considered in their paper. Instead, all seven hospitals relied on the experience of veteran members of staff to determine stock levels and the amount of stock to order. It was also found that human resources play a vital role in ensuring low wastage rates. This was achieved through adequate staff (laboratory and medical) training, effective and transparent internal communication and awareness amongst all staff of the greater impact of wasting a unit of blood. Issuing polices employed by the hospitals are in line with the academic literature, with all hospitals issuing oldest blood first (FIFO), it was also noted that larger hospitals benefited from the use of electronic cross-matching. Stanger et al. (2012a) recommend the implementation of each of these findings to improve efficiency, however it would be of considerable interest to determine if the use of operations research techniques, to set inventory levels and ordering quantities, could lead to further improvements, especially as the majority of wastage occur at the hospital level.

Pitocco and Sexton (2005) examine the efficiency of 70 American blood centres, accounting for a third of all red blood cells produced, with the aim of identifying inefficient centres and exploring the affects any improvements would have on the overall supply of blood in America. Pitocco and Sexton (2005) used Data Envelopment Analysis (DEA) to create an efficient frontier, centres found to be on the frontier and not dominated by another centre are considered efficient, with centres below the frontier or dominated by another considered to be inefficient.

Pitocco and Sexton (2005) recognised that site characteristics (population density and production targets for example) can affect the productivity of any given centre - rendering comparisons across all centres meaningless. To overcome this challenge each site characteristic was tested for correlation with the DEA model, if no significant correlation was found, as was the case, then site characteristics do not impact the centres efficiency and comparisons amongst all centres were possible.

Out of the 70 centres analysed, 34 centres (43.6%) were found to be inefficient and that vast improvements could be made. However, Pitocco and Sexton (2005) took a more conservative route and reported that if inefficient centres only eliminate 50% of their inefficiencies then all 70 sites combined could save approximately $62 million (or 4.5% of their current budget) and more importantly, increase the number of red blood cells available for transfusion by 375,465 units or a 7.3% increase. Pitocco and Sexton (2005) note that if all other blood centres are of a similar efficiency level to the ones in this study, then the efficiency improvements would go a long way to eliminating the large blood unit shortage currently being experienced in the United States.

Pitocco and Sexton (2005) recommends that inefficient blood centres identify similar efficient centres, in terms of input and outputs, and use these centres as benchmarks and improve efficiency by implementing their operating policies and procedures. The main drawback to this paper is that it does not report on any inefficient centres becoming more efficient and the resulting impact on the supply chain.

Heddle et al. (2009) sought to determine benchmark outdating rates for hospitals in Ontario, Canada, as well as, investigate, which, if any factors influenced the number of units outdating. Using logistic regression Heddle et al. (2009) were able to identify that the hospital's distance from the blood bank, the average number of units transfused, the month of the year and interactions between the distance and the number of units transfused all significantly impacted the outdate rate.

It was found, that hospitals that were closer to a blood centre were more likely to have a much lower wastage rate compared to more distant hospitals. The reasons for this are two fold, first, the lead time for a hospital close to the blood bank will be a lot lower, resulting in the closer hospital not needing to hold a significant amount of stock. Second, the more remote hospitals often only transfuse a small amount of blood each year and thus holding a smaller stock, so any one unit expiring would represent a much larger percentage. To ease comparison, hospitals were split into several different categories based on distance to the nearest blood centre (close, moderate and farther away) and the average number of transfusions per year (low, medium and high). All hospitals within each of the categories were then plotted onto bubble charts (the bubbles representing an individual hospital and the size of the bubble representing the number of transfusions) to identify over and underperforming hospitals. To set the benchmark outdating rate that each hospital group should aim for, Heddle et al. (2009) used the first quartile of outdates within a group. So hospitals nearest to the blood bank and transfusing a large number of units should aim for an outdate rate of less than 0.4% while, on the other hand, more remote hospitals with low transfusion rates should aim to outdate no more than 20.3% of units received.

2.2.7 Forecasting demand and usage of blood components

Forecasting plays an important role when planning inventory levels. Accurate forecasts can help blood bank managers determine when shortages are more likely to occur as well as inform them when inventory levels are too high for the current demand. As demand for blood is a highly stochastic process, it can be difficult to accurately predict future inventory requirements. As a result of this, research into blood forecasting has primarily used time-series analysis, looking at past data to identify possible future statistical patterns.

Cohen et al. (1981) applied time series forecasting to data collected from a single hospital blood bank with the aim of understanding daily and monthly blood demand. However, their Box-Jenkins model could not accurately forecast daily blood demand due to the very high variation in patients needing transfusions. The same problem was found when observing data over a monthly period; the forecast demand was a lot flatter than actual demand. Cohen et al. (1981)

conclude that it is difficult to predict with any accuracy the number of patients requiring blood transfusions at a single hospital. Consequently, other researchers have focused on forecasting demand at the regional and national levels, to try and aggregate demand.

Forecasting future blood needs over larger areas and for longer periods of time is vital to healthcare planners. It allows planners to predict overall demand levels across a region and to anticipate seasonal shortages, which in turn feed into both the blood collection strategy and hospital resource planning.

Frankfurter et al. (1974) developed a short-term planning tool for regional blood bank managers. The tool was developed to provide advanced warning to managers of any potential shortages or excessive surplus in the blood inventory. To achieve this, the model predicts daily inventory levels for up to 14 days into the future, using exponential smoothing, exponential functions and historical empirical data. The model proved to be very beneficial to the bank's management team, helping them decide when to cease blood collections due to significant surpluses being projected, when to continue operating as normal, and when to increase collection efforts, due to predicted shortages. These adaptive changes resulted in a much smoother inventory level resulting in fewer expiries and shortages. At the same time, the model, demonstrated that although it is very difficult to predict accurately daily inventory levels, it is possible to identify trends for a region over a relatively short time period (in this case two weeks).

Seeking to understand what blood is used for, Wells et al. (2002) undertook an observational study that gathered data on all blood transfusions in the North

East of England. Though the data collection periods were relatively short, only 28 days in total and the data is somewhat dated (collected in 1999 and 2000), important trends and observations can be drawn from this study. In the first instance, Wells et al. (2002) observed a significant rise in the use of blood transfusions for people over the age of 45 years. Secondly, they observed that in the North East of England blood transfusions are primarily used for medical purposes as opposed to surgical needs. Although this finding goes against older American and European studies (Vamvakas and Taswell (1994), Mathoulin-Pélissier et al. (2000) and Chiavetta et al. (1996)), improvements and advances have been made in numerous surgical procedures, which reduce the need for transfusions. Finally, they predicted that by 2008 the demand for red blood cells would grow by 4.8%.

In a related study to Wells et al. (2002), Currie et al. (2004) investigated overall blood use in a Welsh hospital trust as well as the impact demographic changes could have on blood supply and demand over the next 16 years. Despite only using data from two large hospitals, the collection period was a year-long – much longer than Wells et al. (2002)'s study. Currie et al. (2004) found that 9% of admitted patients over the age of 70 years would require a blood transfusion as opposed to only 1% of patients under the age of 30. The older population accounts for 70% of all blood use, whereas the younger population only account for 10%. These findings align with the conclusions in Wells et al. (2002), Ali et al. (2010) and Seifried et al. (2011) that an older population requires more blood, this is especially pronounced in patients in-between 70 and 80 years old, who are eight times more likely to need a transfusion than patients that are between 20 and 40 years old (Ali et al., 2010).

The fact that an older population requires more blood products is particularly relevant due to the current ageing of the baby-boom generation. Researchers have looked into the effects this older population could have on the demand for blood products. Using demographic data, Currie et al. (2004) produced forecasts of blood demand and supply, in England, up until the year 2026. They found that the use of blood would grow by 29% overall compared with current usage rates and that demand would outgrow supply by around 20%. More recently, Drackley et al. (2012) forecast supply and demand for blood in Ontario, Canada and found that if the Canadian Blood Service did not recruit more donors then demand would exceed supply as early as 2012 and that by 2036 there would be a shortfall of approximately 300,000 units.

Glynn et al. (2003) and Sönmezoglu et al. (2005) explore the impact a national emergency has on the collection rate and safety of blood. Using a statistical approach (chi-square) Sönmezoglu et al. (2005) note that following the 1999 earthquake in Turkey there was a significant increase in the number of first time donors and that despite this increase there was not a significant increase in blood transmitted infections, though the discard rate increased slightly. In a similar fashion Glynn et al. (2003) also note a large increase in first time donors and only a slight increase in the presence of blood transmitted infections following the 2001 terrorist attacks in New York. Though they also note that the number of donors that repeat subsequently is relatively low.

The problem with forecasting blood usage is two-fold. In the first instance, predicting short-term demand at a single hospital with sufficient accuracy is extremely difficult. This is primarily caused by the huge variation in blood usage policies, but also coupled with the impossibility of predicting blood demand, due to the random nature of the demand variables. Secondly, although long-term forecasts provide a general sense of future blood requirements, changes in legislation, screening processes, diseases and medical advances make these projections susceptible to large errors.

Despite these shortcomings, forecasts are still extremely useful for future planning. Short to medium term forecasts are particularly useful to regional blood bank managers as the predictions can alert managers to potential shortages and surpluses, as demonstrated by Frankfurter et al. (1974). While longer term forecasts allow blood services to anticipate the effects of large demographic changes on the supply and demand of blood.

2.3 Blood collection management

This next section looks at the problem of the collection of blood from the donor population. The first subsection examines the literature regarding how collections are scheduled, while the second subsection considers ways to improve the donor experience and the final subsection looks at the problem of locating blood facilities, including during disasters.

2.3.1 Planning for collection

In order to collect blood from donors, a blood service will run blood collection sessions, or blood drives, at specific venues and at preselected times and dates. The majority of blood services use two different types of collection venues: fixed and mobile sites. Mobile sites can be broken down further into two categories: regular mobile collection vehicles and self-contained vehicles. As the name suggests a self-contained vehicle can carry out all the necessary collection activities from within the vehicle and is usually a van or articulated lorry. However, this type of vehicle is limited to the number of donors it can process in a session. On the other hand, a regular collection vehicle will transport all the materials needed to a venue, typically a school, church hall or a local business, where it is set up for the session. The primary advantage of this system is that it can process more donors, however each venue must meet specific eligibility requirements. Fixed venues are generally found in large cities which can justify the running costs and have sufficient donor numbers to allow for a year round operation.

The blood collection schedule is usually fixed 6 - 18 months in advance and is influenced by numerous factors including venue and donor availability. As blood has a limited storage life and can only be collected from donors, blood is susceptible to seasonality effects, these effects usually appear around holidays (Christmas) and during the summer months, often resulting in non-optimal collection schedules where demand outstrips supply and emergency appeals and collections are oftentimes necessary.

Below, we outline the standard methods used to create a collection schedule, while at the same time pointing out their drawbacks. The first method that is often used to create a collection schedule consists of visiting each venue at a fixed interval. This method has the advantage of allowing donors to know far in advance when a collection schedule will be held. However, little effort is made to select the most cost efficient venues given the current resources available and demand levels, nor does this method account for any demand increases or donor changes to the donor makeup or any generic behaviour patterns.

Another method for creating schedules, that is commonplace in the United States, is to invite blood drive venues to send in their list of preferred dates for running a collection session. If a conflict arises, say there are not enough crews to visit all venues on a particular date or that visiting those venues is likely to result in over-collecting blood units, then the venues are scored based on a set of criteria, such as previous collection histories, and the highest scoring venue is selected first. While this approach ensures that highly successful venues are always accommodated for, it can leave smaller venues being left with undesirable dates. Furthermore, this approach suffers from the same drawbacks as the method discussed above but with the additional disadvantage that very unpopular dates could see no collection activity taking place.

Owing to the inefficiencies in the blood service's methods of creating collection schedules, researchers have, characteristically, sought to make improvements in this area. Initial works were undertaken by Pegels et al. (1975) and Cumming et al. (1976). Their works consisted of creating a model to aid blood collection planners to evaluate and analyse current and alternative schedules. However, these works suffer from drawbacks which meant that their usefulness was limited, which is highlighted next.

Firstly, the models have assumed that a predetermined number of units will be collected at a given session. More importantly, these solutions also assume that the number of units that can be collected remains constant despite previous collection efforts. The models fail to accurately take into account that donors must wait a relatively large time between donations, which will affect the subsequent number of donors available. Secondly, these models only seek to collect a total number of units and do not provide collection targets for each blood group. This can result in significant over and/or under collecting of certain blood groups. Moreover, accepting donors earlier than they are needed could result in significant shortages later on in the schedule. Thirdly, and finally, the models proposed by Pegels et al. (1975) and Cumming et al. (1976) are remarkably iterative, requiring planners to modify and re-run the models in order to assess if any improvements have been made. While these models do allow collection planners to evaluate minor changes to the schedule (opening/closing a venue for example) they are impractical when trying to examine whether large changes would be more beneficial. For these reasons, it is effortless to understand why blood service planners choose to adopt simpler methods, such as those discussed above.

Slightly later Pegels et al. (1977) sought to minimise the demand and supply gap caused by seasonality, this was explored by allowing blood planners to visualise their current collection schedule. Their model displayed the inventory levels, units collected, units outdating as well as the number of transfusion for the trial schedule. Once the trial schedule had been created, planners could then identify weeks in which they were over or under collecting. It reported that this method of generating the schedule resulted in a much smoother inventory level, as well as, far fewer outdates and emergency collections, when compared to the initial collection schedule.

However, this model has several important drawbacks which reduce its usefulness. Firstly, the scheduler must edit the schedule manually, there is no algorithm to help suggest the best changes to implement. Secondly, the model needs to know how many units can be collected at a specific venue and finally, the model does not take into account collection costs.

The paper by Cumming et al. (1976) formulates the Markovian population model used by Pegels et al. (1977). The Markovian model is used to calculate the number of units in inventory, the average age of blood at transfusion and the number of outdates given the expected collection schedule and demand rates. The paper also validates the model's performance against past schedules and performs a sensitivity analysis on the model's inputs.

Prastacos (1984), in his overview on blood inventory management states that scheduling blood collections has been treated as a complex vehicle routing problem. The blood bank manager must optimise which locations are visited as well as the routes taken to visit them in such a way that transportation costs are minimised, while at the same time still meeting collection targets. Prastacos

(1984) assumes that all collection points are mobile and that the mobile vehicles can move from one location to another throughout the day. As such, the problem can be formulated as a knapsack problem, where the objective seeks to minimise the transportation costs. However, the paper does not report on the implementation nor does it present any results.

Furthermore, their model makes two assumptions which limits its usefulness. Firstly, it is assumed that routes do not change over time. Secondly, the model also assumes that the expected number of units collected for a given route does not change over time. Therefore, ignoring the fact that donors must wait approximately 16 weeks before re-donating.

Other papers have looked at ways to optimally allocate collection nodes to blood centres to minimise transportation costs, these include Jacobs et al. (1996), Cerveny (1980) and Or and Pierskalla (1979). However, as the focus of these papers is to locate blood facilities and then allocate collection centres to them they are presented below.

Decision support system papers often include a collection module. These modules allow planners to identify when demand is outstripping supply in the short term so that corrective action can be taken, as is the case in Frankfurter et al. (1974) forecasting paper. Kendall (1980) allowed planners to estimate the scale of the blood collection effort based on regional objectives. However, these papers do not provide planners with any insights as to the locations at which they should focus their collection efforts, nor do they provide any detailed costing.

During periods of national disasters, epidemics and emergencies the collection schedule must often be changed. An et al. (2011) create a simulation model that evaluates the impact of epidemics on the availability of blood in the United States. They note that the key to reducing shortages during periods of higher demand is to plan collections in advance and to plan the intensity of the collection effort.

Relating to how much blood should be collected at a donation session Lowalekar and Ravichandran (2010) develop a simulation model to determine whether only collecting a set number of units would reduce outdate rates. Their paper explores two different cut off policies and shows that by introducing a cut off level that total collection costs could be reduced as well as wastage rates.

In much more recent papers Alfonso et al. (2015) and Alfonso and Xie (2013) formulated several mathematical models to generate mobile collection schedules. The first two models formulated annual collection plans, where the weeks a venue should be visited are returned. In both of these models the objective function is to minimise the number of units a region has to import from any other region. The first model is a deterministic MILP, though crucially this model assumes that the number of collections at each site is given and that the number of units that will be collected is known in advance and is fixed. However, the second model relaxes these constraints and a probabilistic MILP is formulated. This new model allows the frequency of collections at a given site to change, although the model can only choose from one of the collection frequencies available in the set. Furthermore, the number of donors that will attend this site is a probabilistic function which is based on the potential number of donors and their willingness to donate.

Based on the first two models Alfonso et al. (2015) also formulated a mathematical model for determining which day of the week a session should be held and how the human resources of the centre should be assigned in order that the number of hours worked does not exceed the given constraints.

Gunpinar (2013) formulates an integer programming model that seeks to minimise the daily distance travelled by bloodmobile, while also ensuring that sufficient blood to meet demand is collected. In this model a bloodmobile can visit a maximum of three collection points, where blood is collected before returning to the regional blood centre at the end of the day. The resulting model was solved using the CPLEX solver, a branch and bound algorithm and a column generation algorithm. It was found as the demand for blood increased the bloodmobiles were travelling greater distances and that a greater number of bloodmobiles were required, it was also found that visiting only three (out of a possible 60) collection points a day was sufficient.

2.3.2 Improving donor sessions

Donation sessions all over the world follow a very similar system flow, which is presented in Figure 2.3. When a donor first arrives, they are greeted, registered for the session and handed a donor health check form to be completed while they wait. Once a qualified member of staff is available, the donor is called into a screening room, where their vitals are taken and they are tested for anaemia (a drop of blood is taken from their finger and put into a copper sulphate solution, if it sinks the donor is eligible to donate). If a donor does not pass either of these stages they won't be eligible to donate at this session. However, if the donor is eligible, a collection bag is issued with the donor's details on and they are sent to the donor room to be bled. A donor can be bled once a collection chair or bed is available, a needle is inserted into the donor's arm and approximately 400ml of blood is drawn, this process usually takes 8 - 10 minutes. After the donor has been bled, they are sent to a waiting room for refreshments and to be observed for roughly 15 minutes. The American Red Cross estimates that the total time needed to donate a unit of blood is approximately one hour. However, if the donors wait is excessive they are more likely to leave, which can negatively impact their likelihood of donating again in the future (Williamson and Devine (2013), Ownby et al. (1999)).

Brennan et al. (1992) developed and validated a simulation model for the American Red Cross in order to evaluate customer service and productivity levels at their mobile donation sessions. The primary aims of the simulation were to reduce the amount of time it took to donate blood. It was observed that under the current process queues built up around the first few stages in the donation process. In light of this, various scenarios were developed including: grouping all screening activities into one step and moving from a modular concept in the donor room to having staff responsible for all beds. The simulation indicated that several scenarios could be beneficial in reducing the total time spent in the system. These scenarios, were then tested at three closed (no walk-in donors) collection sessions. Each of the collection sessions showed a marked decrease in

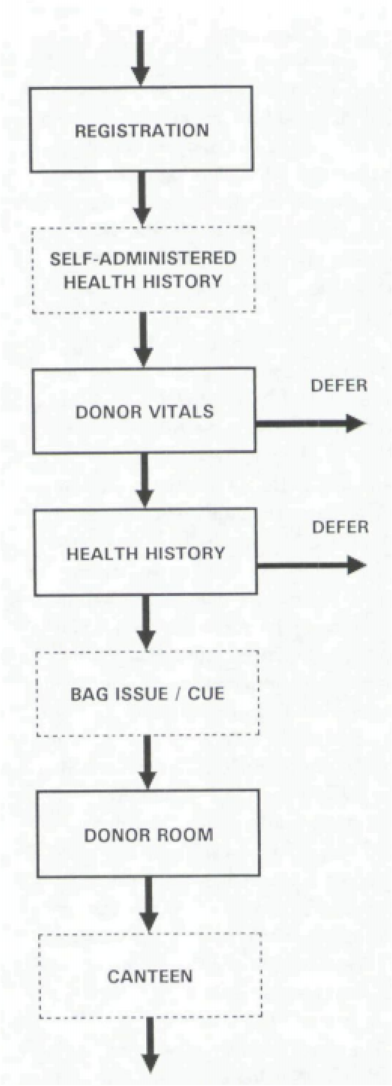

Figure 2.3: Blood Donation Flow Diagram (Brennan et al., 1992)

the amount of time taken to complete the donor screening tasks (50% reduction).

Hoyt et al. (1996) carried out a similar study to Brennan et al. (1992) for the US Navy's mobile collection sessions. The primary objective was also to reduce the total time spent donating and to minimise queues at large donation sessions (70-100 units of blood collected). They found that the number of beds could be reduced from 9 to 6, without negatively impacting the time spent in the system. They also experimented with changing the staffing level at various stages of the process and discovered that by increasing the staffing levels at the interviewing stage, queues and the time required to donate could be slightly reduced. Unlike Brennan et al. (1992) , Hoyt et al. (1996) do not report on having implemented their recommendations.

Determining the optimal configuration and number of resources to assign to each stage of the donation process, in such a way that minimises total donation time, was the focus of a study by De Angelis et al. (2003). Using simulation and optimisation De Angelis et al. (2003) modelled a fixed collection centre in Rome that collected on average 60 units a day. The simulation had two purposes, first was to allow management to explore what-if scenarios, while the second function was to generate training sets for a neural network.

A neural network was created to estimate the functional relation between the average time spent in the system and the quantity of each resource needed at each stage of the process. This function was incorporated into two optimisation models, the first model sought to minimise the average time spent in the system subject to a budget, integrity constraints as well as bounds on the number of each resource. The second model sought to minimise the total budget, subject to a maximum donation time that management deemed acceptable. De Angelis et al. (2003) were able to reduce the average time spent donating blood from 83 minutes to 66 minutes, this was achieved by increasing the number of beds available from 7 to 11 and without significantly impacting the budget. This optimal configuration was then plugged into the simulation model, which reported an average time spent in the system of 69 minutes, providing validation for the optimisation results.

Alfonso et al. (2012) extend on the aforementioned studies by including both fixed and mobile sessions, as well as, modelling the collection of multiple blood products (red blood cells, plasma and platelets). Petri net models are constructed to describe the fixed and mobile collection processes, before being converted into simulation models. The simulation's distributions are generated with data collected by observing donor flows at several fixed and mobile collection venues in France. The model also allows for both walk-in and scheduled donors which other works lack. Alfonso et al. (2012) report that their model is a useful tool to help managers re-evaluate current staffing levels and scheduling strategies. It was also reported that the model was being used to redesign local fixed collection centres, although no details or results were provided.

Being able to predict how many donors will attend a session is essential to help planners determine a session's capacity, to keep queuing times down and importantly avoid over collecting. By creating a logistic regression model Bosnes et al. (2005) were able to determine the likelihood of a donor attending an appointment through 18 explanatory variables. Using donors' previous donation histories, Bosnes et al. (2005) were able to identify which of the explanatory variables positively or negatively influenced the probability of a donor showing up to a session. Older donors were found to be more likely to keep their appointments, as were donors with a good track record of donating, as well as, donors who had kept their appointments over the last two years. These findings collaborate with previous research done by Ownby et al. (1999) and James and Matthews (1996). Bosnes et al. (2005) also identified that how the appointment was made was influential, with donors that experienced personal contact (as opposed to a written invitation) with the blood service having a higher probability of donating. On the other hand, it was found that the longer the time between making the appointment and the appointment itself, likewise, the number of deferrals in the last 2 years negatively affect the probability of attending.

The model was able to produce accurate short-term forecasts, however it suffered from several drawbacks as highlighted by Bosnes et al. (2005). Firstly, the model cannot make accurate future forecasts, as data for the model is only collected one day in advance of a given session. Secondly, the model cannot predict the number of cancelled appointments, although Bosnes et al. (2005) do present a strategy for building a cancellation model. Finally, the model does not provide a way for determining the number of walk-in donors likely to attend a session. It would be of interest to extend the model to address the last two drawbacks, as well as, to incorporate how donors react to social media contact (such as Twitter, Facebook and Email) invitations and appeals. These extensions would result in a much more detailed prediction model and at the same time allow blood services to evaluate the efficiency of their social media interactions.

2.3.3 Locating blood facilities and donor assignment

Cerveny (1980) sought to optimally locate bloodmobile base stations (where a bloodmobile is kept and maintained). A mathematical model is formulated that seeks to minimise the total yearly distance travelled by a blood mobile. The model is applied to two American Red Cross regions, a densely populated region, consisting of a single mobile base and a mixed, high and sparse density region, constituting a single large base in addition to multiple smaller satellite garages. The model was solved using a heuristic method and results indicated an area that the densely populated region would benefit from relocating their current operations too, which was subsequently done. However, at the time of publication it was too early to report on how beneficial this move has been. For the mixed density region, it was determined that they should seek to expand the capacity of their current bases. The model also proved helpful in aiding managers to identify mobile venues which are frequently visited and candidates to become permanent (fixed) venues.

Other studies investigating the relocating of blood facilities include Price and Turcotte (1986). However, whereas most models are extremely data hungry, Price and Turcotte (1986)'s approach is unique in the fact that they had to rely on sparsely available data. Due to the necessary data not being systematically collected and owing to budgetary constraints, Price and Turcotte (1986) relied on readily available census data, public transport data, donor addresses and the centres delivery schedules. From these data sets multiple gravity models were created to identify the respective central point for; the number of deliveries made by the centre, the number of visits to each mobile clinic site, the donor population and finally, the total population. Price and Turcotte (1986) helped management prioritise their objectives, while at the same time, generating a list of suitable sites to relocate the blood centres to. These sites were based on the proximity to the different centres of gravity, road access, public transportation links and venue availability. This study highlights the ability to generate insightful models even with scarcely available data.

After the devastating 1999 earthquakes in Turkey, the Turkish Red Crescent saw the opportunity to re-evaluate and improve their current operations. The overall objective of the project was to help improve the level of service offered to users (both hospital blood banks and donors). Quantitatively, the success of proposed changes would be measured by the total transportation costs and the maximum service response times. Şahin et al. (2007) constructed three mathematical sub-problems to allow various degrees of what-if analysis. The first model was a pq-median location-allocation model with the goal of minimising total population-weighted distances between the various blood service facilities and demand points. The second, a set covering model sought to minimise the number of additional satellite venues needed to ensure demand points were within a certain distance of supply nodes. Finally, the third sub-problem is an integer program that ensures the even distribution of blood mobiles throughout the region based on the population at each demand node. As a result of this study it was suggested that the Turkish Red Crescent promote seven blood centres to be regional blood centres - to oversee regional operations. It was also indicated that to ensure each demand point was within a 150km service radius an additional 11 satellite facilities were needed. The novelty in this study lies in the fact that it is the first study to adopt a hierarchal approach to locating blood services, whereas, previous papers had only considered a single level approach.

In a related paper Jabbarzadeh et al. (2014) develop a robust two stage stochastic optimisation model that can help design a blood supply chain during a disaster. In the first stage the model locates all of the permanent facilities (before a disaster occurs), while the second stage seeks to locate all temporary facilities and allocates donors to these facilities based on the disaster's impact. The model's overall objective is to minimise the total costs of the supply chain and is solved using branch and bound.

Jacobs et al. (1996) were asked by the American Red Cross to evaluate the impact of relocating a distribution and staging facility. Using integer programming Jacobs et al. (1996) were able to determine the optimal allocation of collection sites and customers to their current facilities, such that transportation costs were minimised. Other facility locations were considered in their study however, it was found that the American Red Cross could achieve substantial savings by simply reassigning current collection locations and customers. Jacobs et al. (1996)'s model demonstrates the need for blood services to evaluate and optimise their current operations before embarking on more costly initiatives.

Similar to Jacobs et al. (1996), Or and Pierskalla (1979) develop a comprehensive model to determine the optimal allocation of hospitals to blood centres. Or and Pierskalla (1979) also take into account the number of emergency and routine deliveries, furthermore they consider the delivery trucks' capacity and distance constraints. As in Jacobs et al. (1996) the overall objective is to minimise transportation costs. Or and Pierskalla (1979) present several heuristic methods to solve their Blood Transportation-Allocation Problem and present results for the Chicago area indicating improvements are achievable.

2.3.4 Common solution approaches

This subsection provides a summary of some of the methods used when solving the problems discussed in the blood collection literature. This subsection also seeks to classify the blood collection problem in terms of the classical OR problems. Initial work in creating collection schedules used highly iterative approaches, these works include those by Pegels et al. (1975) and Cumming et al. (1976). Naturally, this is not an ideal approach, with planners needing to re-run their schedules numerous times with no guarantee of finding an improving, or even optimal, solution. This iterative approach was due to the very limited computer capacity in this era.

More recent works in this area have often adopted either simulation (An et al. (2011) and Lowalekar and Ravichandran (2010)) or mathematical modelling (Alfonso et al. (2015) and Alfonso and Xie (2013)). The simulation methods are used to model how inventory levels vary over multiple time periods and how an increase, decrease or limitation of donors can impact these levels. The solution methods used to solved the mathematical models discussed in this literature use an off the shelf solver, though as Alfonso et al. (2015) mentioned that the problems with more variability take a significant amount of time to solve - in their case the models would take more than 5 hours to solve.

There are several works (Gunpinar (2013), Cerveny (1980), Jacobs et al. (1996) and Or and Pierskalla (1979)) that have considered the problem of minimising the distance travelled by blood mobiles. These models are commonly solved by heuristics (Cerveny (1980) and Or and Pierskalla (1979)), as well as branch and bound as in Jacobs et al. (1996) and Gunpinar (2013). This latter work also makes use of a column generation algorithm.

Papers that consider locating and allocating donors to blood centres often formulate the problems mathematically as in Jabbarzadeh et al. (2014) and Şahin et al. (2007), both of these problems are solved through the use of branch and bound. A problem that takes a different approach is that of Price and Turcotte (1986) where centres of gravity are considered.

Depending on what the researchers are exploring the blood collection problem can be classified under several of the more classical OR problems. The model can be classified under the variable routing problem when considering minimising the distance travelled between collection points and blood centres. Other problems that have examined where blood centres should be located and to which centres donors should be assigned naturally falls under the location-allocation problem. Works that have considered generating a collection schedule can often be classified, unsurprisingly, under the scheduling problem. This is the case for the work carried out by Alfonso et al. (2015) and Alfonso and Xie (2013), where they determine at which fixed interval each of the mobile collection sites should be visited.

However, the work presented in this thesis does not neatly fit into any of the aforementioned categories as classically understood. While the first part of the problem of deciding which venues to open is a location problem and assigning donors to these locations is an allocation problem, this does not take into account the changes in demand levels and donor levels across the multiple periods and that crucially each period is linked to every other period. Therefore, instead of being a spatial location-allocation problem (as classically understood) the problem considered in this thesis is considered to be a temporal location-allocation problem.

2.4 Conclusion

This chapter has provided an overview of the major literature relating to the blood supply chain. The first section considered works relating to inventory management, including issuing policies, crossmatching and blood redistribution. The age of blood, benchmarking and best practices and forecasting future demand for blood were also considered in this section. The second section explored the literature relating to the blood collection problem, looking at how schedules have been iteratively generated, how the donor experience can be improved and finally different methods for locating blood centres. This section also summarises the methods used and attempts to classify the blood collection problem in term of the classical OR problems.

Despite the large amount of literature considered works regarding collecting blood, specifically when a venue should be visited and how many donors should be invited are still sparse. The work that is closest to ours is that of Alfonso et al. (2015), however there are several important differences between their work and the model presented in this thesis.

Firstly, their paper treats all demand as homogeneous, it does not distinguish between the different blood groups, while this simplifies the mathematical model this does not reflect the need for rare blood groups, which can be more difficult to collect. The demand for each blood group has been incorporated into our model. Secondly, the cost of collecting the blood is not considered, the objective only seeks to ensure that the region is self sufficient regardless of cost, whereas our model specifically seeks to minimise collection costs. Thirdly, the paper seeks to determine the interval between each of the sites visits. This interval remains fixed for each specific site irrespective of future demand levels and donor availability. Furthermore, the model can only choose from a fixed set of intervals and each site can be visited a maximum of five times throughout the entire planning horizon, these limitations are not present in the model we propose. Fourthly and finally, the model does not distinguish between male and female donors, nor are the number of donors that donate at each session tracked, this could result in donors not resting sufficiently between donations leading to over estimating the number of units collected. We have distinguished between male and female donors and track donors, knowing exactly how many are eligible to donate in a given period.

Osorio et al. (2015) review highlights several gaps in the literature relating specifically to collecting blood. These include, the need for future research that considers different collection alternatives, the cost of collecting blood, locating mobile blood collection centres and taking into account periodicity in regular donors (Osorio et al., 2015).

The proposed model goes some way to filling these gaps. Firstly, our model seeks to minimise the total cost for collecting blood in the given planning horizon and secondly, the model takes into account and tracks the amount of time donors have to wait between donation sessions, this is done for both male and female donors. Therefore it is believed that the proposed model is an important addition to the blood collection literature.

In the first part of the following chapter the blood collection problem is mathematically formulated for England and Wales, while the second part of the chapter discusses solution methods.

Chapter 3

Blood Collection Model

This chapter starts by providing a description of the issues facing a blood service. It also describes the primary characteristics of the English and Welsh blood service, most notably the panel system. This chapter then moves on to mathematically formulate the blood collection model. The model's assumptions along with its limitations are also discussed. Finally, this chapter considers ways to solve the model. A simulated annealing heuristic approach is presented, which did not prove particularly successful during testing. A much more successful variable bounding heuristic is also presented.

3.1 Problem description

A blood service's primary goals are to collect, process and distribute blood to the hospitals it serves, either nationally or regionally, in the most efficient way possible, while also ensuring that the blood it distributes is free from disease. One of the fundamental tasks a blood service undertakes is the drawing up and execution of blood collection schedules. The blood collection schedule is usually determined several weeks or months in advance and tells the blood service which venues will be visited during the current planning horizon. Larger (national or multi-regional) blood services will often develop unique collection schedules for each region it serves, with each individual schedule incorporating the overall blood needs of the entire region. Any collection schedule should allow the blood service to accurately estimate how many units will be collected, in order to meet the hospital's demand over the planning horizon, as any shortfall could result in surgeries being postponed at hospitals. Furthermore, an additional advantage of creating a blood schedule is the ability to advertise collection sessions to donors in advance.

Naturally, a blood service will only have a finite number of resources available to it, which must be taken into consideration while developing a collection schedule. The most significant resource constraints arise from the human resources a blood service has available, more specifically, the number of crews available to the blood service each week, which place a limit on how many sessions the service can run in a given period. A blood session cannot operate without a medical practitioner (usually a nurse) present throughout the entirety of the session. A blood service will also have limited physical resources, such as blood mobiles, that further limit collection efforts. Furthermore, a venue owner may also decide to limit the number of times a blood drive can be held at their facility each week, preventing a blood service from continually visiting popular locations, as well as locations with high footfall.

Though a blood service might not impose strict financial constraints on collection activities, costs should be kept to a minimum. If the blood service is a publicly funded organisation or charity, then excessive collection costs would, understandably, spark public backlash.

A further difficulty is the presence of different and often incompatible blood groups present throughout the population. This can often culminate in over or under collection of units for a particular blood group, which results in higher wastage and shortage rates, respectively. An added complication when predicting how many units are likely to be collected during a session are walk-in donors - donors that attend a session but have not been invited or made an appointment - who can negatively impact a session's throughput rate and influence collection targets.

Owing to the constraints mentioned above, blood services recognise the importance of needing a blood collection schedule and the impact it can have on their operations. Nevertheless, the standard methods used to develop the collections schedule do not necessarily result in the most efficient schedules.

The English and Welsh blood system

As mentioned in the first Chapter blood collection is run by the NBS in England and Wales. The NBS have set collection targets for each period and part of their job is to create a collection schedule that meets these targets. Each regional planner creates a collection schedule for their region based on these demand targets, and the planner decides which venues are visited and on which day based on the venues availability and donor availability. There are no specific tools that aide the planner to generate the schedule and it relies largely on the planner's knowledge of the region he oversees. Below the specific characteristics of the English and Welsh blood system are presented, including the planning horizon, the panel system, the types of venues operated and the inter-donation times. These characteristics are crucial to formulating a collection model that reflects current operations.

In England and Wales, blood collection schedules are drawn up for a 40 week period, though there is no operational reason for this length of time. The blood collection model produces schedules for 40 weeks where each time period is a single week.

The blood service operate a panel system, these are based on the initial filing system used by the blood service, where panels were stored in index boxes. Every venue and donor belongs to a single panel which are not defined by geographic region (area size) nor by the number of donors or venues present. A donor is assigned to a panel based on the venue where they first donate and it is to this panel's venues that they will be invited to donate in future periods, though it is possible for a donor to change panels. A donor can be invited to any of the venues that are within their panel, this donor-venue assignment is not fixed and can change every period depending on the venues that are open and the demand levels for that period.

The blood service operate collection sessions from three different venue types. The first type is that of fixed sessions, these are not discussed here, instead they are presented in Chapter 4, owing to the fact that this venue type is relatively uncommon and is not modelled. The other types of venues include temporary venues and mobile venues. Temporary venues are venues that are set up in village halls, churches, schools and community centres with all necessary equipment transported to the venue and set up in the venue. These venues can process larger quantities of blood compared to mobile venues and are the most common type of venue operated by the blood service. There are two subclasses of mobile venues, small vans and larger trucks, each of which contain all the necessary equipment to collect blood from donors. These trucks are deployed from a regional base and travel to a location that is suitable for a mobile vehicle. As such, these trucks cannot be deployed to temporary venue locations. The larger trucks can process a greater number of donors during a given session and logically there is a limited number of each mobile vehicle available in any given period. There is a cost associated with operating a collection schedule at any venue and a crew must be present at the venue for a venue to operate. These venue types and their specific parameters are discussed in detail in Chapter 4.

The last characteristic concerns how the blood service regards donors. Donors are classified into two distinct categories, registered donors, those donors that have donated at least once before and are therefore registered with the blood service and walk-in donors. Walk-in donors are donors that show up to a session without having been invited to donate. Registered donors are invited to donate at sessions held within their own panel and each donor has a reliability score that determines how likely they are to respond to an invitation based on their donation history. Furthermore, the blood service know which blood group a registered donor belongs to, whereas for walk-in donors this is unknown. It is important to note that not all sessions will accept walk-in donors, though the large majority do. Crucially, once a donor has donated a unit of blood in England and Wales they must be rested, currently a male donor must rest at least 12 weeks (periods) while a female donor must rest at least 16 weeks (periods). This limitation drastically reduces the number of donors available at any one time.

The proposed model, which is presented in detail below, seeks to overcome some of the challenges presented to blood planners. Significantly, this model accommodates for the fact that not all donors that are invited, or say they will attend, actually attend a session. It has also been acknowledged that donors are unable to donate for several weeks after their last donation. This has a significant impact on the number of actual eligible donors available and can influence the success of subsequent sessions held at the venue over the planning horizon.

Furthermore, the model takes into consideration the various blood groups present throughout the population. It allows collection planners to set specific collection targets for each blood group, thus, enabling planners to avoid unnecessary over and/or under collecting in individual periods. It is believed, that this should reduce overall wastage rates. Additionally, it provides a useful tool for planners to efficiently analyse and explore different collection schedules.

3.2 Model Assumptions

Naturally, various assumptions have been made about how blood services undertake the scheduling tasks. These are discussed below, as well as, other, more generic assumptions that have been made. The model assumes near zero lead times this assumption is a common assumption made when modelling higher level blood systems. This is primarily done as the lead times involved are often of the order of hours so therefore do not have any major impact on systems over larger timeframes.

It is also assumed that the number of resources (crews, mobile vehicles) available in each period are known at the time the schedule is generated. This implies that the planner has already allowed for staff rest and sickness and that the number of mobiles available in each period reflects the current maintenance program.

Another assumption made is that the demand targets set by the planner reflect any major national and/or regional events that may be taking place over the planning horizon (for example special steps needed to be taken to ensure sufficient supply during the 2012 Olympics). More generally, demand targets should accurately reflect future hospital demand. Furthermore, it is assumed that the targets set for each period are significant and that units must be collected at that time, in other words, each period takes into account the useful shelf-life of the blood units. As a consequence of the demand target assumption, any blood that is over collected in one period cannot be carried forward and count against demand in future periods. If this assumption was relaxed careful tracking of the age of each unit throughout the planning horizon would be needed to ensure that the blood reached the hospital with a useful shelf-life remaining. This would raise the complexity of the model (Cohen, 1976).

The final assumption made is the most significant and concerns how the quality of the schedules are measured. Previous works(Gunpinar (2013), Cerveny (1980), Jacobs et al. (1996) and Or and Pierskalla (1979)) in this area have sought to minimise the distance travelled by the collection teams. While, minimising distance travelled is an effective way to build a schedule it has several drawbacks. Firstly, it assumes that the collection base is optimally located so that each collection site is as close as possible. Collection venues that are far away are penalised regardless of how cost efficient they are. Secondly, an objective that minimises distance travelled expects that a blood service's major costs are reflected in each mile travelled. Whereas, in reality, the total collection costs and any penalty costs are far more significant. Therefore, it is believed that collections costs, rather than distance travelled, are a much better way to evaluate various schedules. Hence, this model seeks to minimise total collection and penalty costs. Collection costs are tangible costs and can be further broken down into variable and fixed costs.

Variable costs consist of all costs associated with storing, testing, processing and distributing a unit of blood. Inventory costs should include the cost of the bag the blood is stored in, as well as any sundries. Testing costs should reflect all costs associated with testing the unit for diseases. Finally, processing and distribution costs should include the production costs for componentising a unit of whole blood and all transportation and refrigeration costs. Fixed costs, can be summarised as containing all costs associated with setting up and running a blood session, such as crew, transportation and the cost of hiring the venue. On the other hand, the penalty cost is not effectively a real cost, nowhere would it appear on the blood service's balance sheets. However, the penalty cost is necessary to capture the effects of, as well as how undesirable, shortages are. The penalty charge must be sufficiently large to discourage significant under collection throughout the schedule, while also allowing for the fact that collection targets for the region may be unrealistic for the current donor population to meet. Any shortfall would result in the blood service needing to source the units elsewhere. Furthermore, it is also assumed that this penalty cost is the same for each of the blood groups. Although it is possible to argue that rarer blood groups should carry a much higher penalty cost, most blood services (including the UK) charge the same amount per unit regardless of the rarity of the blood group. Therefore, it is assumed that the cost of sourcing any unit from a different hospital, region or country would be similar across the blood groups.

In the remainder of this chapter the formulation of the model will be presented and discussed, before moving on to discuss the model capabilities and limitations, in the final part appropriate solution methods are considered.

3.3 Model Formulation

Key indices and parameters are discussed next, with the complete list of the parameters and indices used to formulate the model presented in Table 3.1 and likewise, the model's decision variables are presented in Table 3.2. To formulate the problem of preparing an effective blood collection schedule mathematically, consider the following notation. It is assumed that the schedule is determined for a fixed number of periods, with each period indexed by t. Moreover, all venues and donors are assigned to a specific panel which is indexed by i. Each venue represents a site that a blood session can be held at and are indexed by j. For each venue j, it costs either f_{ij} or f_{ik}, depending on the type of venue visited (see Table 3.1). Visiting a venue allows a certain number of blood units to be collected, however a venue is limited by capacity and can only accept a certain number of total donors b_{ij}^{tv} or b_k^{mv}, for temporary or mobile sessions respectively. Moreover, a venue can only accept walk-in donors as a function of the number of registered donors invited to donate, furthermore a venue can be visited at most β_{ijt}^{tv} or β_{ijt}^{mv} times in period t.

Symbol	Definition
i	index for a panel
j	index for a site
k	index for a blood mobile vehicle type
t	index for a period (week)
g	index for a blood group
p_t	number of venue slots available in period t
m_k	number of bloodmobiles of type k available
f_{ij}^{tv}	cost of opening a temporary venue j in panel i
f_{ik}^{mv}	cost of opening a mobile venue of type k in panel i
b_{ij}^{tv}	capacity of temporary venue j in panel i
b_k^{mv}	capacity of blood mobile of type k
θ_i	fraction of walk-in donors that can be invited to panel i
α_i^m	fraction of male donors from panel i expected to show up at a session
α_i^f	fraction of female donors from panel i expected to show up at a session
β_{ijt}^{tv}	number of sessions that can be scheduled at a temporary venue j in panel i during period t
β_{ijt}^{mv}	number of sessions that can be scheduled at a mobile venue j in panel i during period t
a_{igt}^m	number of newly eligible male donors from group g in panel i during period t
a_{igt}^f	number of newly eligible female donors from group g in panel i during period t
γ_g	fraction of the general population with blood group g
c	cost of collecting and processing one unit of blood
w_g	penalty cost for each unit below the collection target of blood group g

Table 3.1: Model notation.

Finally, the following decision variables are included:

Variable	Definition
x_{ijt}	number of sessions scheduled at venue j in panel area i during period t.
y_{ijkt}	number of bloodmobiles sessions of type k scheduled at site j in panel i during period t.
z_{igt}^m	number of male donors in panel i from group g requested to donate during period t.
z_{igt}^f	number of female donors in panel i from group g requested to donate during period t.
z_{it}^w	number of walk-in donors in panel i requested to donate during period t.
u_{igt}^m	stock of male donors in panel i of blood group g during period t.
u_{igt}^f	stock of female donors in panel i of blood group g during period t.
s_{gt}	shortfall in units of blood group g during period t.

Table 3.2: Model decision variables.

Assuming the above notation, a mixed integer linear program model (MILP) for optimally scheduling blood sessions is given below and shall subsequently be referred to as the *blood collection problem.*

$$
\begin{aligned}
\min \underbrace{\sum_t \sum_j \sum_i f_{ij}^{tv} x_{ijt}}_{(i)} \quad & + \underbrace{\sum_t \sum_j \sum_k \sum_i f_{ik}^{mv} y_{ijkt}}_{(ii)} \quad + \underbrace{\sum_t \sum_g \sum_i c(\alpha_i^m z_{igt}^m + \alpha_i^f z_{igt}^f)}_{(iii)} \\
+ \underbrace{\sum_t \sum_i c z_{it}^w}_{(iv)} \quad & + \underbrace{\sum_t \sum_g w_g s_{gt}}_{(v)}
\end{aligned}
\tag{3.1}
$$

s.t

$$
\begin{aligned}
\sum_i \sum_j x_{ijt} + \sum_i \sum_j \sum_k y_{ijkt} \leq \quad & p_t && \forall\, t && (3.2) \\
\sum_i \sum_j y_{ijkt} \leq \quad & m_k && \forall\, k, t && (3.3) \\
x_{ijt} \leq \quad & \beta_{ijt}^{tv} && \forall\, i, j, t && (3.4) \\
\sum_k y_{ijkt} \leq \quad & \beta_{ijt}^{mv} && \forall\, i, j, t && (3.5) \\
\sum_g (\alpha_i^m z_{igt} + \alpha_i^f z_{igt}) + z_{it}^w \leq \quad & \sum_j b_{ij}^{tv} x_{ijt} + \sum_j \sum_k b_k^{mv} y_{ijkt} && \forall\, i, t && (3.6) \\
z_{it}^w \leq \quad & \theta_i \sum_g (\alpha_i^m z_{igt}^m + \alpha_i^f z_{igt}^f) && \forall\, i, t && (3.7) \\
z_{igt}^m \leq \quad & u_{igt}^m && \forall\, i, g, t && (3.8) \\
u_{ig1}^m = \quad & a_{ig1}^m && \forall\, i, g && (3.9) \\
u_{igt}^m = \quad & u_{ig(t-1)}^m - \alpha_i^m z_{ig(t-1)}^m + a_{igt}^m && \forall\, i, g, t = 2, \ldots, 12 && (3.10) \\
u_{igt}^m = \quad & u_{ig(t-1)}^m - \alpha_i^m z_{ig(t-1)}^m + \alpha_i^m z_{ig(t-12)}^m && \forall\, i, g, t \geq 13 && (3.11) \\
z_{igt}^f \leq \quad & u_{igt}^f && \forall\, i, g, t && (3.12) \\
u_{ig1}^f = \quad & a_{ig1}^f && \forall\, i, g && (3.13) \\
u_{igt}^f = \quad & u_{ig(t-1)}^f - \alpha_i^f z_{ig(t-1)}^f + a_{igt}^f && \forall\, i, g, t = 2, \ldots, 16 && (3.14) \\
u_{igt}^f = \quad & u_{ig(t-1)}^f - \alpha_i^f z_{ig(t-1)}^f + \alpha_i^f z_{ig(t-16)}^w && \forall\, i, g, t \geq 17 && (3.15)
\end{aligned}
$$

$$\sum_i(\alpha_i^m z_{igt}^m + \alpha_i^f z_{igt}^f) + \sum_i \gamma_g z_{it}^w + s_{gt} \geq d_{gt} \qquad \forall\, g, t \tag{3.16}$$

$$x_{ijt} \in \mathbb{Z}^+ \qquad \forall\, i, j, t \tag{3.17}$$

$$y_{ijkt} \in \mathbb{Z}^+ \qquad \forall\, i, j, k, t \tag{3.18}$$

$$z_{itg}^m,\, z_{igt}^f \in \mathbb{R}^+ \qquad \forall\, i, g, t \tag{3.19}$$

$$z_{it}^w \in \mathbb{R}^+ \qquad \forall\, i, t \tag{3.20}$$

$$u_{igt}^m,\, u_{igt}^f \in \mathbb{R}^+ \qquad \forall\, i, g, t \tag{3.21}$$

$$s_{gt} \in \mathbb{R}^+ \qquad \forall\, g, t \tag{3.22}$$

An efficient collection schedule should ensure that sufficient blood is collected to meet the demand moreover, this blood should be collected for the least possible expense. The blood collection problem should take into account both of these aims, which results in two separate objective functions - one that minimises collection costs and another that minimises demand shortfall. Owing to this, the blood collection problem is a multi-objective problem. However, by weighting the two different objectives it is possible to combine them into a single general objective function (a process known as linear scalarisation or weighted-sum method). In the blood collection model, and in light of the fact that collection expenses are already weighted by costs, this is achieved by applying a cost (weight) to every unit of shortfall. Therefore, when the cost per unit of shortfall is significantly greater than the other costs, the objective function prioritises meeting the demand levels before reducing collection costs. The combined objective (3.1) minimises the fixed costs, for temporary (*i*) and mobile (*ii*) venues, variable costs, for registered (*iii*) and walk-in (*iv*) donors and penalty costs (*v*) of collecting sufficient blood for the current planning horizon.

Owing to the staffing requirements needed to run each session, the blood service can only operate a certain number of collection sessions each week. The number of sessions that can be held varies week on week depending on each team's availability. Constraints (3.2) ensure that the maximum number of sessions that can be operated, in each period t, is not exceeded.

In a similar fashion the blood service has a limited number of blood mobiles available each week and constraints (3.3) ensure that the total number of mobile venues, across all panels, opened during period t, and using a bloodmobile of type k, does not exceed the maximum number of bloodmobile of type k available.

Naturally, there is a limit to how often venue j, can be visited in each period t, this is determined by the availability of the venue in that period. To ensure that only venues that are available in that period host sessions, constraints (3.4) place upper bounds on the number of sessions that can be scheduled at temporary venue j, in panel i, during period t. Likewise, constraints (3.5) place an upper bound on the number of mobile sessions that can be scheduled at venue j, in panel i, during period t.

Each venue will have a fixed amount of capacity - the number of donors that can be processed in a given session, which needs to be accounted for. However, as the model does not assign donors to a specific venue, it is only necessary that the total number of expected donors (for a panel and period) be less than or equal to the total capacity of the venues open in that panel and period. Note that donors are only ever invited to venues that are within their own panel. As such constraints (3.6) assure that the total number of expected donors (registered and walk-ins) in panel i, that show up to (a) session(s), held during period t at venue(s) j, in panel i, are strictly less than or equal to the sum of all the open session's capacity during that period.

Importantly, the number of walk-in donors that are allowed to attend a given session must be limited, there are two reasons for this. Firstly, the blood service does not want to run a session where only walk-in donors attend, some registered donors must be expected (and invited) to attend. Secondly, if the model were allowed to open sessions with only walk-in donors the solution could result in a significant over estimation of units that can be collected. This is due to the fact that the proposed model does not keep track of the number of walk-in donors available, therefore constraints are needed to prevent this behaviour and more closely match the actual collection system. As such, constraints (3.7) place an upper limit on the number of walk-in donors that can attend a session held at venue j in panel i, during period t. The fraction of walk-in donors that are allowed to attend is controlled by the parameter θ, which is the proportion of walk-in donors to registered donors that is usually observed within that panel i. Moreover, θ can be set to 0 for panels where walk-in donors are not allowed to donate.

This model must also keep track of the number of donors available and requested to donate in each panel and in each period. Constraints 3.8 to 3.15 provide this functionality. More specifically, constraints (3.8) and (3.12) ensure that the number of male donors (3.8) and female donors (3.12) requested to donate in panel i, of blood group g, in period t are strictly less than the respective number of available donors in panel i of blood group g for the same period t.

Owing to the fact that the blood collection schedule is a rolling schedule donors who donated in the last periods of the previous collection schedule are not eligible from the start of the new schedule, this is due to the fact that donors must rest between donations (12 periods for male donors and 16 periods for female donors). It is therefore necessary to split the number of donors available into distinct phases. The first phase tracks the number of eligible donors taking into account the later donations made in the previous collection schedule, while the second phase tracks the number of eligible donors once all donors from the previous schedule have been re-added to the donor stock/pool. Note, however, that the second phases start in different time periods for male and female donors. The number of donors who become eligible in the first phase is tracked by the a^m and a^f parameters for male and female donors respectively.

At the start of each collection schedule it is necessary to know how many donors are eligible in the first period and to assign these donors to their respective pools u^m and u^f for each panel i. Equalities (3.9) and (3.13) set the number of eligible male and female donors, from panel i, available in the first period, respectively.

After the first period and during the first phase of donor tracking the number of donors available in each of these periods is calculated as follows. The previous period's, $t-1$, donor stocks, $u^m_{ig(t-1)}$ and $u^f_{ig(t-1)}$, in panel i and for blood group g, are reduced by the number of donors $\alpha^m_i z^m_{ig(t-1)}$ and $\alpha^f_i z^f_{ig(t-1)}$ of blood group g that are expected to have donated in panel i during period $t-1$.

Note that not all donors that are invited to attend a session will actually show up and therefore parameters α_i^m and α_i^f present the fraction of donors from panel i that are expected to donate at a session within that panel for male and female donors, respectively. The number of newly eligible donors (from the previous schedule) a_{igt}^m and a_{igt}^f from panel i of blood group g in period t are then added to the current period's t donor stock, u_{igt}^m and u_{igt}^f, for panel i and of blood group g. The number of eligible male and female donors during the first donation phase are updated through equalities (3.10) and (3.14) respectively.

The constraints for the second phase of donor tracking are very similar to the first phase, the only difference is that instead of adding the newly eligible donors from the previous collection (a_{igt}^m and a_{igt}^f), the number of donors that donated and were invited either t-12 (for male donors) or t-16 (for female donors) periods ago become eligible to donate again. Equalities (3.11) and (3.15) are responsible for updating the number of eligible donors available.

Constraints (3.16) verify that sufficient blood, for each blood group g, has been collected to satisfy the predetermined demand for period t. The total number of units of blood of group g collected in period t is made up from three components, the number of registered male and female donors that have donated, the number of walk-ins that have donated and the decision variable s_{gt} . These three components must be at least equal to the demand d_{gt} for that period t and blood group g. The decision variable s_{gt} allows for under-collecting of blood group g in period t, however for every unit or fraction thereof that is under-collected a large penalty is applied in the objective function (3.1 part v).

The blood group for every walk-in unit collected is determined γ_g, which is the prevalence of that blood group in the general population. While, the number of registered donors' units collected is simply the sum of all donors that donated in that panel i of blood group g in period t.

Lastly, constraints (3.17) and (3.18) ensure the decision variables x_{ijt} and y_{ijkt} respectively, take on non-negative integer values, while constraints (3.19), (3.20), (3.21) and (3.22) ensure z_{igt}^m and z_{igt}^f, z_{it}^w, u_{igt}^m and u_{igt}^f, and s_{gt} are non-negative, respectively.

3.4 Model Limitations

No mathematical model can perfectly describe the real-world. Below is a discussion of the two major limitations and their impact on how the model will perform in real-world situations.

The first limitation is that the exact date a session should be held on is not provided. Instead, the model only provides the week in which a session should run - the week number is given by the t parameter in the x_{ijt} and y_{ijkt} integer decision variables. The primary advantage to adopting this approach is a reduction in model complexity. As such, a fine level of modelling is not required for each venue; we instead only need to know that it can be visited 3 times in the entire week. This leads to a direct reduction in the overall size of the problem as well. For a typical UK schedule, this reduces the number of periods t from 280 days to only 40 weeks, thus considerably reducing the size of the integer x_{ijt} and

y_{ijkt} decision variables. Though deciding dates for each venue is a considerable undertaking, the model can benefit from the planner's expert knowledge of the region they oversee. For example, the planner may have a better understanding of the demographic making up the donor population and can schedule sessions on days that best take advantage of their habits.

The second limitation is how the constant changes that occur in the number of available donors is dealt with. As at any blood service, donors are perpetually joining and leaving the available donor pool, which in turn affects how many donors are available in each period. The influences on the model made by this constantly changing donor pool are discussed next.

Blood services are continuously trying to recruit new donors to join the blood pool, either, through national recruitment campaigns (such as World Blood Week), or through word-of-mouth campaigns, such as encouraging current donors to invite friends to sessions. Some of the recruited donors will become regular donors, with the rest either never donating again or donating very infrequently (on average once a year). For these reasons, it is assumed that all newly registered donors that first donate during the current planning horizon are walk-in donors.

Furthermore, the model does not take into account how many active donors become ineligible to donate. There are two reasons behind the decision to not model this characteristic. Firstly, the need to track how many donors and for how long each donor is excluded would significantly increase the complexity of the model. Secondly, it is not known how frequently donors are excluded from the donor pool, nor is it known what the total and regional percentages for exclusion are. Consequently, it is assumed that any donors that are either permanently or temporarily excluded during the planning horizon can effectively be compensated for by walk-in donors. It is fair to assume that any blood service should seek to maintain a constant, or slightly increasing, overall donor population. Therefore, it is believed that even though the model does not incorporate all of the characteristics mentioned above, the net impact on the overall donor population should be negligible.

3.5 Solution Methods

This section considered the two heuristics that were built to solve the blood scheduling problem, which being an MILP is extremely difficult to solve. The first heuristic discussed is a basic simulated annealing heuristic, however, this discussion is relatively limited, with many of the details appearing in Appendix B. This is due to the fact that the initial results obtained by using this heuristic were very unpromising (see table 5.2) and was not used to solve the model using real world data. It is important to stress however that this was only a very basic implementation of the simulated annealing algorithm and that possible extensions, as presented in the final chapter, could drastically improve the algorithms performance. The second heuristic presented is a variable bounding heuristic, this approach was much more successful in solving the blood collection problem.

3.5.1 Simulated annealing

In this subsection, the implementation of the blood scheduling heuristic is considered. How a separate donor subroutine was created is presented first, before moving on to present how the initial solution was created and looking at the local search moves available. Finally, the acceptance, cooling and stopping parameters are presented.

Creating an algorithm to allocate donors effectively is a very difficult task, primarily due to the large number of variables involved and also how any changes in donor value can have large impacts on the overall performance of the schedules produced. As such, donor allocation and venue assignment were separated into two different problems.

A separate linear program (LP) was formulated to assign donors to a given set of venues. This model's objective (3.23) seeks to minimise variable costs, which consists of the cost of bleeding the donors, plus any penalty costs caused by a shortfall in the collection targets.

$$\min \sum_t \sum_g \sum_i c(\alpha_i^m z_{igt}^m + \alpha_i^f z_{igt}^f) \quad + \sum_t \sum_i c z_{it}^w \quad + \sum_t \sum_g w_g s_{gt} \tag{3.23}$$

The constraints for this model consist of the donor constraints found in the blood scheduling model, notably inequalities 3.9 through 3.15. These constraints are responsible for ensuring that venue capacities are not exceeded as well as tracking which donors are available and in which periods they should donate amongst others.

This subroutine requires the x_{ijt} and y_{ijkt} values to have been determined in advance by the venue allocation routine and are passed in as part of the function call. The subroutine returns the cost of assigning donors to the open venues as well as the penalty cost. This LP was implemented using the CPLEX OPL C++ libraries and every call was solved using the standard CPLEX LP solver.

The cost function is a crucial element in this heuristic's construction and allows the quality of the different schedules produced to be easily compared. The cost function utilised by this heuristic uses the components of the objective function found in the blood scheduling model. The venue, donor and penalty costs are all updated as changes are made to the solution allowing costs to be quickly evaluated.

The initial solution was created by using a greedy approach that is randomised on venue selection. As well as all the general model parameters the construction function requires the user to provide an assignment factor. This factor can be any value between zero and one and is used to determine the fraction of available venue slots to be filled in each period. A value of one would attempt to have all venue slots filled, while a value of zero would produce an empty initial solution.

The construction phase starts by calculating the current shortfall for every period in the planning horizon. The period with the highest shortfall is then selected to be explored. Next the number of venue slots to fill is determined based on the total number of slots available in the selected period and the assignment factor's value. The number of venues to open is always rounded up to the nearest whole number. A list of available venues is then created and venues are randomly added to the solution until either no more venues are available or the desired number of venue slots have been filled. Finally donors are assigned to the opened venues and a cost for the current solution is calculated. This process is repeated until all the venues in the planning horizon have been visited.

Having created an initial solution the next phase in the simulated annealing heuristic is to explore neighbouring solutions, while at the same time cooling the system. To create a neighbour the heuristic can make three different moves: an add move, a drop move and a drop/add combination (a swap move). Before any new neighbour is created it is necessary to select a period to explore, this is done by simply selecting one of the periods from the planning horizon. A move is randomly selected from the above list, though if the selected period has no open venues an add move is automatically selected and likewise if the period is full then the heuristic can only chose between making a drop or a swap move.

The moves available are fairly self-explanatory with the add move simply attempting to add a new venue or increase the number of times a current venue is visited in the solution, while a drop move seeks to remove a venue or decrease how frequently it is visited in the solution. The swap move is a two-step process first a venue or venue visits are decreased in the current solution then a new venue or the number of times a current venue is visited are increased - though the heuristic is prevented from re-adding the venue that has been dropped from the solution.

The first parameter that needs to be considered is the acceptance probability for a non-improving move. As discussed earlier in this chapter, the acceptance probability determines how frequently a non-improving move is accepted. The acceptance criteria used by the simulated annealing heuristic is essentially the same one used in the Metropolis-Hastings algorithms.

The temperature and cooling schedule play an important role in how well the algorithm performs; however selecting a starting temperature is more of an art than a science. Despite this there are several guidelines in place to help select appropriate parameters. In order for the simulated annealing heuristic to work the starting temperature needs to be sufficiently high to allow the search space to be explored, however it should not be so high as to turn the search into a random exploration. Therefore, it is suggested that the initial starting temperature should accept a relatively high number of worse solutions. The starting temperature utilised in this implementation was set to 10,000 degrees.

The second temperature related parameter is the cooling factor, which determines how quickly the system cools. A linear cooling schedule is often used and it is suggested that the cooling factor should have a value between 0.8 and 0.99. It is important to note that the higher the value the longer it will take the system to cool. Several different cooling factors were tested, with the lower values cooling the system too quickly resulting in poor schedules being produced. After these experiments a value of 0.97 was selected as the cooling factor.

The final parameter to be considered is the heuristic's stopping criteria, which determines when the heuristic should terminate. It was decided to terminate the heuristic after six hours had passed, rather than once the system had cooled.

The reasons for this decision are twofold, firstly all of the other models were allowed to run for this length of time and secondly, once the system had reached a point where accepting non improving moves was very unlikely, improvements in the current solution were still being observed. Furthermore, several different run time lengths were tested, with shorter lengths generating solutions with significant gaps (>10%) to the lower bound and although longer running times did produce improvements the amount of time required was considered excessive for the improvements made.

3.5.2 Variable bounding heuristic

The second heuristic implemented to solve the blood allocation model is a variable bounding heuristic. This solution method was implemented owing to the computational difficulties experienced in solving even small to moderate sized data sets, the results of which are presented in the next chapter.

The primary premise behind a bounding heuristic is to relax the integer variables and solve the underlying relaxed LP. Updated bounds are then placed on the variables and the model resolved as an LP. Once this has been completed a number of times, the fractional decision variables are restored back to integer variables and the model is solved using a MILP solver.

The bounding heuristic implemented in this thesis was written in the OPL scripting language, which is part of the ILOG CPLEX Optimisation Studio integrated development environment. The two phases of this heuristic: (1) relaxing and bounding and (2) variable fixing and solution of modified MILP are presented next.

The first phase of the bounding heuristic is to iteratively solve the relaxed MILP. The relaxation is achieved by allowing all the integer values (x_{ijt} and y_{ijkt}) to take on continuous values instead. The relaxed problem is then solved by CPLEX as an LP with the values of the decision variables x_{ijt} and y_{ijkt} being recorded. A value of 0.5 is then added to each variable before rounding down to the nearest integer. These integer values are then used as the new lower bounds to the variables. The process of adding 0.5 and then rounding down can be expressed mathematically as $\llcorner \hat{x_{ijt}} + 0.5 \lrcorner$ and $\llcorner \hat{y_{ijkt}} + 0.5 \lrcorner$. The next step is to attempt to solve this new model with the tighter lower bounds. The process is then repeated until the objective value of the new model is greater than or equal to the objective value of the previous iteration. Once this occurs, it is not necessary to continue as no improving solution can be found without possibly violating one of the constraints.

In the second phase, some of the x_{ijt} and y_{ijkt} decision variables are fixed, new lower and upper bounds are specified for the remaining variables, and integrality restrictions are restored before solving the problem as an MILP. This phase will attempt to either specify a specific value for each decision variable or use an updated upper / lower bound, resulting in a much smaller problem that needs to be solved. In the first instance, the fractional values of the decision variables found in the previous phase are considered. For each decision variable, if its fractional value is less than or equal to 0.05 then the value is fixed to the nearest (lower) integer. A similar rule is applied when the fractional value is greater than or equal to 0.95, however, this time the value is fixed to the higher integer value. For all the remaining decision variables, the heuristic simply sets new lower and upper bounds based on the values of the relaxed decision variables. Lower bounds are set to the nearest integer below the current value and upper bounds set to the nearest integer above the current value.

To illustrate this procedure, consider the following small example. For the first phase of the heuristic, assume that after relaxing the decision variables from integer to continuous values, the values given in the first row of Table 3.3 are obtained. Each of these values is then incremented by 0.5 as shown in row 2 of the Table, while row 3 shows the rounding applied. These new lower bounds are then used in the next iteration.

Current lower bound	0.25	1.78	2.01	1.34	0.78
Increment lower bound by 0.5	0.75	2.28	2.51	1.84	1.28
Round down to nearest integer (new lower bound)	0	2	2	1	1

Table 3.3: Tightening of lower bounds example (phase 1)

For the second phase, consider the values of the decision variables shown in Table 3.4. There are three possible cases at this stage. The first case is that the decision variables less than or equal to 0.05, in which case the decision variable is rounded down to the nearest integer (row 2). In the second case, the decision variable is greater than or equal to 0.95 and the decision variable is rounded up to the nearest integer (row 3). In the third and final case, the decision variable is in-between 0.05 and 0.95 and the lower and upper bounds are set to the nearest integer below and above, respectively (rows 4 and 5).

Decision variable	0.05	1.75	1.99	0.5	2.01	1.02	0.98
If $\leq .05$ (fix variable)	0	-	-	-	2	1	-
If $\geq .95$ (fix variable)	-	-	2	-	-	-	1
Update lower bound	-	1	-	0	-	-	-
Update upper bound	-	2	-	1	-	-	-

Table 3.4: Rounding of decision variables (phase 2)

All the major steps involved in the variable bounding heuristic are presented in Algorithm 3.1. The algorithm takes as an input the mixed integer linear program, which is then relaxed (lines 1 and 2). An infinite cost is applied to the current solution (line 3).

Algorithm 3.1 Variable Bound Heuristic

```
1. Input: $MILP$
2. $sol$ = relaxMILP($MILP$)
3. $sol_{cost} = \infty$
4. **While**($true$)
5.   $sol^{new}$ = solve($sol$)
6.   **If**($sol^{new}_{cost} \geq sol_{cost}$)    break
7.   **Else** $sol = sol^{new}$
8.   $sol_{LBs}$ = updateBounds($sol^{new}_{LBs}$)
9. **End**
10. $MILP_{new}$ = restoreIntegrality($sol$)
11. Solve($MILP_{new}$)
12. **Return** $MILP_{new}$
```

Next, lines 4 through 9 attempt to tighten the lower bounds. Specifically, line 5 solves the relaxed problem and records this in a new solution, lines 6 and 7 record any cost improvements, while line 8 updates the lower bounds of the solution based on the new solutions bounds. The steps carried out in the function $updateBounds()$ are those described above in phase one of this heuristic. These steps are repeated until the new solution's objective is equal to or greater than the previous solutions (line 6). Line 10 takes the relaxed LP and attempts to restore the integer decision variables, the second phase of the bounding heuristic described above. Line 11, simply attempts to solve the new MILP. Finally, line 12, returns the solved MILP and the algorithm is complete.

Summary

This chapter described and formulated the blood collection problem, specifically the objective and constraints were formulated and explained. The problem is a multi-objective problem, seeking to minimise collection cost and collection shortfall, though it was possible to formulate the problem with a single objective by weighting each of these priorities, with minimising collection shortfall being considered the more important of the two. This chapter also considers the model's limitation as well as the assumptions made during the formulation phase. The final part of this chapter discusses solution methods, a basic simulated annealing approach was first attempted though this did not prove fruitful as such a different heuristic was considered, the variable bounding heuristic. The next chapter introduces and describes the data set provided by the blood service and how the parameters of the blood collection model were populated with data.

Chapter 4

Current Practice

This chapter will explore how the dataset that is used throughout the remainder of this thesis was constructed. First, a general overview of the data available, as well as the format in which this data were provided is presented. Second, specific data regarding donors and venues are presented and discussed. Panels are not discussed in a dedicated section but are covered in the donors and venues sections. The final section of this chapter provides a brief summary.

The data used in this study were provided by the National Blood Service and covers the eastern county of Kent in England. Owing to the nature of the data, strict confidentiality rules were put in place. Accordingly, as such, it is not possible to publish the entirety of the dataset. Any data pertaining to donors must be treated in the same way medical records are treated and are, therefore, strictly confidential. However, other elements are subject to less strict controls, such as the price NHS hospitals buy units from the blood service (publicly available). Though as much detail as possible is given when discussing the dataset, there are limits imposed by confidentiality clauses.

The collection scheduling model proposed in the previous chapter requires data pertaining to four key areas: panels, donors, venues and general. The categories represent a logical separation of the data. Each of the categories can be considered individually, however, if needed, donors and venues can be linked through the use of the panel category. The database tables received from the NBS were split in a similar fashion. This explains why the dataset has been presented using these categories throughout the remainder of this thesis.

4.1 Overview

As mentioned above, the dataset for this project was obtained in part from the NBS, with any missing data being obtained from other reliable sources or estimated based on data provided by the NBS.

The data provided covered the months from December 2012 to October 2013. The data covers the English county of Kent. This region and time frame were selected for several reasons. Firstly, Kent includes both larger cities and more rural villages. Secondly, the period in question was an unexceptional period with no major events taking place in the region and covered multiple seasons. Finally, the dataset consists of over 140 panels in the region, which operate over 200 venues. The total demand for the period was 65,677 units.

The primary source of data for this project was a database provided by the NBS, called the PULSE database. The PULSE database was received in Microsoft Access 2003 format, which needed to be updated to work with more modern database systems. Although the format was changed, the data remained intact. Another source for data was face-to-face and email interviews that were held with NBS staff. These interviews were informal and were primarily to ensure the assumptions made, for both the model constraints and any missing data, were in line with what the NBS actually observe. The final source for data was government whitepapers and reports, as well as NBS annual and strategic reports that are all readily available to the public.

The PULSE database tracks important information pertaining to all the panels, venues, and donors the service is responsible for. The database also has modules that allow managers to run complex queries, create reports for performance analysis, and send out invitations to invite donors to upcoming sessions. However, the system does not track financial transactions, nor does it track inventory levels at local hospitals or regional blood centres.

NBS staff and have an intimate knowledge and understanding of the organisation's systems and processes. This is especially evident with regard to the scheduling of blood collections. As far as it is known, the collection scheduling process is not shared by the NBS. However, there is a "red book" that is published by the Department of Health that covers best practices with regards to transfusions. This book includes a chapter on selecting appropriate venues for holding collection sessions. It can be argued that expanding this chapter to include the entire collection scheduling process would be highly beneficial to the service.

Reliance on individual expertise does raise the question of how effective the NBS are with regard to knowledge transfer between employees and whether the systems in place are efficient. However, these questions are beyond the scope of this thesis.

Before computers were in widespread use the blood service stored details about donors on flashcards and grouped donors in boxes that represented the panel they belonged to (similar to an old library index system). Though the majority of this system has been digitised, the panel system remains in place. Venue costs are currently not accurately tracked on the PULSE database, and according to our NBS contact, are stored in a similar paper record fashion. This reliance on paper is not only an issue for the blood service but for the entire NHS. Medical records are still predominately kept in this fashion and all blood transfusion records are paper based and legally must be stored for over 10 years.

4.2 Donors

Donors are the most important part of any blood program. Without their continued donations, many life-saving hospital procedures could not be carried out. In the vast majority of countries, donors are volunteers, though some countries will pay people to donate blood. A donor will donate one unit of whole blood, approximately 430ml, which can later be separated into different blood products. Male donors are able to donate every 12 weeks in the UK, while female donors must wait at least 16 weeks between donations.

Blood services often differentiate between registered donors and walk-in donors. Registered donors are considered regular donors; they have previously registered with the blood service and have a donation history. These donors are, therefore, the main focus when developing a collection schedule. In contrast, walk-in donors are donors who are donating for the first time. Their exact numbers can be highly unpredictable. Walk-in donors can pre-register with the NBS and book to attend a specific session in advance. Despite being able to register, the model presented above does not distinguish between pre-registered walk-ins and normal walk-ins, in other words both will count towards the walk-in limit.

As well as walk-in donor numbers being highly unpredictable, the blood service does not know a walk-in donor's blood group in advance, regardless of whether he/she is pre-registered or not. This poses an operational challenge to the blood service; if too many walk-ins are accepted at a venue, then over-collection of specific blood groups can be high.

In view of the potential for blood born diseases, blood services collect and track large amounts of data regarding donors. Donor data is collected electronically and on paper, with every donor's donation, from collection and processing all the way through to the patient who receives the transfusion, being closely tracked. The donor dataset received from the NBS is, therefore, extensive. However, not all the data needed to accurately model the collection process was present. In the following sections, the data requirements are discussed before presenting the donor dataset received from the NBS.

A list of all the donors in Kent was requested and for each donor the following information was required. First, the panel the donor belonged to is needed, which would simply allow for grouping of donors into individual panels so that it is possible to know which donors are available to be called upon for donation at a given venue. A reliability score of some form was requested for either each individual donor or an entire panel. This score reflects how reliable a donor is. In other words, if a donor was called to attend a specific session, based on the donor's past donations, the score would determine how likely the donor is to attend the given session. This score could be a percentage or an ordinal score that could easily be transformed into a percentage. The score would be used in the α_i^m and α_i^f parameters of the blood collection model as described in Table 3.1.

The donor's previous donation was also requested, for two reasons. Firstly, to aide in verifying the accuracy of the reliability score and secondly to help in determining the donor's last donation date if this was not available elsewhere.

The donor's gender is important in determining how long the donor's donation cycle should be (12 weeks for males and 16 weeks for females, in the basic model). The donor's gender coupled with their last donation date before the start of the new planning horizon are essential for determining in which period the donor becomes eligible to donate again. These two pieces of data help determine the a_{igt}^m and a_{igt}^f parameters in the model described in Table 3.1.

The donor's first registration date was needed to establish whether the donor was a first time donor during the schedule under consideration, or a regular registered donor. If he/she was a first time donor, then the donor needed to be removed from the number of available donors and considered a walk-in donor for that venue. Finally, each registered donor's blood group was requested, in order to track how much blood of each blood group was being collected and whether demand targets had been met or not.

Table 4.1 shows an example of the format of the data, with randomly generated data used as placeholders. Each row in the donor database represents a unique donor.

ID	Blood Group	Reliability (1-5)	Panel	Last Donation
XA123456789P	A+	3	P123	01/01/2012
X7654321345P	O-	1	P042	01/01/2013

Table 4.1: Example donor table

Before any processing or transformations were applied, the database constituted just over 55,000 donors belonging to 160 panels. The most common reliability score for all the donors was grade 1 (reliable). Overall, more than 60,000 units were donated over the time frame, with some donors donating at least 4 times. It was thought to be of interest to compare the distribution of blood groups in our initial database against national standards. The results are summarised in Table 4.2. It can be seen that the distribution found in this dataset more or less conforms to the national trend.

The database contained elements that could be directly imported into the model with only minor formatting adjustments. These elements included the donor's unique ID, their blood group, and the panel to which he/she belongs. Other elements required slight modifications, such as the last donation date, which needed to be converted into a week number. The reliability score also needed to be transformed from a 1-5 scale number to a percentage. Other transformations were much more involved, such as determining whether a donor is a walk-in donor or calculating the total number of eligible donors. As such, these transformations are explained in more detail below.

Blood Group	National Distribution	Sample Distribution
O+	37%	36%
O-	7%	9%
A+	35%	34%
A-	7%	7%
B+	8%	8%
B-	2%	2%
AB+	3%	3%
AB-	1%	>1%

Table 4.2: Comparison of blood groups

The NBS reliability score is a 1-5 scale score that determines how likely a donor is to donate after being invited to attend a session. A score of 1 indicates the donor is very reliable, while at the opposite end a score of 5 signifies a highly unreliable and, therefore, unpredictable donor. As the model requires a percentage score for each panel, the following steps were taken. First, following discussion with the NBS, probabilities were attached to each score as shown in Table 4.3.

Secondly, to arrive at a panel score, the average percentage score across all donors was used for the entire panel. To ensure the reliability score was not unfairly skewed downward, all walk-in donors and donors that have not donated in a while were removed from the panel.

NBS Score	Probability of Attending
1	0.95
2	0.8
3	0.5
4	0.3
5	0.1

Table 4.3: Donor reliability scores and probabilities

The gender of each donor was not supplied and it is unknown how this is stored on the PULSE database. It is assumed that a gender column is present somewhere in the database. However, this column was not present in the dataset provided, despite being repeatedly requested. The reason for this omission is unclear, however, it is likely to be due to the fact that this information, coupled with the other details provided, could potentially lead to individual donors being identified. To overcome this problem, each donor was randomly assigned a gender in line with the overall gender distribution of donors nationally (approximately 60% female and 40% male). Though these ratios might not hold specifically for donors in Kent, no data was found to contradict this assumption.

Calculating when donors are next available to donate and, therefore, populating the a^{m}_{igt} and a^{f}_{igt} parameters, is a two stage process. First each donor's last donation date needed to be converted into a week number to be used within the model. The periods in the model started at week one and ran through to week 40, with any weeks before the start of the planning horizon having negative week numbers. This process was straight forward and involved extracting the donor's last donation date from the donor table and then converting it into a week number.

Once week numbers had been assigned to donor's last donation dates, the second stage of the process could begin. Donors were separated into three different groups: The first group was composed of donors who had last donated before the start of this planning horizon and at least more than 12 or 16 weeks before the start of this schedule, in other words, any (negative) week number less than -12 (for men) or -16 (for women). The second group included any donor who donated before the start of the current planning horizon but during the male/female waiting cycle, namely any donor who donated between weeks -12 or -16 and 0, respectively. Finally, the last group constituted all donors who donated during the current planning horizon, in other words donors who donated during weeks 1 through 40.

All donors who last donated during the first phase (weeks less than -12 or -16) are available to donate from the start of the new planning horizon (i.e. period 1), as they have waited at least the required cycle time for their respective gender.

Any donor who donated during the second phase could donate once he/she have waited for his/her cycle to complete. The most complicated donors to deal with are those donors who had donated in the final phase (i.e. during the planning horizon under consideration). These donors fell into one of two sub categories: either donors who have donated multiple times throughout the horizon or "new" donors. Any donor who had only donated once during the current horizon and had no previous donation records was considered new and was not represented in the a_{igt}^{m} or a_{igt}^{f} parameters, i.e. the donor was not considered a registered donor.

Donors who have donated multiple times over the horizon pose more of a challenge, as their last donation date cannot be utilised to determine when they become eligible again. Instead, every donor's donation was traced, with every venue code and session number being extracted. Afterwards, the dates for each session were retrieved from the session table and sorted into chronological order, before being converted into week numbers. Once this process had been completed, the donors could then be reclassified into one of the three phases. Any donors who were reclassified into phase 3, donating in the current planning horizon, were considered as walk-in donors. Those in phase 1 or phase 2 are treated as described above.

The dataset contained over 55,000 unique donors. However, some donors were assigned to panels for which no further information was provided. As such, these donors were reassigned to the next nearest panel. These "mystery" panels are often small with only a handful of donors. Furthermore, the neighbouring panels were only a short distance away.

The list of unique donors also contained numerous donors that have not been active for several years, some going as far back as 1999. Any donors who had not donated since the year 2004 were removed from the available donor list. The reason for this was that if these donors were invited to attend a session the chances of them actually showing up is extremely small. However, if these donors did decide to start re-donating, the model would capture them as "new" walk-in donors. This decision resulted in 700 donors being removed from the donor pool. In total, the NBS received donations from 15,600 walk-in donors during the 40 week collection schedule under consideration.

4.3 Venues

It is highly probable that donor sessions would be poorly attended if donors have to travel long distances to a few static facilities to make a donation. For this reason, blood sessions must be flexible and are usually held within local communities, close to the target population.

Given the medical nature of blood donation, these activities must be carried out at suitable sites, with specialist staff and machinery available. Fortunately, the staff and machinery for routine donation collections can be easily transported and with the use of antiseptic and sterile packaging suitable venues are plentiful. Furthermore, any potential site must have sufficient space to allow each of the donation stages (registration, interview and testing, donation, and, finally, recovery) to be carried out safely and efficiently.

To achieve these objectives, the NBS makes use of three different types of collection facility: temporary fixed venues, mobile venues, and finally, permanent or static venues, each of which are expanded next.

Temporary fixed venues sessions are usually held over a single day, which includes time for setting up and packing down. These sessions are regularly held at public venues, including churches, school halls, and community centres within a given panel.

Venues of this type offer the flexibility of being able to regularly visit multiple areas where donors are concentrated, without the overheads of owning a permanent facility. They are also useful in reaching donors in more remote regions. Moreover, these venues offer the size needed to run large donor sessions, usually resulting in higher throughput. The additional size of these venues also allows for each stage of the process to be clearly separated into areas.

The disadvantages of this type of venue are the necessity to transport all the staff and medical equipment by truck, the set up and pack down times, restrictions on hall availability and, finally hall hire costs.

Mobile venues are self-contained vans or trucks. These vans and trucks come in two sizes, large and small, which the blood service operates. These self-contained trucks allow for much more flexibility when selecting areas to visit. All that is required is a sufficient space for the mobile venue to be set up. Consequently, the schedulers often select busy car parks, shopping centres, universities, and high streets as venues. While a mobile venue could theoretically be set up and run anywhere in a panel, the NBS operate a list of approved sites.

Although a fixed temporary venue can process more donors than a mobile venue, mobile venues do offer a subtle advantage: they are a lot more visible to the public, with clear branding on the trucks, as opposed to a fixed temporary venue that might only have a small banner outside, making the trucks important marketing tools. However, mobile venues also have drawbacks other than the processing capacity mentioned before. Ease of getting in and out of the trucks can be an issue for elderly and disabled people. There is also the ongoing maintenance costs required to keep the trucks on the road. The most limiting factor is the fleet size and the capital required to add more vehicles.

Permanent venues can be stand-alone buildings, though they are most often located within an existing hospital complex. As the name implies, these venues are static and do not move around a region. These types of venues are frequently located in high density population areas, as they must have sufficient donors to justify the cost of operating every day. Permanent facilities allow the NBS to carry out highly specialised tasks, such as collecting platelets using apheresis or operating a cord blood bank. The model does not incorporate these types of venues for two reasons. Firstly, they are relatively rare, requiring extensive planning and capital to open new sites. Secondly, and more importantly, none of these venues are present in the region under consideration.

The NBS further differentiates sessions by holding either a private or a public session. Public sessions are sessions to which any donor within the panel can be invited and walk-ins are often allowed. Private sessions are normally held on the site for a particular organisation. Other donors may not attend these sessions and walk-ins are not permitted.

It is important to note that the blood service operate a list of locations at which a session can be held. Each location is assigned the type of session that can be held there: either fixed temporary or mobile and whether or not that session is public or private. The assigned type for a given venue cannot be changed or substituted for another type. As such, a fixed venue could not be sent to a mobile location for the obvious reason that there would be no venue to use. Nor could a mobile venue be sent to a fixed location, primarily due to the fact that the blood service cannot guarantee sufficient room outside the building for the trucks. However, it is assumed that a mobile location could accommodate any type of mobile vehicle assigned there.

In order to successfully model the venues the blood service operates, a full list of the locations at which a session could be held was required, as well as the type of venue that could be held there. Moreover, for each venue, the cost of operating the venue, the capacity of the venue, the number of expected or potential number of walk-in donors, and, finally, the number of times the venue could be visited each week was requested.

Three tables containing information regarding the blood service's venues were received. These included a table listing all the venues operating in the region, a table with all the sessions held over the planning horizon, and, finally, a table containing each donor that attended each session. In total, the blood service operates 200 venues in the Kent region, of which 56% are fixed venues. The service operate 5 large mobile vehicles and 7 smaller vehicles. Approximately 750 sessions were held over the 40 week period, resulting in over 60,000 units of blood being collected.

Finally, for each of the sessions that the blood service held, the number of units collected during the session, the date the session was held, and a list of donors that attended that session were provided.

Despite most of the venue data being provided, there were certain elements missing. The list of venues received included some venues that were assigned to panels for which no information was held. Some of these venues had postcodes outside of Kent and so could be safely removed, while others had Kent postcodes but no corresponding panel. In these cases, the postcode of the venue was extracted and nearby venues were identified. From this list, the closest venue's panel was assigned to each of the venues lacking an identifiable panel. This resulted in total 10 venues being reassigned to different panels, with 1-2 venues being added to certain panels.

Although costings for each venue were requested, this type of information is not tracked by the PULSE database. Instead, approximate cost distributions for fixed and mobile venues were derived. The distribution for fixed venues was a triangular distribution with a minimum cost of £127, a maximum cost of £670, and a most frequent cost of £300. Mobile venues also followed a triangular distribution with the following values: small mobiles had a minimum cost of £210, a maximum of £560, and a most likely cost of £355; large mobiles had a minimum of £313, a maximum of £660, and a most likely cost of £455. This data was used to determine values for the f_{ij}^{tv} and f_{ik}^{mv} parameters presented in Table 3.1.

The cost distributions presented above are used in the model to represent the fixed costs of setting up and running a venue at the given location. These costs include: the venue hire cost, the personnel costs, and transportation costs. After discussing this with the NBS, it was deemed that the distribution given would be fair representations of all the fixed costs involved.

In order to calculate the capacity of each venue, the maximum number of donors that had been processed at that facility over the planning horizon was determined. This is equivalent to how the blood service determines a venue's maximum throughput and should provide a relatively accurate capacity figure.

It was also essential to establish the number of times a venue could be visited throughout a period. This was done by examining the number of times a venue was visited each week in the dataset. It was found that private venues were visited no more than once in a period, i.e. the sessions were only held over one day. Therefore, all private venues in the model cannot be visited more than once in any period. This would appear logical, as these private venues often have a limited number of registered donors that can be invited and an employer is unlikely to host a blood drive over a prolonged period.

With reference to public sessions, the maximum number of sessions was determined in a similar fashion. However, it was necessary to exclude a couple of weeks from the analysis as these weeks contained unusual events, for example National Blood Week. With these weeks accounted for, it was determined that no venue is visited more than twice in a given week. Therefore, any public venue in the model could not be visited more than twice in a week.

The final piece of information that needed to be determined was the number of potential walk-ins that were available to each of the venues. This proved to be a particularly difficult task as the blood service did not provide the date a donor first registered with them anywhere in the dataset, only the date of their last donation.

In order to approximate the number of walk-in donors, all donors were extracted from the dataset and the total number of times they had donated over the planning horizon was calculated. Any donor that had donated more than once was considered to be a regular/registered donor. Any donor that had not donated was considered dormant and any donor that had only donated once was considered to be a walk-in donor. At this stage all donors were split into their respective panels. This enabled us to calculate the approximate fraction of walk-in donors in the entire panel. Overall, it was found that around a third of donors attending would be walk in donors, with the rest being registered donors. It was assumed that each venue within a panel would follow this proportion and from this the potential number of walk-ins available for each venue was calculated. This consisted of multiplying the ratio of walk-in donors by the average number of donors that had attended the venue over the planning horizon.

Furthermore, it was assumed that the fraction of walk-in donors available was constant over the entire planning horizon. In reality, however, the number is likely to vary. In particular, it is believed that weather could have a significant impact. Although further analysis in this area would be useful to the blood service, insufficient data currently being collected on walk-in donors limits what is possible.

One of the most important parameters to consider is the schedule duration. The National Blood Service utilises a 40 week time frame, for its planning horizon, which dates back several decades. Despite the current male and female donations intervals not evenly fitting into this time frame there are no plans to extend or reduce it. As such, the planning horizon can be more or less seen as a constant and therefore all analysis and discussion below are carried out with a 40 week planning horizon in mind.

Another parameter that can have a significant influence on the collection model is the number of crews available each week. Crews can vary in size depending on the number of donors that are expected to pass through the session. Although crew size may vary, the limiting factor is that a registered nurse must be present at all times to oversee the well-being of the donors.

Although the exact number of crews available in each period was not explicitly present in the dataset provided, it was possible to extract the number of crews each week based on the number of venues that were opened during the same week. From this it was established that the blood service seek to operate 20 collection sessions every week, meaning that 20 crews are available each week. However, it is not always possible to find sufficient staff to operate the desired 20 sessions a week. As such, some weeks may be understaffed and this is reflected in the dataset. The model was configured to reflect this fact. Most of the weeks had 20 sessions operating but some had just 17.

While no overall collection target data was provided by the National Blood Service, it was possible to extract target data for each session. Each session has an ideal number of whole blood units to be collected. Based on the planned and actual visits, it was possible to estimate each periods collection targets and, in turn, determine an overall collection target for the planning horizon. In total, the blood service sought to collect over 65,000 units of blood over a 40 week timeframe, which is slightly over 1,600 units each week. However, the blood service did not provide a breakdown of collection targets by blood group. This is important as ideally a patient will only receive a transfusion from a donor with the same blood group. In order to determine the quantity of blood to be collected for each of the eight main blood groups, the total demand was split following how prevalent the blood groups were throughout England. This resulted in the most common blood type (O+) requiring over 24,000 units with the least common (AB-) requiring only 650 units to be collected throughout the 40 weeks.

There are certain drawbacks to this approach. Firstly, it is assumed that the prevalence of the blood groups mimic the usage of each blood group at the hospital level. Secondly, the blood distributions used are national statistics and do not reflect regional differences. Despite these drawbacks, it is believed that this approach is the most accurate way to determine demand by blood group without being provided with additional data.

The variable cost for each unit collected was set at the cost the blood service sell every unit to the NHS. While normally this approach would also include a profit margin, the blood service is a non-profit government organisation and as such sells every unit at cost. It is possible, however, that this cost includes the staffing and venue costs, which may lead to double counting. On the other hand, the use of this cost ensures that all the variable costs needed are accounted for.

As for the penalty cost, the blood service currently do not directly factor in the cost for any shortfalls when considering collection schedules. However, there is a real cost associated with needing to source additional units to meet the demand. Consequently, after discussing this matter with the blood service, a suitable penalty charge was agreed upon - the blood collection model would charge £1,000 for every whole unit below the current demand levels.

As mentioned in earlier chapters, blood donors must wait a set period of time between donations, known as the donation interval. Therefore, this is a crucial parameter that will influence the number of donors available at one time and how many cycles they can get through in the planning horizon. The donation interval in England is set by the National Blood Service and there is currently a minimum interval of 12 weeks for men and 16 weeks for women. However, this is not consistent around the world, as no consensus within the medical community can be reached on the length of time needed for a donor's body to recuperate after donating. Since this model is exploring how the National Blood Service perform, it will initially use the current 12 and 16 week intervals, though later on in the analysis these will be modified.

The number of panels in operation is not fixed and is updated as needed, based on the recommendations of the schedulers. Throughout the course of the current planning schedule, however, the number of panels remains fixed, with any updates being made before the next schedule is produced. After preprocessing the data, a total of 144 panels were used.

4.4 Summary

This chapter has presented the dataset that will be used in the subsequent analysis. The dataset contains data pertaining to December 2012 to October

2013. In total the final dataset contains 200 venues, 144 panel, and over 55,000 unique donors of various blood groups. This chapter also considered how each of the parameters in the model were set, including the costs (c, f_{ij}^{tv}, f_{ik}^{mv} and w_g) and reliability parameters (α_i^m and α_i^f). It was found that it would be useful if the Blood Service recorded the date a donor first donated as this would aid in determining the number of walk-in donors. The following chapters make use of this dataset. Chapter 5 presents the computational results, first for a subset of the data and then the whole dataset, while Chapter 6 presents the analysis of several different scenarios.

Chapter 5

Computational Results

This chapter presents results and analysis obtained from the *blood collection model* described in Chapter 3. The first section presents computational results for both small test datasets and the much larger real dataset. The next section presents the NBS's current schedule and discusses how well collection targets were met and where improvements can be made. The third section analyses a schedule generated wholly by the *blood collection model* looking at how panels, venues, and donors are utilised. The final section of this chapter presents a summary of the differences between the current and proposed schedules.

5.1 Computational Results

There are a total of eight smaller test instances which enable the efficiency of each of the solution methods to be examined. All of these test instances are

Instance	Panels	Fixed	Mobile	Venue Sessions	Demand
Sub1	10	11	6	1	4,595
Sub2	10	11	6	2	4,595
Sub3	20	18	12	2	9,193
Sub4	20	18	12	3	9,193
Sub5	40	38	24	4	18,388
Sub6	40	38	24	5	18,388
Sub7	80	66	50	14	36,777
Sub8	80	66	50	11	36,777
Average	37.5	33.25	23	5.25	17,238.25

Table 5.1: Small data instances

subsets of the much larger dataset produced for Kent. The size of the test instances ranges from small instances with only 10 panels up to much larger instances consisting of 80 panels. Each of the instance's venues, demand, and the number of sessions that can be run each week are proportional to the number of panels. Finally, all datasets have a planning horizon of 40 weeks.

Table 5.1 provides a summary for each of the instances, while Table 5.2 presents computational results from the test problems for all three solution methods (CPLEX, a basic simulated annealing heuristic, and a variable bounding heuristic). All experiments were carried out on the same 2.8GHz quad-core (i5) laptop with 8Gb of RAM.

CPLEX performs the best on the smallest datasets, however, as problem sizes get larger, the optimality gap starts to grow significantly. For the 80 panel datasets, CPLEX could only produce solutions with a 53.72% to 62.32% gap relative to the LP lower bound. Whereas, the variable bounding heuristic performs very poorly on several of the smaller sets, as the problem size increase, this method outperforms the others. It would appear that problem Sub3 is particularly difficult to solve with all three solutions having a gap greater than 10%.

Overall, CPLEX has the worse average gap at 23.48% and the variable bounding heuristic performing the best with only a 9.12% gap. Looking at the simulated annealing heuristic, it generally under performs on these test datasets when compared to the variable bounding heuristic. This solution method is not considered further.

The computational results for both CPLEX and the variable bounding heuristic for the actual dataset are presented in Table 5.3. A total of nine scenarios are considered and are all based on the data provided by the National Blood Service. Each of the scenarios contained 144 panels and were able to choose from a total of 208 venues, which consisted of 117 temporary fixed venues and 91 mobile venues. Furthermore, for each of the mobile venues there is the choice of sending a large or small blood mobile to collect the blood. The planning horizon remained constant for all scenarios at a total of 40 weeks.

Considering the CPLEX results first, it is clear that these instances were extremely difficult to solve using a branch and bound solver. The average optimality gap was 76.4%, with the vast majority of solutions being closer to 86%. Curiously, CPLEX performed slightly better on two of the scenarios. The first of these was scenario 3, where it managed to produce a gap of 55.02%. On the other scenario, scenario 6, it performed even better managing an optimality gap of 28.37%. However, in each of these cases, the gaps are still far too large to be of much practical use.

	CPLEX				Simulated Annealing			Variable Bounding Heuristic		
Instance	Lower Bound(£)	Objective(£)	Gap(%)	Time(sec)	Objective(£)	Gap(%)	Time	Objective(£)	Gap(%)	Time(sec)
Sub1	614,274.07	614,335.69	0.01	16435.42	631,423.76	2.72	21600.2	770,633.00	20.29	5.55
Sub2	582,143.75	584,698.51	0.44	-	598,628.41	2.75	21600.2	596,704.00	2.44	18.01
Sub3	1,166,877.65	1,335,866.60	12.65	-	1,450,434.01	19.55	21600.2	1,528,035,52	23.64	39.74
Sub4	1,164,008.18	1,175,488.08	0.98	-	1,203,057.66	3.25	21600.2	1,216,181.30	4.29	7214.32
Sub5	2,329,858.48	4,275,100.20	45.50	-	3,028,052.42	23.06	21600.1	2,476,792.06	5.93	21600.2
Sub6	2,324,500.86	2,648,681.09	12.24	-	2,647,129.79	12.19	21600.1	2,496,141.14	6.88	21600.3
Sub7	4,653,397.70	12,350,357.09	62.32	-	5,734021.47	18.85	21600.3	4,820,523.14	3.47	21600.2
Sub8	4,653,364.74	10,055,453.88	53.72	-	5,482,142.33	15.12	21600.2	4,953,192.29	6.05	21600.3
Average	2,186,053.18	4,129,997.64	23.48	-	2,596,861.22	12.18	21600.2	2,357,275.31	9.12	11709.82

Table 5.2: Computational Results: small dataset

		CPLEX				Variable Bounding Heuristic		
Instance	Demand	Lower Bound(£)	Objective(£)	Gap(%)	Time	Objective(£)	Gap(%)	Time(sec)
Base	65,677	8,288,185.88	61,163,991.51	86.45	*	8,325,659.03	0.45	21600.3
Reliability Increase (20%)	65,677	8,288,618.50	61,550,821.35	86.53	*	8,318,740.87	0.36	21600.2
Pay Donors £20, 70% reliability	65,677	9,091,981.16	20,214,700.00	55.02	*	9,602,258.52	5.31	21600.2
Pay Donors £20, 90% reliability	65,677	9,162,673.55	61,893,700.00	85.20	*	9,193,600.55	0.34	21600.2
Changing Interval 10 male, 14 female	65,677	8,285,641.73	61,178,059.50	86.46	*	8,298,556.77	0.16	21600.2
Changing Interval 8 male, 12 female	65,677	8,282,514.77	11,562,210.82	28.37	*	8,572,701.13	3.39	21600.2
Changing Interval 10 male, 14 female, 5% extra demand	68,961	8,700,734.83	64,469,479.50	86.50	*	8,719,614.77	0.22	21600.1
Changing Interval 8 male, 12 female, 5% extra demand	68,961	8,697,255.69	64,167,112.23	86.45	*	8,755,830.96	0.67	21600.2
Changing Interval 10 male, 14 female, 20% extra demand	78,812	9,947,223.30	74,312,369.50	86.61	*	10,027,485.14	0.80	21600.3
Average	67,866	8,749,425.49	53,389,271.60	76.40	N/A	8,868,271.97	1.3	21600.2

* Indicates the 6-hour time limit was reached.

Table 5.3: Computational Results: actual dataset

The variable bounding heuristic achieves far superior results, with an average gap relative to CPLEX's lower bound of only 1.3%, with the vast majority of scenarios having gaps below 1%. The two under-performing scenarios are scenario 3 and 6 which had gaps of a 5.31% and 3.39%, respectively. Although these gaps are significant, the solutions are still useful for practical and analytical purposes.

As with CPLEX, all of the variable bounding heuristic results are reported after the heuristic had been run for 6 hours. Although the heuristic performs well on the majority of the scenarios above, the two scenarios with a larger gap suggest that care should be taken when applying this heuristic to other datasets.

Interestingly scenarios 3 and 6 happen to be the worst performing scenarios for the variable bounding heuristic and the best performing scenarios for CPLEX, though it is not clear why this is the case.

5.2 NBS Current Schedule

To allow for meaningful comparisons between each of the schedules generated by the *blood scheduling model,* it must be possible to compare each of the models and to have a standard to compare them against. All scenarios will be compared using the schedule's costs. Furthermore, using the schedule currently utilised by the blood service should provide a good reference point for comparison. This section describes two scenarios derived from the blood services current schedule, which will form the basis of comparison.

The first scenario is simply the blood service's actual schedule. This includes when the venues were visited, and how often, which panels were visited and the number of units collected in each period. However, as there is currently no cost for this schedule (the blood service do not assign costs to the schedules they operate), one needs to be calculated for comparisons to be made against the schedules produced by the *blood scheduling model*. The second scenario is similar to the first with an important difference - the number of units collected, and therefore, the donor assignments are ignored. Instead, the venues and panels visited are fixed in the model (the x_{ijt} and y_{ijkt} decision variables are set so that they match the blood services current schedule) and the model is then solved to determine the assignment of donors to open venues. These two scenarios allow comparisons to be made concerning how efficiently the blood service utilise venues and the available donor pool.

The data used in the scenarios described above is based on the actual schedule that the NBS ran between December 2012 and October 2013. The cost parameters employed were the same as the costs described in Chapter 4. As such, a unit of blood costs £124.34 to collect and a penalty of £1,000 is incurred for every unit below the collection target. Venue costs are specific to each venue.

The overall costs for these scenarios are presented next, before moving on to a more detailed examination of which venues are open and which panels are visited in the base models. Finally, a break down and comparison of how well demand targets are met and any shortfall costs are discussed.

Scenario Costs

The overall cost of a schedule is made up of three sub-costs. The first of these costs relates to the cost of opening and operating venues within the schedule. The second cost relates to the cost of bleeding donors and finally, the third cost is a penalty cost to discourage shortfalls. This cost structure forms the objective function 3.1 of the blood scheduling model presented in Chapter 3.

For all but the first scenario, determining the cost of the schedule is simply a case of reporting the objective produced by the model. However, owing to the fact that only the total number of units collected in each period was available, it was not possible to input the current NBS schedule's exact details into the model. Therefore, a different approach was taken, that of calculating each of the sub-costs as best as possible to approximate the schedule's actual cost.

Calculating the cost of opening the venues was straightforward and was just a case of taking the cost of the opened venue and multiplying it by the number of times the venue was visited. Calculating the cost of bleeding donors was also trivial and consisted of taking the total number of donors bled and multiplying them by the unit cost. Finally, although at first glance the cost of calculating the shortfall penalty appears straightforward, it was simply taking the number of units collected away from the total demand, this approach masks the possibility of over-collecting in one period and then under-collecting in another. The model assumes that units over-collected cannot count towards a future period.

Instead, to calculate the shortfall cost, the total number of donors invited in each period were distributed amongst the eight blood groups according to how frequently they appear in the UK population. The resulting table was then compared to the demand table for each period from which the actual shortfall cost could be calculated.

After calculating a cost for the blood service's current schedule (scenario 1) and running the set of venue openings through the allocation model (scenario 2), the two schedules had comparable costs attributed to them, both of which are presented in Table 5.4 along with a breakdown of each of their sub-costs.

	Scenario 1		Scenario 2	
	Cost(£)	%	Cost(£)	%
Objective:	13,410,514.05	100.00	8,664,787.19	100.00
Fixed Cost:	272,189.61	2.03	272,189.61	3.14
Variable Cost:	7,804,324.44	58.20	8,134,141.79	93.88
Penalty Cost:	5,334,000.00	39.78	258,455.80	2.98

Table 5.4: Cost breakdown for Scenario 1 and 2

It is immediately obvious from the table above that there is a significant difference in the cost of the two schedules. The current NBS schedule costs over four and a half million pounds more to operate than that proposed by the second scenario. Given that the only difference between the scenarios is how the donor pool is allocated to venues, this cost difference must be down to an inefficient allocation of donors.

The pair of scenarios have similar variable costs: £7,804,324.44 for the first scenario and £8,134,141.79 for the second. However, these costs represent notably different percentages of their overall costs. Scenario 2's variable costs account for over 90% of the schedule's cost, while the variable costs for scenario 1 represent just shy of 60% of the total cost. This percentage difference is easily explained by taking into account the percentage contribution of scenario 1's shortfall costs are vastly different. The current NBS schedule incurs a penalty cost more than 20 times that of scenario 2's shortfall cost. Suggesting that herein lays the primary reason for the differences between the two scenarios; the blood service is simply not meeting their demand target. Therefore, this is where the most significant cost saving opportunities for the blood service could occur.

Although the objective function provides a way to compare the quality of each schedule, further analysis is needed to better understand the exact causes behind the differences in schedule costs, specifically a more thorough look at the shortfall costs is required. However, before analysing the shortfall costs, it is important to gain an appreciation of the overall schedule, the venues and panels visited, which will be discussed next.

Venues

While constructing a schedule the blood service has a choice of over 200 venues to choose from that are made up of 117 fixed location and 91 mobile locations. In each 40 period time horizon the NBS can visit a maximum of 759 venues, owing to limitations on the number of venues that can be opened in a given period. This equates to an average, once rounded to the nearest whole number, of 19 venues being operated each period with the possibility of opening a maximum of 22 venues in periods 3, 7, and 24 and a minimum of 16 venues in periods 21 and 32.

Out of a possible 759 venues that could be opened, 758 are opened, comprised of 504 fixed venues and 254 mobile venues, making fixed venues (accounting for 66% of the total) the favoured venue type.

Despite the fact that there is no concrete reason for the blood service not having opened an additional session, there are several possible explanations which include adverse weather, making it impossible to get to the venue, staffing issues, or even that the venue was not in a functional state, e.g. ongoing building works or, one of the most common reasons for a session not being held, broken heating.

The following Figures 5.1 and 5.2 display the distribution of the total number of times each fixed venue (5.1) and each mobile venue (5.2) are visited in the current schedule. Slightly over 90% of the available fixed venues are visited resulting in only 11 never being visited during this planning horizon. A slightly higher usage rate is observed for mobile venues, with around 96% of the venues being used and only three venues never being visited.

It is observable on both charts that there are several frequently visited venues (visited more than 13 times for fixed venues and more than 10 times for mobile venues). For fixed venues, there is a greater variety in how often the venues are visited, while for mobile venues there is a distinction between the popularity of venues.

A large majority of fixed venues ($\sim 65\%$) are visited four times or less, similarly a sizeable portion of mobile venues ($\sim 60\%$) are only visited two times or less,

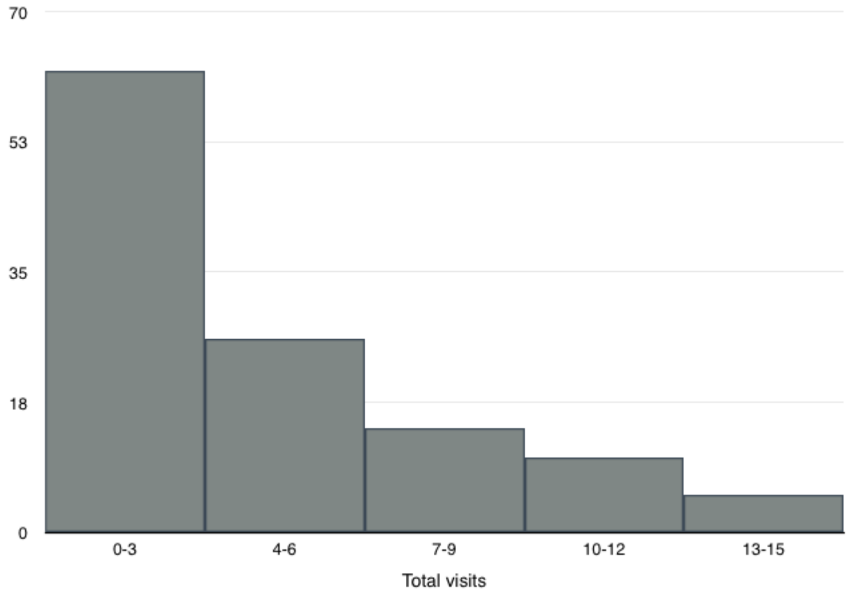

Figure 5.1: Frequency of temporary venue visits

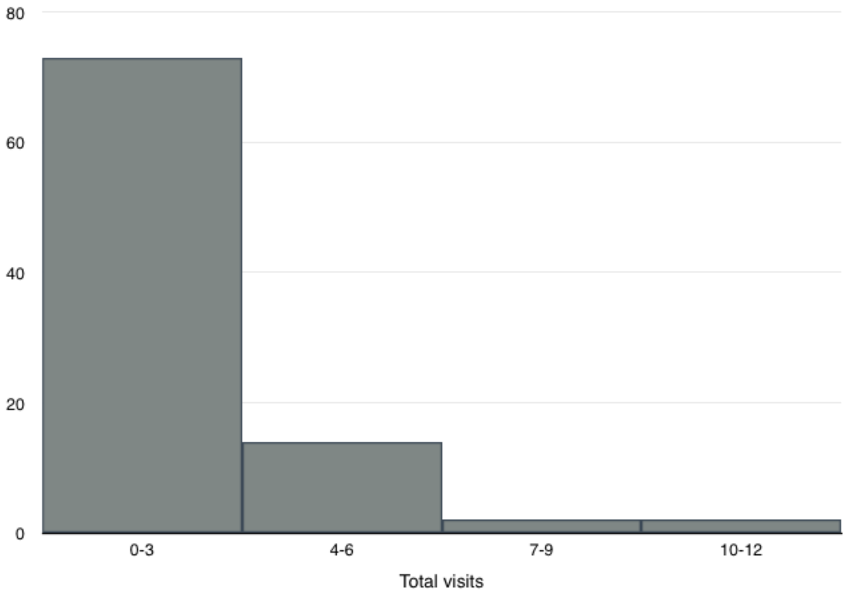

Figure 5.2: Frequency of mobile venue visits

indicating that the less popular venues are only visited sporadically throughout the schedule. It is important to note however that for mobile venues, especially those visited only twice, that they are most likely to be sessions held at small companies in the region. Furthermore, the most popular venues (those visited 10 times or more) make up slightly less than 13% of visits and appear to predominately be located in densely populated areas, with the two most popular venues, being visited 15 times each, located in major towns.

In total, venues visited at least ten times account for only 8% of the overall number of venues available, with the top fifteen fixed venues, with a combined total of 176 visits, accounting for 35% of all visits. Additionally the number of mobile venues, which are visited at least eight times, corresponds to 15% of the total number of mobile visits.

On average, a fixed venue is visited approximately 4 times during a schedule. This average drops for mobile venues which are only visited, on average, 3 times during the schedule. Furthermore, an average length of ten weeks passes between revisiting a given venue. Breaking this down, the most popular fixed venues are visited with an interval of between two and a half and four weeks, while the least popular venues, that are visited more than once, have a much longer interval of up to twenty weeks. Likewise, the most popular mobile venues are visited at least every month, with less frequently visited venues being visited on average at five month intervals. This indicates that the blood service seeks to visit densely populated areas at least every couple of weeks.

Although the majority of the most frequently visited fixed venues are located in densely populated areas, it is important to consider the capacity of each venue, to ensure that the blood service is not repeatedly revisiting these venues due to a lack of space. Figure 5.3 is a scatter graph comparing each of the venues' capacities to how frequently they are visited. While such a comparison would not be possible for mobile venues, as the blood service operates only small and large blood mobiles with fixed capacity, instead Table 5.5 shows the number of large and small blood mobiles sent to the most popular mobile sites.

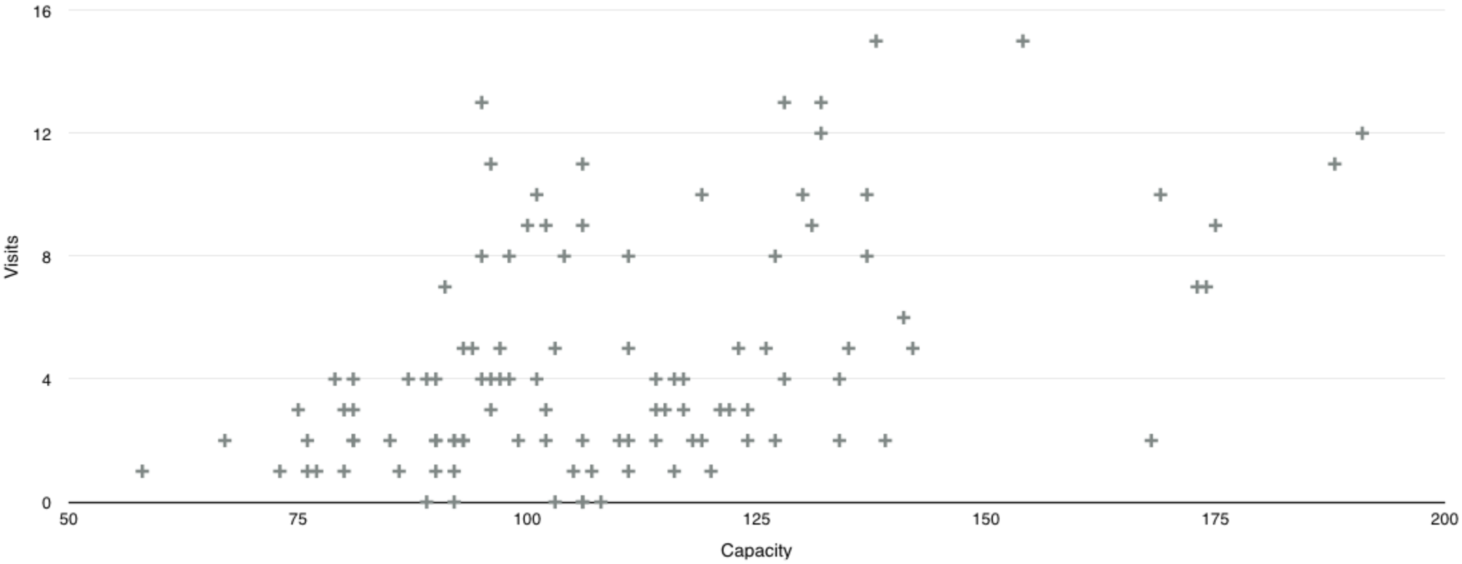

Figure 5.3: Venue capacity compared to number of visits

A fixed venue can on average process 110 donors in a session, with the largest session able to process 191 donors and the smallest session can only process 58 donors. Venues with the largest capacities are visited more than average with the exception of venue 42 which is only visited twice, perhaps owing to the fact that this venue only serves a sparsely populated area. However, the two most frequently visited sessions are not those with the highest capacities, though both are within the top 15.

Interestingly, a small group of venues with below average capacity (between 95 and 106) are visited ten times or more. These venues are located within major towns in the region, suggesting that the blood service should look for venues with greater capacity in these areas. Furthermore, this suggests that a venue's capacity is not always directly linked to the area's population. This is further evident when considering the capacity (168 donors) and population (approximately 7,000 people) served by venue 42.

The scatter graph further highlights the popularity of certain venues, with popular venues being visited at least seven times. However, their capacities range significantly from 91 to 191. The chart also highlights that the majority of venues are visited no more than six times in the planning horizon and their capacities also vary significantly, although there is a greater concentration of venues around the 100 mark.

It could be argued that Figure 5.3 also indicates that the blood service has too many venues to choose from, especially when considering the fact that certain venues are rarely, and some even never, visited. This suggests that the blood service may benefit from seeking venues with higher capacities in densely populated areas and visiting them more frequently.

As for mobile venues, Table 5.5, indicates that the five most frequently visited venues are most often visited by a mixture of both large and small blood mobiles.

Significantly, all of these mobile venues are located in smaller villages adjacent to larger urban populations.

Venue ID	Large mobile vehicle sent	Small mobile vehicle sent	Total visits
84	6	5	11
60	10	-	10
13	4	5	9
89	6	2	8
66	-	6	6
70	3	3	6

Table 5.5: Mobile venues utilised

For venue 66, it is noteworthy that although all of the blood mobiles sent were small, on each occasion the venue was visited twice in each of the periods. For each of these periods there were several large blood mobiles still available. Similarly, for the other venues that received a mixture of large and small blood mobiles, there was very often large blood mobiles available. These points illustrate the fact that the blood service could benefit from more efficiently assigning blood mobiles to venues, especially as one large blood mobile has the same capacity as sending two smaller vehicles.

Panels

This subsection explores which, and how frequently, panels are visited in the current NBS schedule. In total, there are over 140 panels for the blood service to choose from, with each panel containing at least one venue and at most having 6 venues.

Figure 5.4 shows the frequency with which panels are visited, taking both fixed

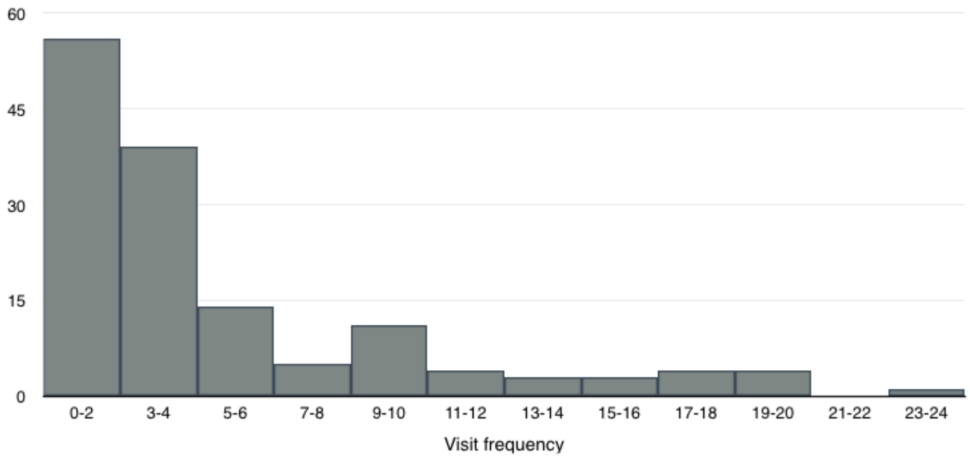

Figure 5.4: Frequency of panel visits

and mobile visits into account, throughout the planning horizon. Every panel is visited at least once, while there is no reason for this occurrence in the dataset provided nor in the discussions held with the blood service, it is possible that this is either a marketing or perceived fairness requirement. The rationale from a marketing perspective would be to enable the blood service to be seen as actively collecting blood from all over the region. From a fairness perspective, it would ensure that the donors, specifically those in remote villages, are not seen as less important. In the formulation of the blood scheduling model there is no constraint to enforce such behaviour, owing to the fact that the constraint is not an operational restriction.

On average, each panel is visited slightly over five times, with the most prevalent panel being visited a total of 23 times throughout the schedule. The large majority of panels are visited less than six times with most panels being visited only twice. Interestingly, panels that are visited once are in the minority with only five panels being visited a single time. Clearly the blood service aims to visit all available panels at least twice. However, panels visited more than ten times make up just fewer than 20% of all visits, with those being visited ten times being the most common in this category.

The most commonly visited panels are located near major towns and cities in the region, although the most frequently visited panel does not include the most populated city in the region. All of these panels often have multiple fixed venue sites. Panels that are visited less than four times in the schedule cover either small villages or hospitals, shopping centres and places of work, where there is a much smaller registered donor base. These panels often do not have any fixed venues in them and consequently are visited by mobile venues.

Despite a panel being frequently visited it is important to verify that this is not caused by the venues in the panel having insufficient capacity. Therefore Table 5.6 presents the most often visited panels, specifically those visited more than 16 times, alongside the average capacity of the venues frequented.

Panel	Times visited	Average venue capacity
P005	23	104
P142	20	122
P050	19	105
P073	19	128
P170	19	156
P030	18	106
P053	17	95
P104	17	151
P151	17	128
P071	16	111

Table 5.6: Average venue capacity for popular panels

All but four of the panels have an average capacity greater than the mean capacity for all venues, of those panels one in particular is well below the average, panel 14. The venues in this panel are the same venues discussed in the previous subsection, those venues with below average capacity and frequently visited. This further suggests that the blood service should be looking for venues with larger capacity served by this panel. This suggests that a few of the panels are frequently visited due to a lack of capacity in that panel's venues, though this is clearly not the case for all of these panels.

It would be of interest to compare the contribution of each panel to the overall collection target and to contrast these results with how frequently the panel is visited. Table 5.7 presents all the panels that fulfil over 1% of the total demand. The average contribution for all of the panels is only 0.69% or around 430 units of blood, with the lowest panel only providing 13 units of blood, from two visits. This is in stark contrast to the highest contributing panel which provides over 2,100 units of blood accounting for just shy of 3.4% of demand.

Unsurprisingly, the most frequently visited panels are all present in Figure 5.5. However, there is one panel which is in the top three of Table 5.7 namely panel P167 that is in fact only the 17^{th} most often visited panel. This is easily explained by the fact that the venue in this panel has the highest capacity of all other venues. Subsequently, this further highlights the fact that the number of times a given panel is visited does not directly correspond to their contribution i.e. the panel visited the most does not contribute the most.

Surprisingly, the contribution of each individual panel is remarkably low with

Panel	Contribution	Panel (contd.)	Contribution
P142	3.39	P015	1.82
P005	3.23	P016	1.81
P167	3.15	P132	1.77
P170	3.11	P088	1.73
P104	3.06	P187	1.72
P151	2.98	P147	1.71
P073	2.97	P171	1.67
P030	2.75	P024	1.60
P050	2.56	P051	1.41
P053	2.28	P057	1.35
P138	2.26	P133	1.35
P076	2.23	P148	1.35
P071	2.22	P063	1.22
P068	2.14	PV61	1.21
P048	1.92	P143	1.10

Table 5.7: Panel contribution (only those panels with ≥1% shown)

each panel accounting for less than four percent of demand. However, when considering the accumulation of units, it quickly becomes apparent that the contribution of each of the panels approximates a logarithmic curve, as show in Figure 5.5. Remarkably, only 21 panels are needed to satisfy fractionally over half of the total demand, moreover only 82 panels, just over half of all available panels, are needed to collect 90% of the demand. Looked at from the opposite end, the final 40 panels only contribute around 5% of demand or around 3,000 units. This suggests that if the blood service could increase donor attendance at the first panels then the total number of panels operated could be significantly reduced.

To summarise the blood service does not use all of the venues it has available and of the ones that are utilised there are venues that are visited significantly more frequently than others. The venues with a higher capacity appear to be

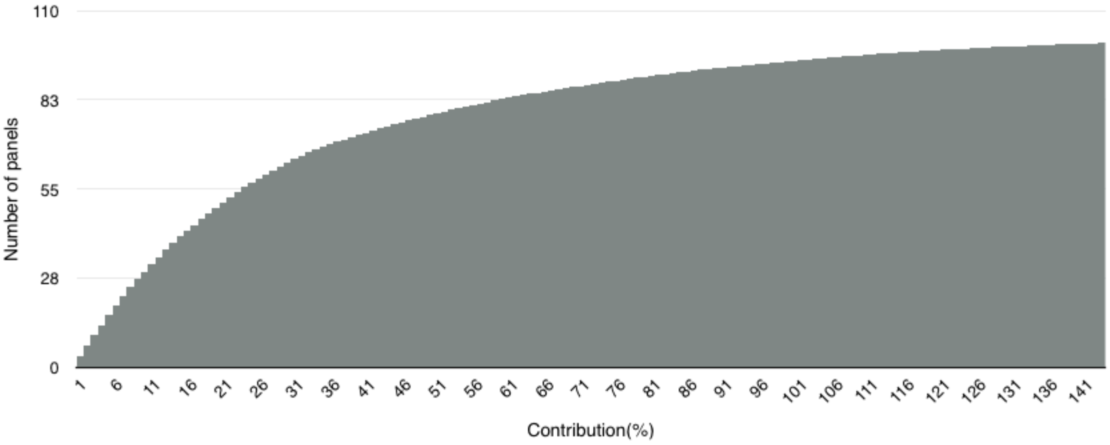

Figure 5.5: Cumulative unit contribution for number of panels visited

visited more often. As for panels, the service seeks to visit all of them at least once, but again there appears to be a significant difference in how frequently they are visited; namely those around densely populated areas are visited more often. Having examined how the blood service currently use the venues and panels available to them, the next subsection focuses on how the service assigned donors to the opened venues.

Donors

This subsection will look at two points pertaining to stratifying the schedule's demand targets. The first point is a reminder of what and how the demand targets were determined. The second analyses how many units were collected and how well the demand was satisfied under scenario 1 and scenario 2.

Understanding when, of what type and how many units of blood the blood service collects is vital as it cannot only have a major impact on how the healthcare service is run but also a huge impact on the final cost of the collection schedule, predominately affecting the penalty cost.

The number of units collected for each period in scenario 1 was determined by summing all of the units collected from each of the venues visited in that period. In total the blood service were able to collect 62,766 units over the 40 week planning horizon. Instantly noticeable is the fact that the units collected falls short of the service's demand target by over 2,900 units. However, a detailed analysis is limited by the fact that a full break down of donors was not provided. Specifically, it is not possible to distinguish between male and female donor, as donor's gender was not provided in the dataset, nor is it possible to precisely distinguish between registered and walk-in donors attending each venue. Furthermore, and most importantly, it is impossible to determine how many donors were actually invited; only information on the number that actually attended was provided.

Comparing the differences between the units collected and the demand produces Table 5.8. In this table any positive numbers represent a shortfall to the desired collection target while any negative numbers signify that an excess number of units were collected. Out of the 40 periods in the horizon just over a quarter, 11 periods, over-collected blood, the remaining 29 periods all under-collected blood. No periods were able to exactly meet the collection targets, highlighting both the difficulties in predicting donor behaviour and at the same time the need for better donor management.

Period	O+	O-	A+	A-	B+	B-	AB+	AB-	Shortfall
1	113.22	21.42	107.10	21.42	24.48	6.12	9.18	3.06	306
2	71.41	13.51	67.55	13.51	15.44	3.86	5.79	1.93	193
3	10.36	1.96	9.80	1.96	2.24	0.56	0.84	0.28	28
4	95.46	18.06	90.30	18.06	20.64	5.16	7.74	2.58	258
5	23.68	4.48	22.40	4.48	5.12	1.28	1.92	0.64	64
6	-48.10	-9.10	-45.50	-9.10	-10.40	-2.60	-3.90	-1.30	-
7	29.60	5.60	28.00	5.60	6.40	1.60	2.40	0.80	80
8	-18.13	-3.43	-17.15	-3.43	-3.92	-0.98	-1.47	-0.49	-
9	53.65	10.15	50.75	10.15	11.60	2.90	4.35	1.45	145
10	83.62	15.82	79.10	15.82	18.08	4.52	6.78	2.26	226
11	98.05	18.55	92.75	18.55	21.20	5.30	7.95	2.65	265
12	92.87	17.57	87.85	17.57	20.08	5.02	7.53	2.51	251
13	-101.75	-19.25	-96.25	-19.25	-22.00	-5.50	-8.25	-2.75	-
14	256.78	48.58	242.90	48.58	55.52	13.88	20.82	6.94	694
15	24.79	4.69	23.45	4.69	5.36	1.34	2.01	0.67	67
16	9.25	1.75	8.75	1.75	2.00	0.50	0.75	0.25	25
17	46.62	8.82	44.10	8.82	10.08	2.52	3.78	1.26	126
18	-148.00	-28.00	-140.00	-28.00	-32.00	-8.00	-12.00	-4.00	-
19	116.92	22.12	110.60	22.12	25.28	6.32	9.48	3.16	316
20	9.62	1.82	9.10	1.82	2.08	0.52	0.78	0.26	26
21	-141.71	-26.81	-134.05	-26.81	-30.64	-7.66	-11.49	-3.83	-
22	117.29	22.19	110.95	22.19	25.36	6.34	9.51	3.17	317
23	113.59	21.49	107.45	21.49	24.56	6.14	9.21	3.07	307
24	-125.06	-23.66	-118.30	-23.66	-27.04	-6.76	-10.14	-3.38	-
25	34.04	6.44	32.20	6.44	7.36	1.84	2.76	0.92	92
26	115.81	21.91	109.55	21.91	25.04	6.26	9.39	3.13	313
27	-39.22	-7.42	-37.10	-7.42	-8.48	-2.12	-3.18	-1.06	-
28	47.36	8.96	44.80	8.96	10.24	2.56	3.84	1.28	128
29	-5.92	-1.12	-5.60	-1.12	-1.28	-0.32	-0.48	-0.16	-
30	162.43	30.73	153.65	30.73	35.12	8.78	13.17	4.39	439
31	16.28	3.08	15.40	3.08	3.52	0.88	1.32	0.44	44
32	7.03	1.33	6.65	1.33	1.52	0.38	0.57	0.19	19
33	3.70	0.70	3.50	0.70	0.80	0.20	0.30	0.10	10
34	-101.75	-19.25	-96.25	-19.25	-22.00	-5.50	-8.25	-2.75	-
35	149.48	28.28	141.40	28.28	32.32	8.08	12.12	4.04	404
36	15.17	2.87	14.35	2.87	3.28	0.82	1.23	0.41	41
37	-32.19	-6.09	-30.45	-6.09	-6.96	-1.74	-2.61	-0.87	-
38	15.17	2.87	14.35	2.87	3.28	0.82	1.23	0.41	41
39	40.33	7.63	38.15	7.63	8.72	2.18	3.27	1.09	109
40	-134.68	-25.48	-127.40	-25.48	-29.12	-7.28	-10.92	-3.64	-
Total*	1,973.58	373.38	1,866.9	373.38	426.72	106.68	160.02	53.44	5,334

*Total does not include units over collected.

Table 5.8: Shortfall by blood group and period

The number of units over-collected throughout the schedule total 2,423 units, with the most units being over-collected in period 18 (400 units) and the least in period 29 (16 units). Disregarding the number of units over-collected the total shortfall for scenario 1 is 5,334 units, over 2,434 units more than the initial 2,900 unit estimate. Although the units that are over collected would not be discarded by the blood service, as they can still be transfused to patients, their presence makes the task of inventory management much more complex. This complication is due to the fact that the unit's shelf life would be different to a newly collected unit, potentially leading to a much higher wastage rate.

In spite of over collecting, the majority of periods do not meet their demand targets, with period 14 being over 600 units short. Surprisingly, even the first few periods in the schedule missed their targets by just under 850 units. Looking at the shortfall column in Table 5.9, there generally appears to be a series of periods that significantly miss targets followed by a single period that over collects units, though the number of units over collected is often only a fraction of the previous period's shortfall. These periods of under collecting followed by these sudden peaks in attendance further highlight the fact that the blood service has encountered issues when anticipating how many donors will attend a given session. This is potentially a result of how the blood service invites the majority of eligible donors in a panel to attend a nearby session rather than specifically targeting donors by their blood groups. Indeed, the session scheduling model does a more effective job of allocating donors, as evident by the drastically reduced shortfall penalty incurred (only £258,000 compared to

Blood Group	Period	Shortfall (units)
B+	14	62.96
A+	14	60.39
A-	14	42.41
O-	35	23.31
AB+	14	22.36
AB+	35	14.58
AB-	35	10.07
AB-	14	9.88
B-	14	8.44
B-	35	4.05

Table 5.9: Scenario 2 shortfall

the previous £5.3 million) for scenario 2.

Table 5.9 presents the shortfall scenario 2 incurred by blood group and period. Over 258 units of shortfall were incurred in this scenario, a reduction of over 5,100 units on scenario 1. Moreover, there were no over collected units throughout the entire schedule. Of the total shortfall, just over 200 units were incurred in period 14 while 52 were incurred in period 35. Intriguingly, period 14 is the period with the highest shortfall in scenario 1, while period 35 is the period with the third largest shortfall. This suggests that either the collection target for these periods is too high considering the number of donors available or that the venues and panels visited up to and during these periods is suboptimal. Furthermore, scenario 2 is able to collect sufficient units to satisfy the demand for O+ blood for each of the 40 periods, with each of the remaining blood groups averaging a 37 unit shortfall.

For scenario 2, unlike for scenario 1, it is possible to give a detailed donor break down. In total 66,083.02 registered donors were invited of which 47,820.67 turned up equating to 7.2 registered donors showing up for every 10 invited. Of the registered donors invited, 25,406.95 were male and 22,413.72 were female, with the remaining 17,597.88 all being considered walk-in donors.

Owing to the significant reduction in demand targets missed in scenario 2, there is a large decrease in penalty cost compared to scenario 1. Furthermore, the higher variable cost seen in scenario 2 is easily explained by the increase in donors bled throughout the session, though this time no donor to over-collecting occurs, as seen in scenario 1.

This section explored in detail the actual NBS schedule used between the end of December and October 2013. It was found that the blood service seeks to visit every panel at least once and that certain venues are visited more frequently than others. Currently, donor management appears to be poor with all periods experiencing either over- or under-collecting of blood.

Throughout this section, the NBS schedule was compared to a second scenario, which differed in how donors were assigned. The cost produced by this second scenario was significantly cheaper than the actual blood service's schedule. This was due to the more efficient handling of donors. Overall, the improved donor assignment resulted in a saving of over £4.5 million pounds compared to the actual schedule.

Now that the current schedule has been analysed, the next section introduces the third scenario, in which the model is used to determine which venues and panels are visited as well as managing donor assignments.

5.3 Proposed Schedule

This section presents results from a third scenario in which both venue opening and donor allocation decisions are made simultaneously. This scenario will be useful to see how efficiently the blood service uses the venues available.

To allow for a fair comparison between the schedule produced by this scenario and the baseline schedules discussed previously, all costs are kept identical, as are the demand targets, the number of donors available, the likelihood of donors attending, and the proportion of walk-in donors to registered donors attending. This also applies to the number of venues and panels available and the number of mobile vehicles and venues that can be opened in each period.

The overall cost for the schedule produced by scenario 3 is £8,325,659.03, with more detail regarding venue, donor, and penalty cost being presented in Table 5.10.

	Scenario 1		Scenario 2		New schedule	
	Cost(£)	%	Cost(£)	%	Cost(£)	%
Objective:	13,410,514.05	100.00	8,664,787.19	100.00	8,235,659.03	100.00
Fixed Cost:	272,189.61	2.03	272,189.61	3.14	135,953.07	1.63
Variable Cost:	7,804,324.44	58.20	8,134,141.79	93.88	8,162,954,41	98.05
Penalty Cost:	5,334,000.00	39.78	258,455.80	2.98	26,751.54	0.32

Table 5.10: Schedule cost comparisons

Unsurprisingly, this new schedule is significantly better than scenario 1, primarily due to better donor assignment. Scenario 3 also presents a saving of approximately £300,000 when compared to scenario 2's total cost (£8,664,787.19).

At first glance the major cost difference between both scenarios is in how much of a penalty each scenario acquires. Scenario 2 incurs a penalty of £258,455.80, while scenario 3 has a penalty of only £26,751.54 - less than a percent of the overall schedule cost. This meaningful reduction signals that scenario 3 is even better at allocating donors and that the total shortfall has dropped from around 258 units to only 26 units - a reduction of nearly 90%. As a consequence of the reduction in shortfall costs the total variable costs have increased marginally.

Scenario 3 also achieves a further reduction of £136,236.54 compared to the previous scenarios' fixed costs. In percentage terms this represents a reduction of more than 50%, taking the total venue cost for scenario 3 to £135,953.07. The new venue cost makes up only 1.63% of the total scenario cost down from 3.14% in scenario 2.

Although the overall savings achieved by scenario 3 equates to only 3.91% compared to scenario 2, any savings that can be made by the NHS should be encouraged, especially considering the current austerity measures in place. Furthermore, these savings can be used to fund research projects or could even be used to fund, at least, an additional seventeen nurses. Moreover, the considerable reduction in the number of units below the collection target levels, indicate that scenario 3 is of a much better quality from a supply chain perspective - the closer the blood service are to meeting their demand targets the less pressure there will be on the healthcare system over the planning horizon.

Venues

In scenario 3's schedule a total of 629 venues are visited, interestingly all of which are fixed venues - no mobile venues are ever visited throughout this schedule. This is in stark contrast to the current schedule used by the blood service where 34% of the venues visited were mobile and all but one of the available venue sessions were used.

Table 5.11 presents the number of venue sessions filled in each period as well as the slack available and whether or not the collection target was met in a given period. In total only 629 of the 759 venue sessions available were filled resulting in a total slack of 130 venue spaces. The average utilisation rate of the venue's sessions over all periods is over 83%, though only 4 periods have no slack at all. At the lower end of the utilisation rates is period 40 that only uses half of the available sessions, with period 24 only using slightly more. This potentially suggests that the blood service currently operates too many sessions in certain periods and that reductions can be made. This goes some way to explaining the significant reduction in fixed costs observed between the blood services schedule and this one.

In spite of the reduction in venues used, this schedule still incurs shortfalls in various periods, notable in period 14 (the same period that has a shortfall in scenario 2). Interestingly, in each such period, there is still a venue slot available, indicating that the cause of the shortfall in these periods is due to a lack of available donors rather than a venue availability issue. As no mobile

Period	Crews available	Venues open	Slack	Utilisation(%)
1	17	14	3	82
2	20	18	2	90
3	22	14	8	63
4	18	14	4	77
5	18	18	-	100
6	19	16	3	84
7	22	21	1	95
8	18	14	4	77
9	19	19	-	100
10	18	16	2	68
11	17	16	1	94
12	18	16	2	88
13	19	13	6	68
14	17	16	1	94
15	20	17	3	85
16	20	16	4	80
17	19	16	3	84
18	20	14	6	70
19	18	15	3	83
20	20	16	4	80
21	16	11	5	68
22	19	17	2	89
23	20	17	3	85
24	22	12	10	54
25	20	16	4	80
26	18	16	2	88
27	17	13	4	76
28	18	16	2	88
29	21	18	3	85
30	20	20	-	100
31	19	19	-	100
32	16	14	2	87
33	20	18	2	90
34	19	12	7	63
35	20	19	1	95
36	20	19	1	95
37	19	16	3	84
38	18	14	4	77
39	18	13	5	72
40	20	10	10	50
Total:	759	629	130	83

Table 5.11: Crew utilisation

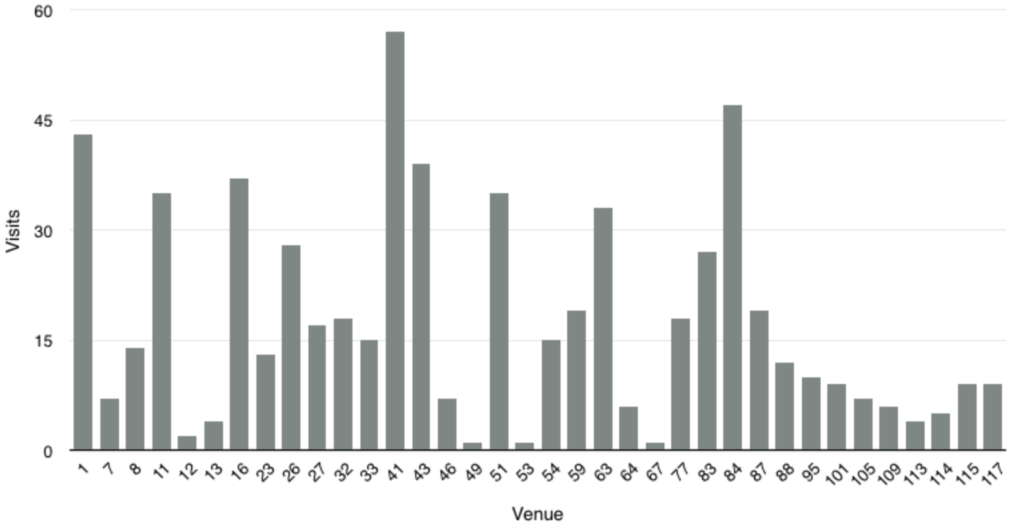

Figure 5.6: Popular temporary venues

venues are visited in scenario 3's schedule, only the frequency of visiting fixed venues are considered. The total number of visits for all fixed venues visited at least once is shown in Figure 5.6, while Figure 5.7 shows the distribution of the frequency of the visits.

Schedule 3 visits certain venues decidedly more frequently than the blood service's schedule, with more venues being visited more than 26 times. The proposed schedule also visits far less venues, in total only 36 out of the 117 available fixed venues are visited, far less than the 105 visited by the blood service. Similarly, an open venue in scenario 3 is visited on average 17.4 times while in the base scenarios a venue is only visited on average less than twice (for both fixed and mobile venues). Furthermore, the most frequently visited venues in scenario 3 are visited at least 25 times while in the base scenario this is only 15 times - or about 1.7 times less often. Naturally, as the most popular fixed venues

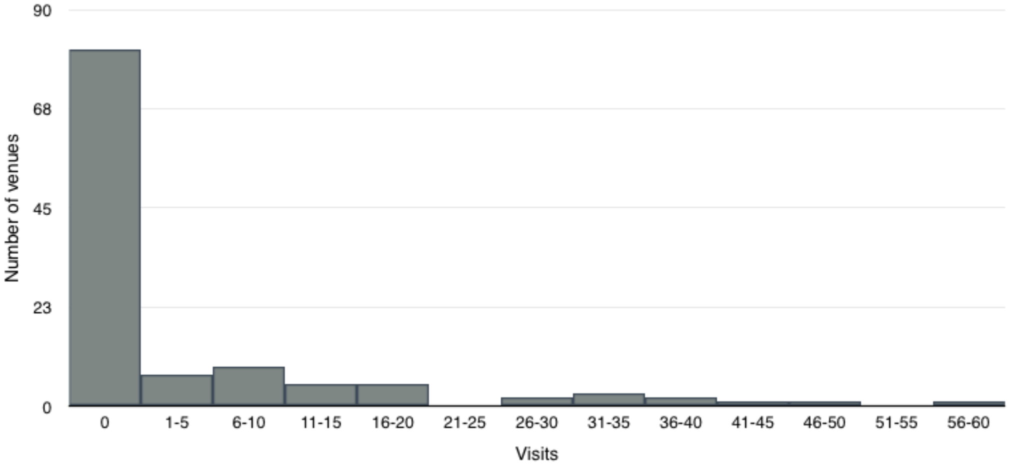

Figure 5.7: Frequency of temporary venue visits

are visited more frequently, the number of weeks between each visit has fallen considerably from an average of two and a half weeks in the base schedule to just under every week in the current one.

As with the base scenario, all of the most popular venues are located in densely populated areas, although not all of the most popular venues in the base schedule appear in the new schedule. Out of the 15 most popular venues in the base schedule only six appear in the new one.

Interestingly, the most visited venue in the current schedule, venue 41, which is visited a total of 57 times throughout the schedule, is only visited twice in previous schedules. Other venues that are visited significantly more times in the new schedule compared to the previous one include venues, 51, 16, 11 and 1 to name a few. Curiously, venue 43 is the fourth most visited venue in the new schedule but is not visited once in the blood services schedule. There are

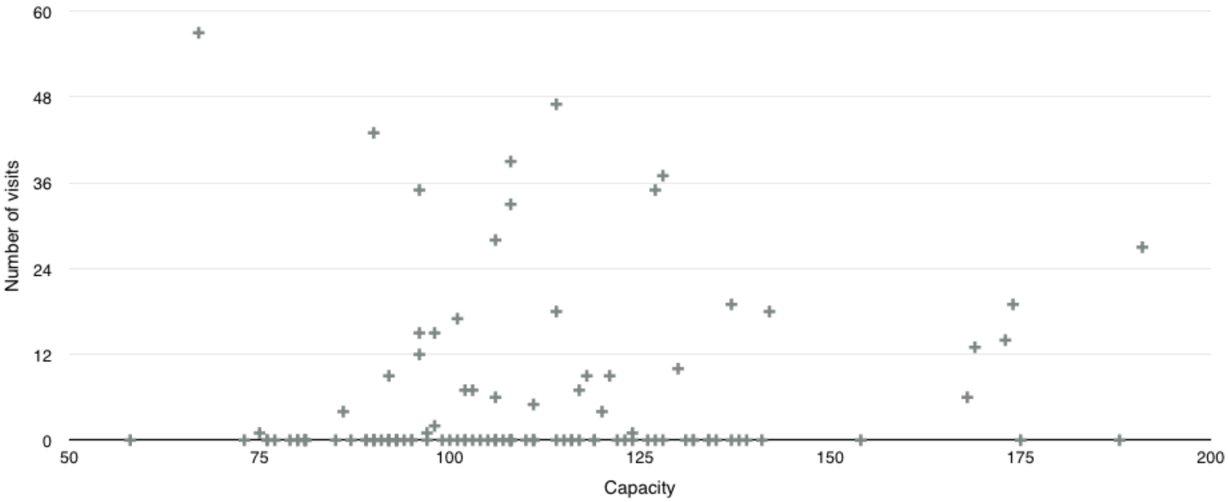

Figure 5.8: Venue capacity compared to number of visits

only two other venues - venue 63 and 64 - that only appear in the new schedule and not the base ones. Notably only one venue in the new schedule is visited exactly the same number of times as in the previous schedule - venue 95 which is visited a total of 10 times and that serves a medium sized town.

To analyse the capacity of the venues and how this relates to how often they are visited a scatter chart was produced to map these two variables. Looking at Figure 5.8 it appears that the venues are either visited less than 20 times or more than 30 times, with only two venues being visited between these limits.

The eight most frequently visited venues (visited more than 30 times each) have an average capacity of around 105 donors. Venue 41 however appears to be a noticeable outlier as this venue has the second lowest capacity - can only process 67 donors, though is by far the most visited venue in this schedule. Furthermore, venues that can handle over 150 donors are visited no more than

27 times, although for the majority of them this is far more than they were visited in the base scenario. Interestingly, three of the higher capacity venues are never visited, including one of the most frequently visited venue in the base scenarios. The other two venues 86 and 66 were visited 11 and 9 times respectively in the base scenario and all three of these venues serve populated areas.

Clearly a venue's capacity is not the main driving force behind how frequently that venue is visited; to understand why some of the lower capacity venues are visited so often, the venue cost to venue capacity ratio needs to be examined.

The venue cost to capacity ratio is simply the cost of utilising a venue divided by its capacity. This ratio also allows for easy comparisons to be made across all of the different venues on a per donor basis. The average cost for all fixed venues is £2.87 per donor and for all mobile venues, assuming only the smaller blood mobiles with a capacity of 45 are used, is £7.95 per donor and, for the larger vehicles with a capacity of 90, is £5.09 per donor. The average mobile cost per donor is significantly higher than the fixed venue average. Furthermore, even though some of the small and large mobile venues have a similar cost to capacity ratio, there are sufficient fixed venues at an equal or cheaper cost to process all invited donors, explaining why no mobile venues are visited throughout the entirety of scenario 3's schedule.

Table 5.12 presents the top 20 venues in terms of cost to capacity ratio available in scenario 3, as well as the number of times they were visited, their capacity and finally the venue's cost to donor ratio.

Interestingly, five of the venues in Table 5.12 (venues 17, 34, 42, 52 and 62) are never visited during this schedule, despite other more expensive venues being visited numerous times. The reason behind this is straightforward, these venues belong to panels that operate multiple venues within them and each has a corresponding venue that is cheaper. As such, there is no need to visit these venues to collect blood from donors within their respective panels.

Venue ID	Times Visited	Capacity	Cost/Capacity Ratio
83	27	191	1.60
43	39	108	1.63
117	9	92	1.63
59	19	137	1.64
16	37	128	1.65
87	19	174	1.69
51	35	127	1.75
84	47	114	1.75
33	15	96	1.76
32	18	114	1.78
63	33	108	1.79
17	0	134	1.85
101	9	121	1.86
34	0	90	1.87
62	0	132	1.89
41	57	67	1.90
42	0	139	1.91
88	12	96	1.91
52	0	154	1.95
109	6	168	1.96

Table 5.12: Top 20 most frequently visited venues

The top 10 venues with the cheapest cost relative to capacity are visited on average 27 times, well above the average for all fixed venues (5 times) and the average number of times fixed venues are visited in the current blood service schedule (17 times). The most expensive venue visited costs approximately £2.75 per donor. This venue is visited only once despite being located in a very densely populated region. None of the other venues in this panel are visited as they all have higher per donor costs. Curiously, this panel has the most visited venue in the base schedule, indicating that the cost to capacity ratio plays an important role in determining which venues the model chooses to visit.

Panels

Figures 5.9 and 5.10 provide an insight into how panels are visited in this schedule, specifically Figure 5.9 shows the number of times a panel is visited, while Figure 5.10 shows the distribution of how frequently all panels are visited. Unsurprisingly, due to the relationship between venues and panels, popular panels, those visited more than 10 times, are visited more frequently in this schedule than in the blood service's schedule. This increase in visits takes the average number of visits to each panel from 5 times for the base schedule to over 17 times in the new schedule. Furthermore, the most popular panel in the blood service's schedule figures prominently in the new schedule with the number of visits more than doubling. However, the second most popular panel does not figure at all in the new schedule, owing to the high capacity to donor ratio of its venues as examined above.

More significant, however is the fact that the proposed schedule in scenario 3 only visits 25% of panels whereas the base schedule visited every panel at least once. As mentioned previously this may be a marketing or donor satisfaction strategy although it does indicate that there is a possible excess of panels being operated by the blood service.

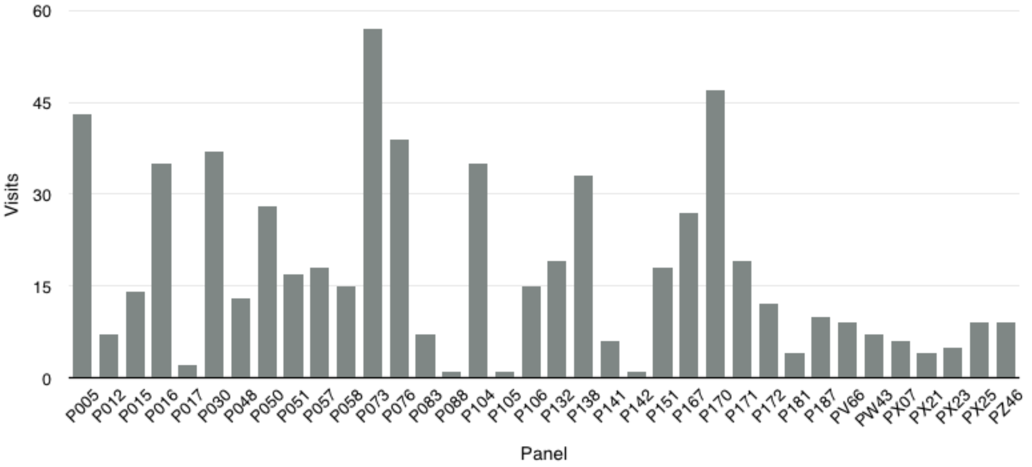

Figure 5.9: Popular panels

The vast majority of the panels not visited have a relatively small donor base with each panel having on average 145 registered donors (compared to 520 registered donors in the panels visited). These panels are constituted mostly of small business parks, shopping centres, and small villages and as such over 73% of these panels are served exclusively by mobile venues. Therefore, the reason a significant number of panels are not visited in this schedule is due to the fact that the venues available in these panels all have a much higher cost to capacity ratio than panels containing fixed venues.

The venue cost to capacity ratio also has an impact on how many venues are visited within multi-venue panels. While there are a total of 39 panels that have two or more venues, of which 23 are actually visited, none of these panels make use of the multiple venues available. This is in stark contrast to the blood service's schedule where 95% of panels that have multiple venues visit at least

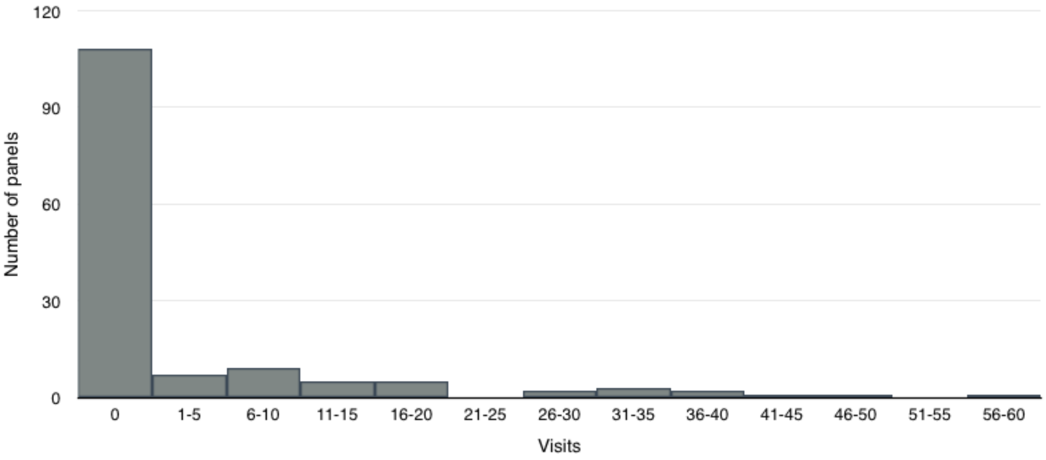

Figure 5.10: Frequency of panel visits

two of the venues.

Focusing on the number of units collected in each panel, Table 5.13 presents the highest collecting panels. The top 10 panels are responsible for just below 60% of all units collected considerably more than the top 10 in the base schedule - which contributed slightly less than 30%. Furthermore, the average contribution of each of these panels has increased from 2.95% to almost 6% with the highest contributing panel satisfying around 7.6% of the demand compared to only 3.4% in the blood service's schedule. This translates to each of the visited panels collecting on average 1,823 units of blood in this schedule.

Figure 5.11 presents as a percentage the cumulative contribution to the number of units collected for the blood service's current schedule and scenario 3's schedule. As each visited panel in scenario 3's schedule collects more blood, the base schedule's curve naturally lags behind the proposed schedule's, using more than double the panels to get the number of units collected, while the other schedule

Panel	Units Collected	Contribution To Demand (%)
45	5,010.95	7.63
44	4,587.78	6.99
9	4,424.01	6.74
27	4,148.07	6.32
21	3,904.59	5.95
1	3,867.31	5.89
20	3,646.86	5.55
34	3,535.03	5.38
12	2,967.36	4.52
Average:	3,930.98	5.98

Table 5.13: Panels where most blood is collected

required only 36 panels to reach 100% of units collected. It should be noted that as well as using more panels, the base schedule collects 2,884.27 units of blood less than scenario 3's schedule. This graph further demonstrates the surplus of panels available and that an efficient collection schedule could be operated with less than 40 panels in the entire region.

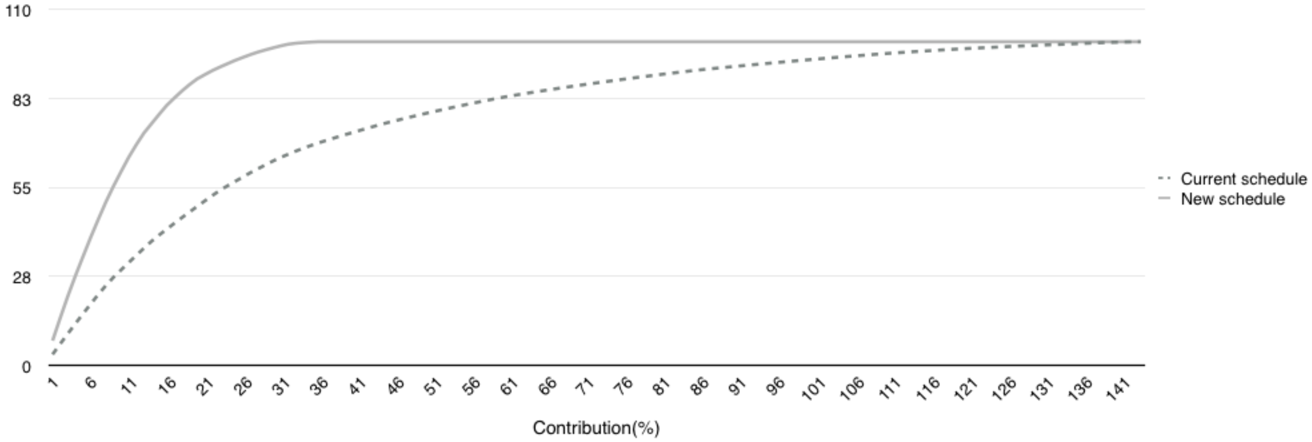

Figure 5.11: Cumulative contribution for new and current schedules

Donors

Having analysed the venue and panel selections of scenario 3's schedule, this final subsection analyses how efficiently donors are managed and explores the possible cause of the shortfalls. Scenario 3 will be compared against scenario 2's results. There are two main reasons for this. Firstly, it is difficult to make comparisons against scenario 1's donor results as a significant amount of over- and under-collecting happens. Secondly, scenario 2 does not suffer from the over-collecting issue and has the added benefit of having donors assigned by the blood allocation model, in the same way as scenario 3.

Throughout the schedule, 65,650.27 units were collected. However, this number is still short of the overall demand target of 65,677 units resulting in a minor shortfall of 26.78 units. This compares favourably to the 62,766 units collected in scenario 2's schedule which has over 5,000 units of shortfall. Table 5.14 provides a complete breakdown of the difference between the demand and the number of units collected, in scenario 3's schedule, by period and blood group.

Out of all the units collected in this schedule a total of 41,434.39 units were donated by registered donors, which were made up of 23,262 male and 18,172 female donors. The remaining units, making up approximately 37 % of units collected, were all donated by walk-in donors. More than 57,204 registered donors were invited to collect the units donated by them, this represents a turnout of 72.4% which is similar to the turnout in scenario 2. On average this schedule collects 1,641 units of blood per period slightly more than the average in scenario 2 (1,569 units per period). Interestingly, period 35 collects the most units (2,099) in this schedule and meets the period's collection target however, in scenario 2 this period is one of two that incurs a penalty - this period incurs a 52 unit shortfall over multiple blood groups. Demonstrating that not only is efficient donor assignment needed but also careful venue (and panel) selection to minimise shortfalls.

Overall this schedule performed well with six out of the eight blood groups meeting their respective demand targets. There appears to be some minor over-collecting happening, though as this totals only 0.03 of a unit it is most likely introduced by numerical imprecision and as such is not analysed further. The two blood groups that underperformed during this schedule were the AB rhesus positive and negative blood types which are amongst the least common, although for the B- blood group, which is also very uncommon, sufficient blood was collected to meet its demand target.

Table 5.14 shows that the shortfall for the AB+ blood group is only 0.68 units in total and that shortfalls are only incurred in three consecutive periods (periods 27-29). The total shortfall for the AB- blood group is much larger at 26.07 units of blood. This shortfall is also incurred over more periods, 13 in total, which can be split into two broad groups; penalties incurred in the first half of the schedule (which accounts for the shortfalls in periods 11 through 15), and those incurred in the second half of the schedule (periods 24, 25, 27, 28, 31, 32, 33 and 39).

To better comprehend the cause behind the shortfall for this schedule it is

Period	O+	O-	A+	A-	B+	B-	AB+	AB-	Shortfall
1	-	-	-	-	-	-	-	-	-
2	-0.01	-	-0.01	-	-	-	-	-	-0.02
3	-	-	-	-	-	-	-	-	-
4	-	-	-	-	-	-	-	-	-
5	-	-	-	-	-	-	-	-	-
6	-	-	-	-	-	-	-	-	-
7	-	-	-	-	-	-	-	-	-
8	-	-	-	-	-	-	-	-	-
9	-	-	-	-	-	-	-	-	-
10	-	-	-	-	-	-	-	-	-
11	-	-	-	-	-	-	-	2.31	2.31
12	-	-	-	-	-	-	-	2.10	2.10
13	-	-	-	-	-	-	-	2.91	2.91
14	-	-	-	-	-	-	-	2.57	2.57
15	-	-	-	-	-	-	-	1.51	1.51
16	-	-	-	-	-	-	-	-	-
17	-	-	-	-	-	-	-	-	-
18	-	-	-	-	-	-	-	-	-
19	-	-	-	-	-	-	-	-	-
20	-	-	-	-	-	-	-	-	-
21	-	-	-	-	-	-	-	-	-
22	-	-	-	-	-	-	-	-	-
23	-	-	-	-	-	-	-	-	-
24	-	-	-	-	-	-	-	0.33	0.33
25	-	-	-	-	-	-	-	2.62	2.62
26	-	-	-	-	-	-	-	-	-
27	-	-	-	-	-	-	0.30	1.44	1.74
28	-	-	-	-	-	-	0.34	2.26	2.60
29	-	-	-	-	-	-	0.05	-	0.05
30	-	-	-	-	-	-	-	-	-
31	-	-	-	-	-	-	-	0.92	0.92
32	-	-	-	-	-	-	-	0.93	0.93
33	-	-	-	-	-	-	-	3.22	3.22
34	-	-	-	-	-	-	-	-	-
35	-	-	-	-	-	-	-	-	-
36	-	-	-	-	-	-	-	-	-
37	-	-	-	-	-	-	-	-	-
38	-	-	-	-	-	-	-	-	-
39	-	-	-	-	-	-	-	2.95	2.94
40	-	-	-	-	-	-	-	-	-
Total	-0.01	-	-0.01	-	-	-	0.69	26.07	26.73

Table 5.14: Scenario 3's shortfall

necessary to present the number of donors available and the number of donors that donated for both the rhesus AB negative and positive blood types. In order to satisfy the demand targets for these two blood groups 1,970.31 units of AB positive and 656.77 units of AB negative blood need to be collected. The blood service had 930 registered donors of type AB+ and 275 registered donors of type AB- available throughout the schedule. These donors contributed 1,243.16 and 388.53 units to the AB- and AB+ collection targets respectively, with the remaining units coming from walk-in donors.

The AB+ blood group is considered, as this group has the least shortfall which simplifies the analysis. The first observation is that in each of the three periods that incur a shortfall there are still venue sessions available and seeing as not all venues have been utilised it would be possible to open additional sessions. This would indicate that the shortfall is not caused by a lack of venues or venue sessions in these periods. Moreover, when considering the total number of donors that all of these open venues can process there is still sufficient spare capacity to invite additional donors. This suggests that there are in fact insufficient registered donors with the AB+/- blood types to satisfy the collection targets.

By examining the number of registered donors that could be invited to attend any open session in these periods, it becomes evident that there is a lack of donors with the AB+/- blood types, with the sessions either having no donors available or having invited the donors already. However, there are still donors of the correct blood type available in these periods, although they are located in different panels. The reason these donors are not invited falls into two main categories. Firstly, the additional cost of opening a venue and inviting the donors outweighs the actual penalty cost and secondly due to the number of periods a donor needs to rest; inviting donors to these periods would cause a larger penalty to be incurred later on in the schedule.

There is a very similar story behind the shortfall for the AB- blood group. Despite there being sufficient spare capacity at some of the open venues there are insufficient eligible registered donors of the correct blood type that can be invited to attend these sessions.

While it would be possible for the model to open additional sessions and allow walk-in donors to make up the shortfall this approach has three important drawbacks. In the first instance this approach would lead to over collecting blood, secondly the cost of opening the additional venues would be uneconomical. Finally and most importantly, both of these blood types are very rare, requiring numerous walk-in donors to be accepted before meeting the targets, an additional 2,600 donors need to be invited to satisfy the 26 units of shortfall. As such without recruiting more donors of the AB blood types or reassigning the donors to different panels, the blood service will be unable to match their collection targets for these groups. Consequently, any schedule construed with these parameters, no matter how efficient, would still incur a shortfall charge.

Summary

This chapter started off by presenting how CPLEX, the simulated annealing and the variable bounding heuristics performed when solving small data sets. Based on the variable bounding heuristic's performance, it was selected as the heuristic to be used when solving the real data set. Using this heuristic, a number of different scenarios were run. The current schedule operated by the blood service, as well as a new schedule generated by the blood collection model were considered.

The current schedule was split into two further scenarios, the first considered the blood service's current schedule and the units that were collected. The second used the blood service's schedule, however donor assignment to the open venues were left up to the scheduling model. A cost was assigned to both schedules and their usage of venues and panels was explored, as well as how well they performed in meeting the collection targets. It was found that in their current schedule the blood service visit every panel at least once and that no venue was visited more than 15 times. Furthermore, both fixed and mobile venues were visited. In the first scenario, there was a significant penalty cost due to missing collection targets, though this was significantly reduced in the second scenario it remained relatively high.

The new schedule was created by allowing the model to determine in which periods and how often each of the venues were open and how donors should be assigned to these venues. This schedule did not visit all of the panels. In fact, only a relatively small number were visited. Similarly, only a few fixed venues were visited, however these venues were visited much more frequently, with one being visited almost 60 times throughout the planning horizon. Moreover, no mobile venues were visited. In terms of cost, this schedule performed better than the blood service's current schedule.

The next chapter considers some additional scenarios. In particular, it explores the potential impact of compensating donors for their donations and analyses how changes to the inter-donation times could aide in collecting more blood.

Chapter 6

Scenario Analysis

This chapter uses the blood collection model to explore several what-if scenarios. These scenarios can be split into two main categories. The first set of scenarios consider how a blood service could improve donor attendance and reliability. The second set of scenarios consider how changing the inter donation interval could impact the number of units collected and overall donor availability. As in the previous chapter, the venues and panels visited in each schedule are analysed as are the number of donors invited and the number of units collected. Similarly, the primary means to measure how efficient a schedule is, will be based on the overall cost of operating that schedule.

6.1 Improving Donor Attendance

Blood services are always seeking ways to improve donor reliability. One such way to encourage donations is by paying donors every time they donate, though this is currently not done in the UK. This section considers the impact of paying donors on the overall cost of the collection schedule and the units collected.

Three scenarios are considered in this section, the first two involve paying donors for their donations, while the third considers the impact of increased donor reliability. The first two scenarios pay registered donors £20 per unit of blood donated, as donors are compensated it is assumed that they will be more reliable, so the two scenarios consider different increases in donor reliability. The first scenario will pay donors £20 per unit and assume that donor reliability is increased to at least 70% and the second scenario will pay donors the same amount per unit but assume that reliability is increased to at least 90%. Importantly only registered donors are paid in these scenarios, first time donors and walk-in donors do not receive any compensation. Finally, the third scenario will increase each panel's reliability by 20% without compensating donors for their donations.

The total amount paid to donors is a variable cost - dependant on the number of donors invited, furthermore from the blood service's perspective paying donors is akin to increasing the cost of processing each unit of blood. In light of this the cost of processing a unit of blood is simply increased by £20 in the scenarios where donors are compensated for their donations. As such, it is necessary to make a minor adjustment to the blood scheduling model objective function, to allow paid registered donors and non-paid walk-in donors to be distinguished. To achieve this, a new cost parameter (c') is introduced, this parameter will be increased to include the cost of collecting the blood and the amount paid to donors for their donations. For first time and walk-in donors the cost parameter (c) is unchanged. The new objective function (6.1) presented below replaces the old one with all other constraints in the model remaining the same.

$$\min \underbrace{\sum_t \sum_j \sum_i f_{ij}^{tv} x_{ijt}}_{(i)} + \underbrace{\sum_t \sum_j \sum_k \sum_i f_{ik}^{mv} y_{ijkt}}_{(ii)} + \underbrace{\sum_t \sum_g \sum_i c'(\alpha_i^m z_{igt}^m + \alpha_i^f z_{igt}^f)}_{(iii)}$$
$$+ \underbrace{\sum_t \sum_i c z_{it}^w}_{(iv)} + \underbrace{\sum_t \sum_g w_g s_{gt}}_{(v)} \tag{6.1}$$

The data for the α_i^m and α_i^f parameters (which control how reliable donors are) were updated as follows. For the scenarios where donors are paid the new parameter values for each period in the schedule is either that period's current value or the new reliability score, whichever is higher. This takes the average reliability for the first scenario (at least 70% reliability) when donors are paid to 73% for male and female donors, which is an increase of 2% compared to current levels. The increase for the second scenario (at least 90% reliability) is more pronounced, increasing the average α_i^m and α_i^f values by 19%. In the third scenario described above each period's α_i^m and α_i^f values are increased by

Cost	Paying donors £20 per donation, increasing donor reliability to 70% (£)	Paying donors £20 per donation, increasing donor reliability to 90% (£)	Increasing donor reliability by 20% (£)
Objective	9,602,259.44	9,193,600.61	8,318,740.70
Fixed cost	184,905.22	135,514.00	130,961.44
Variable cost	8,973,848.82	9,051,534.39	8,163,227.26
Penalty cost	443,505.40	10,552.52	24,552.00

Table 6.1: Scenario costs - improving reliability

20%, increasing the average values to 77% and 76% respectively.

As in the schedules examined above the objective function for each of the scenarios are considered first, before moving on to explore the venues visited and finally examining how well the collection targets were met.

Table 6.1 presents the cost for each of the scenarios described above as well as their component costs. The third scenario in this section, where donors are not compensated, is by far the cheapest scenario. Furthermore, this scenario is cheaper than the schedule presented above in which the blood scheduling model was able to select and allocate venues and donors. However, despite the 20% increase in donor reliability, this schedule only manages to achieve a slight reduction in shortfalls compared to the same schedule. The main cost difference between these two schedules can be explained by the venues visited, with the new schedule reducing overall fixed costs by just under £5,000.

As for the scenarios in which donors are compensated, the scenario with a 70% reliability rate produces the most expensive schedule. The primary cause for this is due to the extremely large shortfall incurred - increasing the schedule's costs by £444,505.40. Moreover, it is worth noting that this schedule proved difficult for the heuristic to solve, creating a solution with a gap of over 5% to the lower bound. Owing to the poor performance of this schedule it shall not be considered further in this analysis.

The schedule with the much higher reliability percentage (90%) performed significantly better than the above schedule, with the overall cost of this schedule standing at £9,193,600.55. Furthermore, this schedule incurred a penalty cost of only £10,552.52 signifying that the collection target was missed by fractionally more than 10 units. In total this schedule collected 44,328.4 units of blood from registered donors, as such the total compensation paid to donors amounted to £886,568, or more than 9.5% of the overall cost.

For at least one of these scenarios compensating donors significantly reduces the number of units under collected. However, this requires a high reliability rate and more importantly increase the schedule's overall cost. In the current economic climate it would be very difficult to justify such a large increase in collection costs. This is especially true when considering the fact that when compared to scenario 3's shortfall of 26 units the cost for every unit reduced is more than £540,000 (total cost of paying donors over the additional units collected). Therefore, given the current cost structure it would be unwise for the blood service to consider paid donations.

Venues and Panels

In the schedule where donors are not paid, a total of 610 venues are opened; as in other scenarios created by the blood scheduling model, all of the venues open are fixed. The number of venues opened in this schedule is less than the number opened in scenario 3 - there is a reduction of 19 venues in total. Owing to this reduction and the increase in units collected in this schedule, the average number of donors processed at every venue has increased from around 104 donors to 107 donors. Only period 28 has utilised the maximum number of venue sessions available, though this period does not produce a shortfall. This in turn points to the fact that the shortfall produced by this schedule is likely to be caused by a shortage of donors with the correct blood type.

The venues opened in this schedule are very similar to the venues opened in scenario 3, with only two new venues being visited, these new venues (venues 90 and 106) are each visited once. The other changes are: how frequently the venues are visited, with venue 1 seeing the largest increase in visits, raising from 43 to 47 visits. While, venues 83 and 101 see the largest decreases with both venues having three less visits. For this schedule a total of 38 panels were visited, which is less than a third of all available panels. On average each panel collects 1,726 units of blood with the highest contributing panel collecting more than 5,000 units and the lowest contributing panel only collecting 25 units over the entire planning horizon.

The venue allocation in the scenario where donors are paid and the reliability rate is increased to 90% is considered next. A total of 609 venues are opened in this scenario - only one less than the non-paid donor schedule, however, this schedule allocates the maximum number of venues in more periods. A total of three periods have fully utilised their available venue sessions, although as in the other schedule, none of these periods incurs a shortfall. Furthermore, both scenarios process, on average, the same number of donors per venue (107 donors).

Although the venues that each of these scenarios visits are remarkably similar there are some significant differences in how frequently they are visited, explaining the differences in venue's costs discussed previously. The largest difference between these schedules is how frequently venue 41 is visited, in the current schedule it is visited 48 times, which is nine times less often than in the other scenario. Furthermore, venue 1 also sees a significant decrease in visits as does venue 58. On the other hand, venues 11, 53 and 63 all see large increases in the number of times they are each visited, with venue 11 increasing from 34 visits up to 41 visits. One additional panel is visited in this scenario compared with the previous one and despite collecting more units the average contribution per visited panel is slightly less standing at 2.56% compared to 2.62%, which equates to a reduction of approximately 40 units per panels.

Donors

Considering the non-paid donor scenario first, a total of 65,652.46 units were collected, which is only an additional two units compared to scenario 3. However, only 54,591 registered donors needed to be invited to collect 42,589 units; which represents an average turnout rate of 78% aided by the increase in donor reliability. This compares favourably to scenario 3's turnout rate of 72% requiring 57,205 donors to be invited to collect 41,434 units of blood. In other words, without the increase in donor reliability to collect ten units of blood it is necessary to invite on average 14 registered donors, whereas with the increased reliability only 12 registered donors need to be invited.

In the paid donor scenario the turnout rate improves further due to the significant increase in donor reliability. In this scenario a total of 49,254 registered donors are invited of which 44,329 attend a session and donate, this equates to needing to invite on average 11 registered donors to collect ten units of blood.

However, despite the increase in the number of units collected neither of these scenarios were able to collect sufficient blood to satisfy the overall collection target. As mentioned above the non-paid scenario missed the collection target by just over 24 units, while the paid scenario under-collected slightly over 10 units of blood. Interestingly, both scenarios are able to satisfy the demand targets for seven out of the eight blood groups, with only the AB- type falling short. Table 6.2, 6.3 and 6.4 displays the shortfall incurred for each scenario and is broken down into blood groups and periods.

In the schedule where donors are not paid a total of 10 periods incur shortfalls, while for paid donors the number reduces to only 6 periods. For both of these scenarios period 15 incurs the highest shortfall. In the first scenario, the collection target is missed by 5.43 units, while in the second scenario it is missed by 3.7 units; moreover both scenarios have venue sessions still available in this period. Upon analysing the number of donors available in this period it becomes evident that the cause of the shortfall is due to the schedule needing to rest donors. While, reducing the shortfall in this period is possible it would inevitably cause new shortfalls to appear in other periods.

Comparing the shortfalls for the non-paid schedule and the schedule produced in scenario 3 the increase in donor reliability has led to some improvements being made. For starters the small shortfall that scenario 3 incurred for the AB+ blood group has been eliminated and secondly the shortfall for the AB- blood group has been reduced by just over one and a half units of blood.

Period	O-	A+	AB+	AB-	Shortfall
7	-3.85	-	-	-	-3.85
9	-	-	-	0.95	0.95
11	-	46.42	15.40	5.97	67.78
12	-36.19	4.08	-	6.20	-25.91
14	-	61.06	4.59	0.50	66.15
15	-20.77	-	-	-	-20.77
16	-108.84	-	-	-	-108.84
17	-88.42	-	-	-	-88.42
20	-	52.98	9.74	0.96	63.68
21	-6.78	-	-	-	-6.78
23	-17.07	3.35	-	2.13	-11.59
25	-0.55	-	-	-	-0.55
26	-	18.77	-	2.72	21.49
27	-	3.05	6.09	-	9.14
28	-	78.82	14.00	4.55	97.36
29	-41.27	-	-	-	-41.27
30	-33.66	-	-	-	-33.66
31	-21.85	-	-	-	-21.85
32	-9.03	64.25	-	3.61	58.84
33	-110.32	-	-	-	-110.32
34	-10.41	-	-	-	-10.41
35	-2.47	10.56	2.12	1.70	11.90
36	-	14.69	4.26	-	18.95
37	-30.77	-	-	-	-
40	-29.33	-	-	-	-29.33

Table 6.2: Shortfall by period and blood group when donors are paid £20 per donation and reliability is increased to 70%

Blood Group	Period	Shortfall (units)
AB-	15	3.71
AB-	32	2.68
AB-	26	1.93
AB-	28	0.91
AB-	24	0.75
AB-	11	0.57

Table 6.3: Shortfall by period and blood group when donors are paid £20 per donation and reliability is increased to 90%

Period	AB-
10	2.69
13	0.20
14	1.81
15	5.43
25	3.23
26	4.37
27	4.53
30	1.37
31	0.63
33	0.30

Table 6.4: Shortfall by period and blood group when reliability is increased by 20%

However, these increases are only small with both scenarios still having shortfalls in eight common periods. Furthermore, even when the reliability of donors is increased to 90% shortfalls still occur, further highlighting the fact that there are insufficient registered donors with the blood group AB-.

6.2 Changing Intervals

The operational impact of changing the amount of times a donor needs to rest between donations is considered in this section. Due to the fact that the human body needs time to restore any lost blood, donors are required to rest between donations; this is often referred to as the donor interval period. While the donor is resting they are unable to attend any donation session. A resting donor differs slightly to a donor who has been deferred, though neither are eligible to donate. A donor can be deferred for having been abroad, getting a tattoo or falling pregnant, among other reasons. Furthermore, any active donor whose blood levels have not recovered in the interval period will be deferred. The length of time a donor is deferred for varies considerably, ranging from a few weeks up to a whole year. As discussed previously, male and female interval times differ in England, with female donors being given longer to recuperate.

Although every blood service recognises the importance of resting donors, there is considerable disagreement on how long they should be rested for. Table 6.5 presents the number of weeks that male and female donors need to wait between donations for various countries. As can be seen in the table most countries differentiate between male and female donors, however four of these countries do not, with Ireland resting donors for a relatively short time (10 weeks) and with Scotland taking a much more conservative interval length (16 weeks). On average male donors need to wait 12.6 weeks between donations and female donors wait an additional week, the shortest number of weeks a male and female donor need to wait are 8 and 10 weeks respectively. In England and Wales, male and female donors are required to wait for 12 and 16 weeks respectively, which for both cases is more than the average resting time compared to other countries.

The reason why there is no commonly accepted inter-donation time is that there has been no research into the medical and operational effects this can have on

Country	Male	Female
England	12	16
Austria	8	10
Finland	8	12
France	8	12
Germany	8	12
Ireland	10	10
Estonia	10	12
Netherlands	10	16
Denmark	12	12
Wales	12	16
Flanders	12	12
Slovenia	12	16
Spain	12	16
Scotland	16	16
Average	12.6	13.4

Table 6.5: Male and female donor intervals in various countries

a donor's health and the blood service's efficiency. As such, each blood service has developed their own rule of thumb. However, with the ageing population requiring more blood, an ever shrinking donor pool and the need to ensure the safety of their donors, the NBS alongside Cambridge and Oxford University are carrying out research into this area.

As the blood scheduling model has been parameterised to allow for different inter-donation times, the model is capable of exploring the impact on the schedule of any proposed changes. However, as the INTERVAL(Moore et al., 2014) study is still on-going, the optimum resting period for donors is still unknown, as such two new resting periods are considered. The first new interval has men waiting 10 weeks and women waiting 14 weeks between donations, reducing the resting time for both genders by two weeks. While, the second new resting time is one of the shorter intervals currently used in Europe, with men waiting only

8 weeks and females 12 weeks.

Although there is a possibility that the INTERVAL study recommends that the current inter-donation time be lengthened, it is considered unlikely. There are two main reasons for this, firstly the current resting times used in England and Wales are relatively conservative and secondly, if the inter-donation times were too short then the blood service would have observed a significant number of deferrals caused by low blood levels. Therefore, no longer intervals will be considered in this section. There is also the possibility that the study recommends different resting times for different types of donors, differentiating by blood group, activity level or age. However, as there is no data available on exactly how donors might be differentiated the changes in inter-donation times considered below are differentiated only by gender.

6.2.1 Current Demand Levels

The results of shortening the interval times to either 8 or 10 weeks for males and 12 or 14 weeks for female donors are presented next, as well as how the new number of registered donors available were determined.

Updating the donor scheduling model to handle the different resting periods is relatively straightforward, the cycle parameters for male and female donors need to be updated, as do the a^m_{igt} and a^f_{igt} parameters for the number of available registered donors. In turn, these updated parameters will affect the donor availability constraints in the scheduling model. To determine the new values for a^m_{igt} and a^f_{igt}, all donors have their inter-donation times updated accordingly

Cost	Changing donation interval to 8 weeks(male) and 12 weeks(female) (in £)	%	Changing donation interval to 10 weeks(male) and 14 weeks(female) (in £)	%
Objective	8,572,701.11	100	£8,298,556.77	100
Fixed cost	125,194.83	1.46	126,791.68	1.53
Variable cost	8,138,803.28	94.94	8,165,497.31	98.40
Penalty cost	308,703.00	3.60	6,297.78	0.08

Table 6.6: Scenario costs for different donor intervals with current demand levels

and by starting from the oldest donation date available in the dataset the week in which they are available to donate for the first time in the new schedule was determined.

The objectives, along with their component costs, for the schedules produced with the new resting periods are presented in Table 6.6. At first glance the overall schedule costs appear counter intuitive with the shortest inter-donation time (8 weeks and 12 weeks) schedule costing more than both the current interval and the shorter - 10 weeks and 14 weeks schedules. However, as discussed above and as evident by looking at Table 6.6, the shortest inter-donation time schedule was particularly difficult to solve, with a reported optimality gap of 3.39% after six hours. This makes direct comparisons of the schedules' objectives unfair. Furthermore, when comparing the schedule's lower bounds, the 8 week and 12 week resting periods is the lowest. This indicates that although the schedule produced by the variable bounding heuristic is worse than the schedules with longer intervals, a better schedule might be found if the heuristics were run for longer. This analysis will include the shorter schedule for completeness and also owing to the fact that the heuristic produces a much better schedule when the overall collection targets are increased slightly.

To allow for meaningful and fair comparisons to be made both of the schedules produced using the shorter interval periods shall be compared to scenario 3. The schedule that requires male and female donors to wait 10 weeks and 14 weeks respectively has a total cost of £8,298,556.77, which is cheaper than scenario 3's objective - resulting in a net saving of £27,102.26. Interestingly, both of the new schedules still incur penalty costs signifying that despite reducing the number of weeks a registered donor needs to rest there are still insufficient registered donors for the current demand targets. One of the main reasons why the schedule produced, when resting donors 8 weeks and 12 weeks, is more expensive is due to the large penalty cost compared to the base schedule's. However, the penalty incurred by resting donors 10 weeks and 14 weeks respectively is more than £20,000 cheaper than scenario 3's penalty cost. This represents a reduction of more than 20 units resulting in the new schedule under-collecting by slightly more than 6 units. A reduction in the fixed costs incurred is also evident, with the shorter interval schedule having fixed costs that are approximately 6%, or just over £9,000, cheaper than scenario 3's. As has been the case in all the other scenarios the variable cost makes up more than 90% of the schedules overall cost, in this case the variable cost has increased owing to the slight increase in the overall number of blood units collected.

Venues and Panels

As in scenario 3, the blood allocation model was allowed to assign donors and decide which venues and panels were opened and visited. In total, the schedule for donors waiting 8 and 12 weeks opened/visited a total of 628 venues and panels, while the schedule with donors waiting 10 and 14 weeks opened/visited 615 venues and panels. The first schedule opens a similar number of venues as in scenario 3 however; the second scenario manages to open 13 less and at the same time to reduce its penalty cost.

Table 6.7 shows the number of venues opened in each of the periods for both the shortest interval (8 and 12 week waits) and the shorter interval (10 and 14 week waits) schedules. In the shortest interval schedule there are a total of 131 unassigned venue sessions resulting in an average utilisation rate of 82.89%. In total there are three periods (weeks 5, 9 and 35) that have allocated the maximum number of venues permitted, two of these - periods 5 and 9, are also fully utilised in scenario 3. None of the fully booked periods incur shortfall costs, though period 9 does over collect several units. As such, the periods within which shortfalls occur all have venue sessions available, indicating that either venues could be opened in these periods or that donors could be assigned to improve this schedule.

For the schedule with waiting times of 10 and 14 weeks there is only one period that is completely full, resulting in 144 unused sessions, and an average utilisation rate of 81.27%. Period 14, which is the full period, does incur a slight shortfall penalty, as does the same period in scenario 3, though this is less than 1 whole unit. The largest shortfalls occur in periods that still have venue sessions available, suggesting that as in scenario 3 there is a registered donor shortage.

Figure 6.1 displays the number of times each of the available fixed venues were

	Number of venues open with intervals of		
Period	8 weeks (male) and 12 weeks (female)	10 weeks (male) and 14 weeks (female)	Maximum number of venues
1	14	16	17
2	17	18	20
3	15	16	22
4	14	15	18
5	18	16	18
6	17	15	19
7	19	19	22
8	15	15	18
9	19	18	19
10	16	16	18
11	14	14	17
12	17	14	18
13	12	12	19
14	16	17	17
15	18	16	20
16	18	17	20
17	18	15	19
18	13	14	20
19	15	16	18
20	15	15	20
21	11	11	16
22	17	16	19
23	19	17	20
24	15	13	22
25	18	16	20
26	15	17	18
27	12	13	17
28	15	17	18
29	15	15	21
30	17	17	20
31	17	15	19
32	15	13	16
33	17	16	20
34	14	12	19
35	20	19	20
36	18	17	20
37	13	15	19
38	15	16	18
39	15	15	18
40	10	11	20
Total:	628	615	759

Table 6.7: Number of venues open in each period

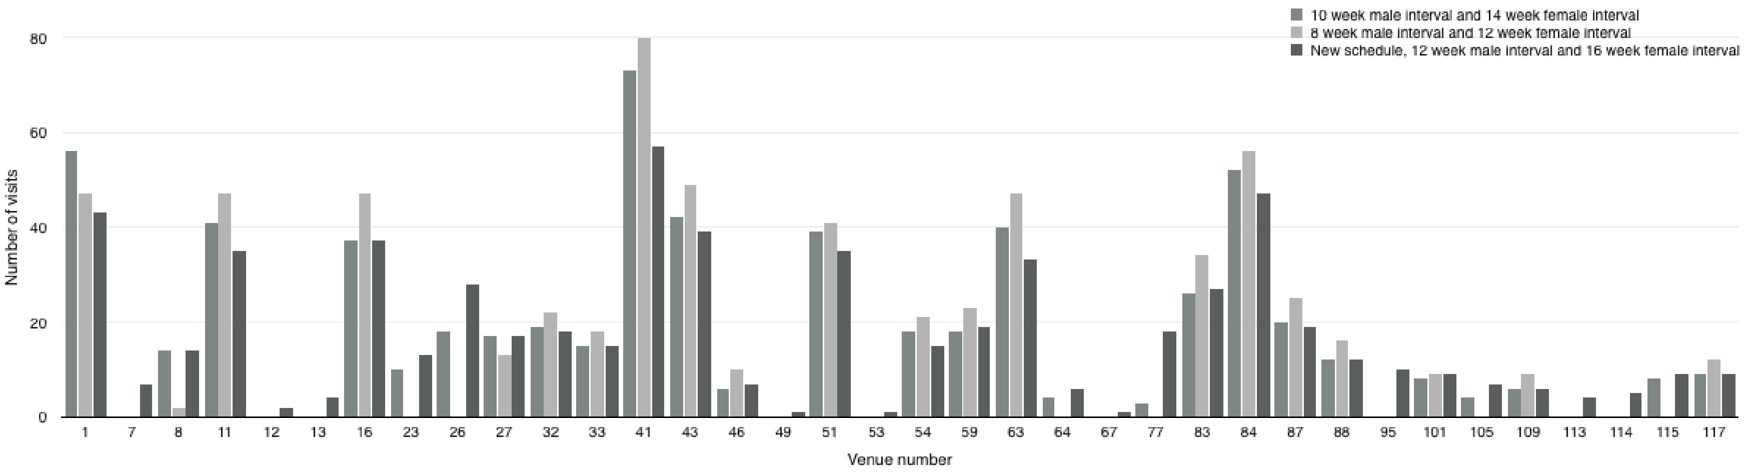

Figure 6.1: Number of venue visits for different interval schedules (Only venues visited at least once are shown)

visited for both interval lengths, as well as for scenario 3. As in scenario 3, neither of the shorter interval schedules visited any mobile venues - owing to their larger cost to capacity ratios. As such, the discussion below only focuses on fixed venues.

Comparing the venues visited, the vast majority are similar to those venues visited in scenario 3, with only nine out of 36 venues visited in scenario 3 not figuring in either of the schedules with shorter intervals. All of the venues not visited in either shorter interval schedules are frequented no more than 10 times, with most only being visited once or twice. The largest reductions in venues visited in scenario 3's schedule and the schedule produced using reduced inter-donation periods are observed in venues 77 and 26, which go from having 18 and 28 visits to 3 and 18 visits respectively. Both of these venues serve large populations and can process over 100 donors in a session. Additionally, neither of these venues are visited in the schedule produced using the shortest donor intervals.

For venues visited more than 30 times throughout the schedule, there is a noticeable increase in visits in at least one of the shorter interval schedules compared to scenario 3's. In particular, the most popular venue in scenario 3, venue 41, shows significantly more visits, with the shortest interval visiting the venue the maximum number of times in its schedule. For the schedule with donors waiting 10 and 14 weeks respectively there is a large increase in the number of visits made to the first venues going from 43 in scenario 3's schedule up to 56 visits.

One possible explanation for why venues present in scenario 3 are no longer present in the current schedules is that each of these venues have relatively high venue cost to capacity ratios. Furthermore, all venues that have significant reduction in how frequently they are visited also have above average cost to capacity ratios. As such, the average cost to capacity ratio for the shorter scenarios go from £2.03 per donor in scenario 3 down to £1.83 and £1.90 for donors waiting 8 weeks and 12 weeks and 10 weeks and 14 weeks, respectively.

For the schedule produced with the shortest inter-donation wait, all of the venues visited with the cheapest venue cost to capacity ratio have shown a significant increase in how frequently these venues are visited. Though it is worth bearing in mind that this schedule's objective was higher than scenario 3's. With donors waiting 10 weeks and 14 weeks the number of visits to cheaper cost to capacity ratio venues has generally increased or remained the same, curiously though the cheapest venue - venue 83 saw one less visit, as did the fourth cheapest venue. It is not immediately clear why this reduction occurred, although it is possibly due to a lack of registered donors.

Naturally, as the number of venues visited has reduced in both of the new schedules so has the number of panels visited. The schedule created using the shortest inter-donation times visited a total of 21 panels, while the schedule produced with donors waiting 10 and 14 weeks visited a total of 27 panels. As expected, there is a decrease in the number of panels visited compared to scenario 3's schedule, where 36 panels were visited. These reductions mean that only 15% for the first schedule and 19% for the second schedule of the total panels available are used, indicating that if the inter-donation time is shortened the blood service could further reduce the number of panels maintained.

As the number of panels visited has been reduced further it is expected that the number of units each panel contributes must have increased. Table 6.8 shows the contribution per panel to the number of units collected as a percentage,

For interval of 8 weeks and 12 weeks	
Panel	Contribution to units collected (% of total)
45	9.64
44	9.29
9	8.72
20	8.19
27	7.82
21	7.61
34	7.34
6	6.63
1	6.46
46	5.71
32	4.44
15	3.70
29	2.96
16	2.20
47	2.02
13	2.01
132	1.46
102	1.28
143	1.09
22	0.90
5	0.53

For interval of 10 weeks and 14 weeks	
Panel	Contribution to units collected (% of total)
45	8.55
1	7.67
20	7.45
44	7.35
9	6.98
27	6.76
34	6.38
21	6.28
6	5.99
46	4.83
5	3.68
32	3.52
15	3.00
12	2.91
29	2.69
13	2.51
11	2.49
16	1.91
47	1.63
139	1.32
132	1.26
102	1.12
143	0.95
22	0.78
4	0.65
35	0.64

Table 6.8: Contribution per panel for intervals of 8 weeks and 12 weeks and 10 weeks and 14 weeks

for both of the schedules with shorter resting periods. On average each of the panels collect 3,116.95 and 2,432.25 units of blood for donors waiting 8 and 12 weeks and for donors waiting 10 and 14 weeks respectively, this represents an increase of approximately 71% and 33% compared to scenario 3. For both of the schedules as well as in scenario 3 the panel that contributed the most is panel 45, contributing 6,308, 5,617 and 5,010 units for each of the schedules in increasing interval length. There is a significant increase in the contribution to the number of units collected by the top 10 panels, just under 60% for scenario 3 but rising to just over 68% for donors that wait 10 weeks and 14 weeks and raising further still to 77% for donors that wait 8 and 12 weeks. This demonstrates that the shorter the interval the more concentrated the collection efforts are in certain panels. Furthermore, as the interval length is shortened fewer panels contribute under 1% - 11 panels in scenario 3, 5 with donors resting for 10 and 14 weeks and finally only 2 panels for donors resting the least.

Donors

The primary motivation behind reducing the resting periods of donors is to try to increase the number of units collected without increasing the overall registered donor pool and at the same time trying to eliminate any shortfall. This subsection analyses donor allocations and how well both of the reduced interval periods did at meeting collection targets compared to scenario 3.

The shortest interval schedule collected a total of 65,456.03 units of blood over the 40 weeks, while the schedule that rested donors 10 and 14 weeks managed to collect a total of 65,670.27 units. However, as mentioned previously both of these schedules fall short of the overall demand target, though the schedule with donors waiting 10 and 14 weeks came relatively close with a shortfall of only 6.28 units.

The average number of units collected in each period when using the shortest interval was 1,636.40 units, while resting donors slightly longer the average increased to 1,641.77 units a period. This last average is very similar to the average achieved in scenario 3 of 1,641.26 units a period. Incidentally, all three of these schedules are able to collect sufficient units to satisfy the period with the highest demand - period 35 which requires a total of 2,099 units. Similarly, all schedules satisfy the demand in the period with the lowest collection target - period 21 which needs 1,078 units.

Comparing the total number of registered donors invited in each of these schedules reveals that the shortest resting period only invited 58,175 registered donors while the schedule with the slightly longer resting time managed to invite 57,341 registered donors, which is an average increase of only 553 donors when compared to scenario 3, which is about an additional 14 donors per week. Out of the donors invited to attend; 42,005 donors attended a session in the schedule with the shortest interval, while 41,432 donors donated in the slightly longer interval schedule. This results in an average turn out rate of 72.2%, which, as expected is very similar to scenario 3's average turn out rate.

Considering the shortest interval first and as with previous scenarios, Table 6.9 presents the number of units under- and over-collected by blood group for each of the periods. As demonstrated earlier when discussing the schedule's objective, this schedule falls short of the collection target by approximately 300 units. In total five out of the eight blood groups meet their demand targets however; the O- blood group does over-collect approximately 88 units throughout the schedule, which would result in the objective incurring an unnecessary £10,000 variable cost. Interestingly, in the periods where O- is over-collected there is no under-collecting suggesting that a more efficient allocation of the donors is possible.

The blood groups that are under collected in this schedule are A-, AB+ and AB- and each have respective shortfalls of 214, 84 and 9 units. Clearly, the majority of the shortfall is caused by under collecting A- blood, though interestingly this blood group was not a problem in scenario 3, whereas both AB blood types did cause penalty costs to be incurred. On the other hand, the shorter interval schedule only incurs a shortfall for 9.83 units of AB- blood while scenario 3 under-collects by approximately 26 units. In periods where shortfalls do occur in this schedule they are most often across all three of these blood groups, with the only exceptions being periods 28 and 29.

In each of the venues which incur a shortfall there are still venue sessions available, meaning that the model could open additional venues. There are several different reasons why the model might choose not to, firstly because there might be no registered donors of the correct blood type available and secondly the cost of opening the venue outweighs the penalty cost - though given how large the shortfalls are in this schedule it is unlikely. It would appear that the causes

Period	O-	A+	AB+	AB-	Shortfall
1	-	6.77	5.29	2.94	15.00
2	-	-	-	-	-
3	-	-	-	-	-
4	-	44.36	14.30	0.34	59.00
5	-	-	-	-	-
6	-	-	-	-	-
7	-	-	-	-	-
8	-	-	-	-	-
9	-2.63	-	-	-	-2.62
10	-	-	-	-	-
11	-0.00	57.86	18.89	3.25	80.00
12	-16.16	-	-	-	-16.16
13	-	-	-	-	-
14	-	-	-	-	-
15	-	-	-	-	-
16	-26.72	-	-	-	-26.72
17	-	-	-	-	-
18	-	-	-	-	-
19	-	69.18	21.38	1.45	92.00
20	-	-	-	-	-
21	-	-	-	-	-
22	-	-	-	-	-
23	-	-	-	-	-
24	-14.32	-	-	-	-14.32
25	-	-	-	-	-
26	-	-	-	-	-
27	-	-	-	-	-
28	-	2.99	1.07	-	4.06
29	-	-	1.65	-	1.65
30	-	33.08	22.06	1.85	57.00
31	-8.09	-	-	-	-8.09
32	-	-	-	-	-
33	-19.78	-	-	-	-19.78
34	-	-	-	-	-
35	-	-	-	-	-
36	-	-	-	-	-
37	-	-	-	-	-
38	-	-	-	-	-
39	-	-	-	-	-
40	-	-	-	-	-

Table 6.9: Shortfall for scenario where male donors donate every 8 weeks and female donors every 12 weeks

behind the shortfalls in this schedule are twofold. In the first instance there is a lack of registered donors especially towards the latter periods and secondly the venue selection and donor allocation are worse than in scenario 3, resulting in the over-collecting of O- and the large amount of under-collecting of A- type blood.

Table 6.10 presents the difference between units collected and the collection targets for the AB- blood group by period for the schedule in which donors are required to rest 10 and 14 weeks. This schedule performs much better than the previous schedule and even the schedule produced by scenario 3, accumulating only 6.68 units of shortfall. All but one of the blood groups (AB-) have collected sufficient blood to satisfy the demand, compared to two blood groups in scenario 3, this is despite opening 14 less venues than scenario 3. The shortfall was accumulated over four periods with two being in the first half of the schedule (periods 12 and 14) and two being in the second half (periods 28 and 39). The period that incurred the highest shortfall was period 39 accounting for just under 50% of units under collected.

Period	AB-
12	1.92
14	0.65
28	0.73
39	3.00

Table 6.10: Shortfall for scenario where male donors donate every 10 weeks and female donors ever 14 weeks

In all but one of the periods that incur a shortfall there is the possibility of opening additional venues. Period 14, which incurs a small shortfall of 0.65

units, is not able to open any additional venues having utilised all 17 of the available venue sessions. However, in spite of the shortfall being small if an additional venue slot was made available in this period the shortfall could be eliminated and the impact on the objective function would be minimal. The reasons why venues were not opened in other periods are very similar to those presented when analysing scenario 3. The first reason is that the venues that are open are either at capacity or there are no more donors with the blood group AB- available. The second reason is that to reduce the shortfall multiple venues would need to be opened in these periods and while this is feasible the cost of opening several venues could negatively impact the schedules overall cost. Finally, due to how intricately all the periods are inter-connected a reduction in one period could cause a larger shortfall to emerge later on in the schedule. As such if the blood services sought to eliminate these shortfalls a better approach would be to create a special panel containing only AB- donors that could be called upon at short notice to attend a special session.

6.2.2 Increasing Demand

The baby-boom generation is starting to age and it is anticipated that this generation will place a large strain on all healthcare services. Furthermore, as this generation ages, the donor pool will start to shrink, requiring the blood service to significantly grow the donor pool in order to keep up with demand. As a consequence of this, it would be of interest to explore the impact of shortening the inter-donation time, while also increasing the demand for blood. This section will compare how well the proposed shorter interval times cope with raising the overall demand targets.

The new demand targets to be considered can be split into two categories. The first will represent a slight increase to the demand, taking the demand from 65,677 units to 68,960 units which represents a 5% increase. This increase will be applied to both of the interval lengths (8 and 12 weeks as well as 10 and 14 weeks). The second category will have a much larger increase, taking the overall demand up to 78,812 units or a 20% increase. This increase aims to stress the registered donor pool in order to examine how successfully the schedule can manage with the increased demand.

This last increase shall only be applied to the longer of the two intervals (10 and 14 weeks). This is due to the computational difficulties experienced when calculating the same increase for the shorter interval. Though this result does not figure in the primary results table, the objective of the schedule proposed by the variable bounding heuristic -with a 20% increase in demand and with donors waiting 8 and 12 weeks- had a gap to the lower bound of over 20%. Therefore, meaningful comparisons cannot be made using this schedule and is not considered any further in this section.

The demand targets for these scenarios were calculated by taking the current demand for each period and blood group and increasing this by the appropriate percentage, this is then rounded down to two decimal places. All other parameters, including the size of the donor pools, are identical to the ones used in the scenarios in the previous section.

All three of the proposed scenarios - increasing both new intervals by 5% and increasing the 10 and 14 week interval by 20% - will be considered next starting by analysing each of the schedules' costs.

Table 6.11 presents the overall costs and the individual costs for each of the scenarios discussed above. The costs for both of the schedules with a 5% increase in demand are similar, with the schedule where donors wait for 10 and 14 weeks between donations costing slightly less. However, the majority of this difference is down to how well the second schedule is at minimising the shortfall cost, missing the collection targets by only 11 units compared to the 55 units under collected by the shorter interval's schedule. Naturally, as the second scenario collects more blood this schedule has a slightly higher variable cost, though as with all of the other scenarios produced by the session scheduling model, the variable costs make up the majority of the schedule's costs, in each of these scenarios variable costs account for more than 95% of the overall cost. Fixed costs are also very similar, making up 1.52% and 1.55% of the overall costs in the first and second scenario, respectively.

Cost	Demand increase of 5%, interval of 8 and 12 weeks (£)	%	Demand increase of 5%, interval of 10 and 14 weeks (£)	%	Demand increase of 20%, interval of 10 and 14 weeks (£)	%
Objective	8,755,830.97	100	8,719,614.77	100	10,027,485.14	100
Fixed cost	132,970.16	1.52	135,487.75	1.55	159,943.66	1.60
Variable cost	8,567,536.81	97.85	8,573,037.12	98.32	9,789,724.48	97.63
Penalty cost	55,324.00	0.63	11,089.90	0.13	77,817.00	0.78

Table 6.11: Scenario costs for different intervals and increases in demand

In spite of these similarities the schedule which rests donors for 10 and 14 weeks each appears to be the better choice of schedule, especially as the overall cost is cheaper and more importantly this schedule performs much better in terms of meeting demand targets.

The schedule with a 20% increase in demand has an overall cost of £10,027,485.14. Interestingly, this schedule performs relatively well missing the collection targets by only 77.8 units in spite of the large increase in demand. This suggests that shortening the inter-donation time along with a large raise in the overall demand is feasible, though there is a large raise in the overall costs and the shortfalls increase by non-negligible amounts. As in the schedules created above the variable costs are the largest cost and fixed costs have risen making up 1.6% of the schedule's overall cost. This raise is caused by the additional venues that are visited throughout this schedule.

Venues and Panels

The schedules where demand is increased only 5%, open 658 and 652 venues throughout the schedule, when donors are rested 8 and 12 weeks and 10 and 14 weeks, respectively. These schedules have an additional 30 venues open when compared to the schedules produced with the same intervals and current demand levels. Furthermore, the schedule that requires 20% more units to be collected opens a total of 740 venues utilising 97.5% of venue sessions available, which is significantly more than the previous schedules. Despite the increase in venues visited and as in the previous schedules looking at the shorter intervals applied to the current demand levels, no mobile venues are visited, as before, this is down to the much higher venue cost to capacity ratio.

Due to the rise in the number of venues open, the average utilisation of venue sessions in each period has naturally also increased. In the schedule produced using an interval length of 8 and 12 weeks and with a 5% increase in demand, eight out of the 40 periods have fully utilised the number of venue sessions available, where out of eight periods three still incur shortfalls. Period five incurred the highest shortfall, suggesting that if additional sessions had been available then the shortfall in this period could be reduced. For the schedule with donors waiting 10 and 14 weeks, there are also eight periods that have fully utilised the available venue sessions. Four of these periods are the same as in the previous schedule, although not all of these periods incur shortfalls. There are three periods in this schedule that incur shortfalls, however these amount to less than two whole units of blood.

As mentioned previously the schedule with a 20% demand increase opens significantly more venues and as such the number of periods in which no additional venues can be opened is higher. In total 34 periods have utilised all of the available venue sessions, of these periods nine incur shortfalls. Unlike the previous schedules this entire schedule's shortfall occurs in periods that cannot open any additional venues, again suggesting that additional venue sessions might alleviate a significant portion of this shortfall.

Figure 6.2 displays the number of times each of the available fixed venues are visited for all three of the scenarios considered in this section. The venues opened and how frequently they are visited for the schedules with a 5% demand

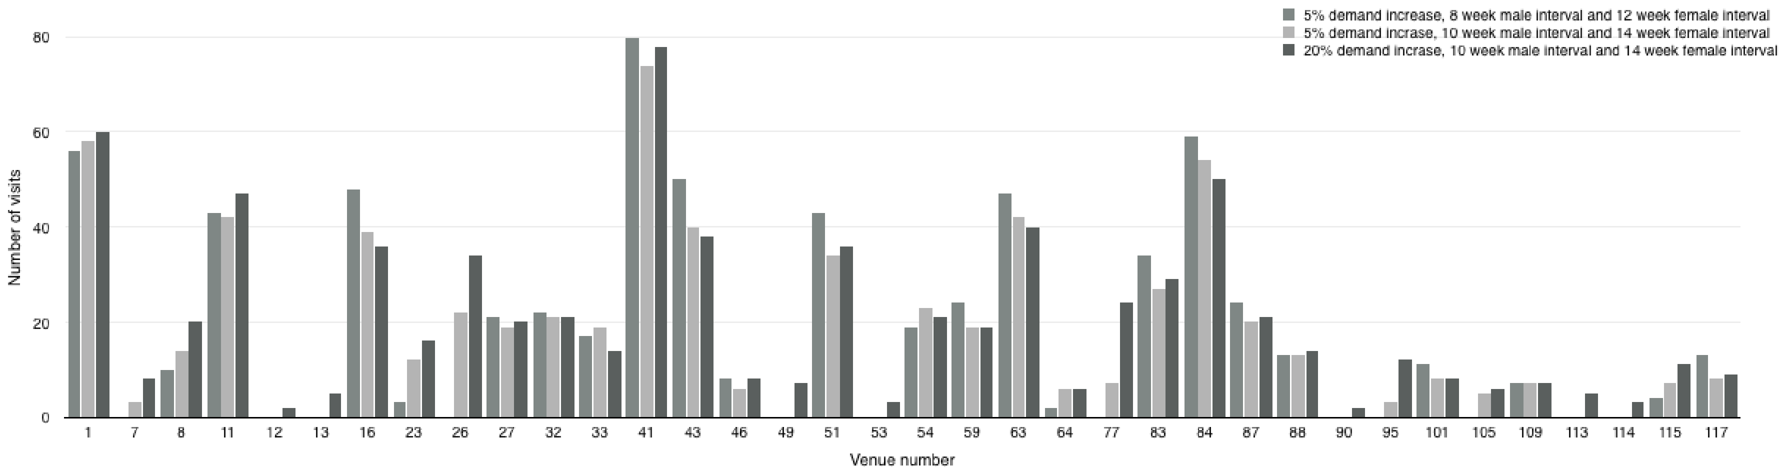

Figure 6.2: Number of venues visited for scenarios with different intervals and increased demand (Only venues visited at least once are shown)

increase have similar characteristics to those opened in the shorter interval schedule without the demand increase.

Venue 41 is visited the maximum number of times in the schedule with the shortest interval and a slight demand increase, although this is not the cheapest venue in terms of the venue cost to capacity ratio. As the schedule with the 20% demand increase opens significantly more venues than the previous schedule large increases in how frequently certain venues are visited can be observed in the aforementioned chart - in particular venues 8, 26, 77 and 95. Furthermore, this schedule opens seven new venues that do not figure in either of the previous schedules, though none of these venues are visited more than seven times. Despite these venues possessing a higher venue cost to capacity ratio they are necessary to collect sufficient blood and minimise the penalty cost.

As with previous schedules the schedule proposed in this section only use a fraction of the panels available, with the schedules produced with a 5% increase in demand using on average only 18% of panels, while the schedule with a 20% increase uses a total of 36 panels or in percentage form 25% of the panels available. In each of these scenarios the most visited panel is panel P073, which in each case is visited more than 70 times. This panel serves one of the most populated areas in the region with over 1,000 registered donors available, contributing to the panel's frequent visits.

The panels that contribute the most to the number of units collected for each of these scenarios is presented in Table 6.12. In the scenario with the shortest interval and a 5% increase in demand each panel on average contributes 2,871

5% demand increase, 8 and 12 week intervals	
Panel	Contribution to collection target(%)
45	9.48
44	8.78
9	8.31
20	7.78
27	7.68
21	7.32
1	7.31
34	6.97
6	5.99
46	5.62
15	3.51
32	3.36
13	2.9
29	2.7
5	2.51
16	2.08
47	1.72
132	1.31
102	1.21
143	1.1
22	0.84
139	0.66
11	0.56
35	0.31

5% demand increase, 10 and 14 week intervals	
Panel	Contribution to collection target(%)
45	8.27
1	7.57
20	7.15
44	7
9	6.91
27	6.23
34	6.22
21	5.88
6	5.85
46	4.45
32	3.45
5	3.42
12	3.38
15	3.15
11	2.86
29	2.81
13	2.75
16	1.95
47	1.58
41	1.44
132	1.27
139	1.12
102	1.1
143	0.87
35	0.8
120	0.76
52	0.75
4	0.45

20% demand increase, 10 and 14 week intervals	
Panel	Contribution to collection target(%)
45	7.02
1	6.84
20	6.32
44	6.03
9	5.69
27	5.57
6	5.43
34	5.3
21	5.06
12	4.5
41	4.24
5	3.98
46	3.96
11	3.26
32	3.08
15	2.72
13	2.52
29	2.49
52	1.98
16	1.49
47	1.44
139	1.41
132	1.18
24	1.1
4	1.05
102	0.99
120	0.85
35	0.74
143	0.71
22	0.67
137	0.63
49	0.55
138	0.38
28	0.37
7	0.25
48	0.2

Table 6.12: Contribution per panel for scenarios with intervals of either 8 and 12 weeks or 10 and 14 weeks and demand increases of either 5% or 20%

units, which is an average of approximately 4.2% of the total number of units collected per panel. For the slightly longer interval with the same demand level this average decreases slightly to 2,378 units or 3.4% of units collected. The scenario that has the much larger increase in demand collects an average of 2,187 units per panel visited, which equates to appropriately 2.78% of the units collected. This decrease in units collected per panel is due to the additional panels that are required to be visited so that the shortfall costs are minimised.

As was discussed in previous scenarios remarkably few panels needed to be visited to account for a large portion of the units collected. This is again the case with the scenarios described in this section, with only nine panels needing to be visited to amass over 53% of the units collected in the scenario with a 20% increase in demand. This result is similar to the other scenarios with eight panels being visited gathering 55% of units collected for donors resting 10 and 14 weeks with a 5% demand increase and the remaining scenario collecting 57% of the units from only seven panels. These points further highlight that in spite of the increase in demand only a fraction of the available panels are actually used.

Donors

This subsection examines how donors are assigned in each of the scenarios that have an increased collection target. First, an overview of the number of units collected for each schedule is presented before moving on to analyse the possible causes behind each schedule's shortfalls.

An overview of the units collected by the two scenarios that included a 5% increase in the demand targets are presented first. In total the schedule with the shortest interval collected a total of 68,904.11 units, of which 43,489.06 units where donated by registered donors and the remaining 25,415.05 were collected from walk-in donors. In order to collect these units a total of 60,356 registered donors were invited, yielding a turn out rate of just over 72% and averaging 1,723.71 units a period. The most units were collected in period 35, with a total of 2,203.9 units collected and all collection targets being met. As for the schedule produced with male and female donors resting 10 and 14 weeks respectively, a total of 60,233 registered donors were invited. From these donors a total of 43,539 units were collected which again represents a turnout rate of around 72%. As in the first scenario the most units collected was in period 35 with all collection targets being met, furthermore both schedules average a very similar number of units collected in each period.

The scenario that included a 20% raise in the overall demand collected a total of 78,733.51 units of blood, of which 28,802 were from walk-in donors. The registered donors totalled 49,933 donors which were made up of 28,535 male and 21,398 female donors. This scenario required each period to collect an average of 1,968.34 units an extra 327 units per period than in scenario 3.

Despite the increases in the number of units collected and as mentioned when analysing the objectives each schedule falls short of their collection targets and as such incur shortfall penalties. The blood groups that cause these shortfalls are the AB blood types, the same blood groups that incur shortfalls in previous scenarios. Although the same blood types incur shortfalls the utilisation rate of venue sessions is higher in these scenarios - especially with a 20% increase in demand. Therefore, it is still necessary to explore the causes of these shortfalls.

For each of the scenarios Tables 6.13, 6.14 and 6.15 present the shortfall incurred for the AB blood types by period. For the scenario with a 5% demand increase and with male and female donors waiting 8 and 12 weeks respectively there is a total shortfall of 55 units, of which the vast majority are for the AB+ blood type. Notably, only periods 5 and 29 incur shortfalls for the AB- blood type, though period 29 incurs a negligible amount - a total of 0.03 units of shortfall. Period 5 also incurs the highest shortfall for the AB+ blood type, accounting for 65% of this blood types total shortfall. Furthermore, period 29 incurs the second highest shortfall with a total of shortfall of 6.29 units.

The maximum number of venues (18), which is slightly below the average, were opened in period five. This suggests that given the majority of this schedule's shortfall occurred in this period, increasing the number of venue sessions available would be beneficial. Incidentally, there are still unassigned donors of the correct blood types available in this period. However, period 29 still has venue sessions available and incurs the second highest shortfall, due to a lack of availability of donors with the correct blood types. Therefore, visiting venues and assigning donors to donate in the appropriate periods is crucial to ensuring large shortfalls do not occur in a schedule. In the schedule with a 5% demand increase and male donors waiting 10 and females 14 weeks, only the AB- blood group produced a shortfall. A total of 11 units were under collected with periods

Period	AB+	AB-	Shortfall
1	-	-	-
2	-	-	-
3	-	-	-
4	-	-	-
5	29.62	9.74	39.35
6	-	-	-
7	-	-	-
8	-	-	-
9	-	-	-
10	-	-	-
11	-	-	-
12	-	-	-
13	-	-	-
14	-	-	-
15	-	-	-
16	-	-	-
17	-	-	-
18	-	-	-
19	0.94	-	0.94
20	0.53	-	0.53
21	3.28	-	3.28
22	0.65	-	0.65
23	-	-	-
24	0.48	-	0.48
25	1.48	-	1.48
26	-	-	-
27	-	-	-
28	-	-	-
29	6.29	0.03	6.32
30	0.51	-	0.51
31	-	-	-
32	1.78	-	1.78
33	-	-	-
34	-	-	-
35	-	-	-
36	-	-	-
37	-	-	-
38	-	-	-
39	-	-	-
40	-	-	-

Table 6.13: Shortfall for scenario where male donors donate every 8 weeks and female donors donate every 12 weeks and with an overall demand increase of 5%.

12 and 31 contributing the most - over three units each, despite having venue sessions available in each period. The cause for these shortfalls are similar to those discussed in previous sections. In some periods there is simply a shortage of donors with the AB- blood type, as such opening additional venues would be fruitless. In others there are donors and venue sessions available though by assigning donors in these periods larger shortfalls occur in other periods due to the number of weeks a donor has to rest.

Period	AB-
9	0.30
11	1.29
12	3.10
14	1.17
18	0.35
27	0.45
29	0.46
30	0.25
31	3.72

Table 6.14: Shortfall for scenario where male donors donate every 10 weeks and female donors donate every 14 weeks and with an overall demand increase of 5%.

When the overall demand is increased by 20%, shortfalls can be observed for both types of AB blood. The AB+ blood has a total shortfall of 34.62 units, with AB- blood having approximately 8 additional units of shortfall totalling 43.19 units. In total this schedule incurs fractionally over 77 units of shortfall, with the highest combined shortfall of just over 17 units occurring in period nine. Interestingly, every period that incurs a shortfall in this schedule has utilised all of the available venues' sessions. Furthermore, by adding one additional venue slot in each of the periods that incurs a shortfall and opening an extra venue

Period	AB+	AB-	Shortfall
1	-	-	-
2	-	3.17	3.17
3	-	-	-
4	-	-	-
5	-	-	-
6	-	-	-
7	-	-	-
8	-	-	-
9	9.49	7.61	17.09
10	-	-	-
11	5.06	6.42	11.47
12	-	-	-
13	-	-	-
14	6.51	7.49	13.99
15	-	-	-
16	-	-	-
17	-	-	-
18	-	-	-
19	-	-	-
20	-	-	-
21	-	-	-
22	-	-	-
23	-	2.51	2.51
24	-	-	-
25	-	-	-
26	-	-	-
27	10.95	2.49	13.44
28	2.63	9.94	12.57
29	-	-	-
30	-	0.37	0.36
31	-	3.21	3.21
32	-	-	-
33	-	-	-
34	-	-	-
35	-	-	-
36	-	-	-
37	-	-	-
38	-	-	-
39	-	-	-
40	-	-	-

Table 6.15: Shortfall for scenario where male donors donate every 10 weeks and female donors donate every 14 weeks and with an overall demand increase of 20%.

in these periods, it is possible to further reduce the shortfall. Therefore, if the blood service wish to significantly increase their collection targets additional venue's sessions would be necessary to avoid large shortfalls developing.

6.3 Summary

In this chapter, two categories of scenarios were explored. The first set of scenarios looked at ways to increase donor reliability and included compensating regular donors for their donations. It was found that when paying regular donors £20 for each unit donated and increasing their reliability to 90%, the overall shortfall could be reduced to slightly over 10 units (compared to over 5,000 units in the blood service's current schedule and 26 units in the schedule proposed by the blood collection model). However, the overall cost of this scenario (£9,193,600.61) is 12% higher than the schedule produced by the blood collection model (£8,235,659.03) for an increase of just 16 units. For this reason, it is currently advisable not to compensate donors for their donations and instead a focus on better donor management, would be far more beneficial to the blood service.

The second set of scenarios examined the possible impact of altering the inter donation interval for donors. Schedules were produced for current demand levels, where either male donors donated every 8 weeks and female donors donated every 12 weeks or where male donors donated every 10 weeks and female donors donated every 14 weeks. Further schedules were produced for these different inter donation times when total demand was increased by 5%. Finally a schedule was produced where overall demand was increased by 20% and donors could donate every 10 and 14 weeks for male and female donors, respectively.

It was found that for the schedule where demand was not increased and donors donated every 10 and 14 weeks for male and female donors respectively, the penalty incurred was less than £7,000. This compares very favourable to the penalty (over £26,000) incurred under the current inter donation interval (12 weeks for male donors and 16 weeks for female donors). When demand is increased by 5%, for a donor interval of 10 and 14 weeks for male and female donors, the penalty increased to £11,000. This indicates that a shorter inter donation time could compensate for a slight raise in overall demand. However, when demand is increased by 20%, for the same inter donation times, then the overall penalty increases to over £70,000. In both of these situations, the shortfalls arise due to the lack of donors with the AB blood group. These results suggest that reducing the inter donation interval could be very beneficial to a blood service, however, there is a need to recruit further donors of blood group AB.

Chapter 7

Conclusions and Suggestions for Further Research

7.1 Conclusion

This section is split into two sub sections. The first provides an overview of the entire thesis and highlights key findings. The second sub section discusses some of the limitations to this research and areas of interest for future research.

At the start of this thesis, the blood banking literature was examined considering both 1) more general works relating to inventory management, which includes issuing and crossmatching policies, blood redistribution, the impact the age of blood has on the supply chain, and forecasting future needs for blood products, and 2) more focused papers relating to blood collection, including schedule creation, the donor's experience, and locating blood facilities. Despite the numerous works in the blood banking area, there are still important gaps in the literature. Among the gaps identified are the lack of research on cost-efficiently scheduling blood collection, locating mobile blood centres, and understanding the impact donation intervals have on collection schedules.

Currently, blood services lack sophisticated models that can help them plan schedules, evaluate how effective those schedules are, and carry out what-if experiments without impacting current operations. As such, the general aim of this thesis was to create a model that enables blood services to better understand their current schedules and enable them to develop cost-effect alternatives that best utilise limited resources.

To meet this need and to help fill some of the research gaps in the literature, a blood collection model was formulated in Chapter 3. This model offers several advances over previous works. The objective is to minimise the total cost of the schedule. The model also keeps track of the donors available in each panel and ensures that they wait a minimum amount of time between donations. Importantly, the model also distinguishes between male and female donors and donors of different blood groups. This allows blood services to understand any shortfalls in more detail. Finally, the model allows what-if scenarios to be easily run and for the efficiency of different schedules to be measured.

The schedule operated by the NBS in Kent between December 2012 and October 2013 was evaluated, using the data provided by the service. This region offered a mixture of both rural villages and more densely populated areas. It was found that the schedule currently operated by the blood service was relatively inefficient. There was significant over collecting in certain periods that led to shortfalls later on in the schedule, indicating that donor management could be improved. It was also found that the blood service visited every panel at least once. While this ensures that all donation areas are visited during the schedule, it is not necessarily cost efficient.

The schedule generated by the blood collection model differed significantly from the schedule currently operated in Kent. The schedule generated by the model was significantly cheaper than the current service. The majority of this reduction is down to the model not over collecting blood and not depleting donors early on in the planning horizon. The newly generated schedule also visited significantly less venues, though those venues were visited more often. The most frequently visited venues were those that were near a large population and had venues with a low venue cost to capacity ratio. Interestingly, this new schedule did not make use of mobile venues. This is due to their high cost and low overall capacity. This suggests that the blood service should focus on venues with lower cost to capacity ratios that are near densely populated areas. The blood service should also consider how efficient blood mobiles are and how frequently they are used.

Having generated a new schedule, several what-if scenarios were considered, a group of scenarios considered if it was worth remunerating regular donors for their donations. These scenarios assumed that regular donors would be remunerated £20 for every unit donated and that with remuneration, regular donors would become more reliable. Scenarios were run with donor reliability levels increased to 75% in the first instance and then to 90% in the second instance. Although it was possible to reduce the shortfall to only 11 units, this required a 90% donor reliability rate and even with such a high reliability rate it was not possible to eliminate the shortfall relating to the AB blood group. These scenarios were also significantly more expensive that the others, owing to the added donor remuneration cost. Based on this, it was judged that remunerating donors was not advantageous to the blood service.

Another group of scenarios considered the impact changes to the inter-donation interval could have on donor availability and how these changes could affect collection schedules. This group of scenarios was further split into two categories, the first set of scenarios considered changes with demand levels unchanged, while the second set considered demand increases of 5% and 20% above current demand levels.

By reducing the inter-donation time to 10 weeks for male and 14 weeks for female donors, under current demand levels, it was possible to produce a schedule that incurred a shortfall of only six units. Reducing the inter donation interval, would therefore be advantageous from an operational perspective. Similarly to the other schedules produced by the blood collection model, this schedule does not visit any mobile venues and visits relatively few panels, though much more frequently, compared to the blood service's current schedule.

When demand was increased by 5% and male donors wait 10 weeks and female donors wait 14 weeks between donations then it was possible to generate a schedule that only had 11 units of shortfall. Furthermore, it was established that with a much larger increase of 20% and the same interval the shortfall would increase to 77 units. This suggest that altering the inter-donation interval could be beneficial to the blood service and that an interval of 10 weeks for male donors and 14 weeks for female donors, would allow for an overall demand increase of 5%, with a shortfall that is less that the current shortfall and importantly without needing to grow the current donor pool.

This thesis raises important questions about the number of venues that the blood service should operate and how frequently they should be visited, it also raises questions about the role of bloodmobiles and highlights the operational advantages of reducing the inter-donation interval.

7.2 Limitations and further research

Naturally there are limitations to the current research and areas that are worthy of further research. This section considers both of these areas in more detail.

The simulated annealing solution method proposed in this thesis underperformed when solving large instances of the blood collection model. This could be down to the fact that the simulated annealing algorithm implemented is very basic. It would be of interest to see if a more sophisticated implementation would perform better. Suggested improvements include incorporating multiple restarts, in which the heuristic is run starting from different initial solutions. Another improvement is to change from a fixed cooling schedule to a dynamic cooling schedule, enabling the temperature to decrease in a non-geometric fashion.

Additionally, it is necessary to have accurate data to ensure that the results being generated by the models are reliable. While the data used in this thesis was either provided directly by the blood service or based on their opinion, there were certain parameters that could have benefited from a finer level of detail. As is recommended in Chapter 4, the blood service should start storing easily accessible data on the cost of collecting blood and donor histories. Improved data accuracy would naturally lead to a more precise model.

There are several extensions that could be made to the model. These extensions either increase the model's accuracy, in terms of modelling the real world, or enable the blood service to consider a different operating structure.

For example, the current model assumes that any blood that is over-collected in one period (a week in this model) cannot be used in subsequent periods. Any over-collected blood is simply destroyed at the end of the current period. However, blood has a 30 day shelf life, so that any over-collected blood in one period could be carried over to the following period for up to a maximum of 4 periods (weeks) before needing to be destroyed. It would be of interest to extend the current blood collection model to incorporate this feature. This would more closely model what happens in the real world and enable better comparison to be made against the current blood service's schedule, where over-collecting is common.

Another extension would be to change the constraints and parameters controlling the number of walk-in donors that could be invited to donate in a given panel. The updated constraints and parameters could incorporate time-dependent probabilities, whereby the number of walk-ins available in each panel diminish depending on the length of time between visits and seasonal effects for that panel, therefore more accurately reflecting what happens in the real world.

Similarly, it would be of interest to explore how the likelihood of donors donating changes with the day of the week. This would enable blood services to better schedule their sessions to ensure the maximum number of donors are available on that day. For example, it might be better to schedule a session in a city on a weekend rather than a weekday.

Another interesting area for future research would be to remove the panel structure from the model. In a panel-less model, a donor would be able to donate at any venue, however, the donor's likelihood to attend venues that are further away would diminish rapidly as the distance increased. The reason this area would be of interest is due to the fact that a panel is only an abstract structure for grouping donors and venues together, a product created to help manage donors. Removing the panel constraint would allow the blood service to invite donors based on how close they are to a given venue and how likely they are to attend. This change would allow the blood service to evaluate the benefits of moving away from a panel based system and to better understand where venues should be located, ideally close to clusters of donors.

Finally, a more general topic for future research that applies to the blood supply research carried out in the UK would be to investigate how widespread the implementation of decision support tools and models are. If, as suspected, there has been little adoption of these methods, it would be of interest to explore the reasons for this and how these obstacles can be overcome.

Appendix A

Simulated annealing

Simulated annealing is a probabilistic heuristic based on the physical process of annealing metals. Annealing in metallurgy is the process of heating up a metal until in a liquid state and then slowly cooling the metal. As the metal cools, larger and defect-free crystals are encouraged to form, resulting in a less ductile metal once returned to its solid state. Put another way, the overall objective of annealing is to bring the material down to a state with the lowest free energy by slowly cooling the material in a controlled way. Annealing is a topic of interest to physicists whom have sought to produce algorithms to allow more efficient analysis of this process; such algorithms include the Metropolis-Hastings algorithm.

The Metropolis-Hasting algorithm is a Markov chain Monte Carlo simulation. It was initially proposed by Metropolis et al. (1953) to calculate the possible states of a material as it was cooled. The major contribution of this paper was the idea that instead of calculating the properties of interest for a large number of random configurations and then weighting them against a probability distribution, a configuration is accepted and evenly weighted based on the probability distribution. Therefore, allowing statistical mechanical problems to be more efficiently analysed. However, their algorithm only considered the system's possible states for rigid-spheres in two dimensional space. The algorithm was later generalised by Hastings (1970), resulting in the algorithm being named after both individuals.

The suggestion that this technique could be used in global optimisation problems came from Kirkpatrick et al. (1983), who adapted the Metropolis-Hastings algorithm to better suit the more general nature of the problems studied in the optimisation field. Simulated annealing has become an oft used heuristic for approximating optimisation solutions, with applications ranging from timetabling and exam scheduling to travelling salesman type problems. Several adaptations to the general simulated annealing heuristic can be found throughout the literature including; adaptive simulated annealing, quantum annealing and simulated quenching.

Algorithm Procedure

A heuristic consists of three primary phases. The construction phase, the local search phase and the stopping phase. These phases are considered for the simulated annealing algorithm next.

An initial solution is a starting solution for which the initial energy (objective) can be calculated and moreover serves as a starting benchmark against which the first local search move can be compared. There are numerous ways to create an initial solution with varying degrees of complexity and time requirements. Perhaps the easiest approach to creating an initial solution, where feasible, is simply to start with an empty solution. Although this technique will cost very little time in the construction phase of the algorithm the search phase could be significantly hindered.

A more prevalent approach when creating starting solutions is to create a partial or complete initial solution. Creating a complete solution has the advantage of allowing the local search to start exploring neighbours close to an already feasible solution. However, when creating a complete solution it is computationally highly taxing or is too restrictive on the local search function, creating a partial solution can be particularly helpful. This approach, allows the local search phase to benefit from starting with a partial solution and without drastically impacting the overall running time of the heuristic.

There are a variety of different ways in which an initial starting solution can be populated, with random, greedy, or a blend of both approaches being amongst the most common. A random approach will simply attempt to add a randomly selected feature to the initial solution and re-evaluate the energy in the system. A purely greedy approach meanwhile attempts to add the best feature currently available to the starting solution.

While, either of these two approaches is sufficient to create an initial starting solution, both approaches have downsides that need to be considered. For example, a random initial solution may make it hard for the local search algorithm to converge to an optimal solution, while a purely greedy approach may suffer computationally from the time needed to evaluate all of the possible features that can be added to the current solution. A hybrid approach, including both the random and greedy elements, allows solutions with sufficient variety to be created comparatively quickly. This is often done by calculating the impact of certain features on the objective and then randomly selecting one of the candidate features to be added to the current initial solution.

Once an initial starting solution has been constructed the algorithm's next phase is to perform a local search based on the newly constructed solution. Each local search iteration consists of creating a neighbour to the current solution and evaluating its impact on the overall objective. Creating a new neighbour to consider can be done by adding or dropping one or more features from the current state or by permuting two or more features already in the solution.

After a new feature has been added to the solution the energy for this state is calculated; an improving move, one which reduces the energy in the system, is always accepted. However, what makes the simulated annealing algorithm unique is that non-improving moves are also accepted with a certain probability. The probability distribution that controls how frequently a non-improving move is accepted is determined by the algorithm's cooling schedule which is discussed in more detail below.

The advantage of allowing worse moves to be accepted into the current solution is that it aides the heuristic to explore more of the search space and crucially helps the heuristic from getting stuck at local optima. For a visual representation of this the reader is referred to the end of this section.

There are several ways to define a stopping condition for this algorithm, which is the final phase of this heuristic before reporting the best solution. The most common stopping conditions for the algorithm consist of halting the search either once a maximum number of iterations have been performed or once the system has reached a certain temperature. Although it is also feasible to use a time-constraint as a crude alternative, or, slightly more sophisticated, would be to halt the heuristic once a certain number of non-improving moves have been performed.

Algorithm A.1 provides the general pseudo code for the simulated annealing algorithm when seeking to minimise the total energy of a system.

The first few lines assign the initial temperature and ensure that the current best solution has infinite energy. Line 4 creates an initial solution; then on line 6 a new neighbour is created. Once a neighbour has been created lines 7 through 14 calculate the cost of the new neighbour solution and if it is an improving move it is unconditionally accepted and then compared against the current best solution. Again, if the neighbour solution is an improvement the best solution is updated. If the new neighbour is not an improving move it is accepted with a certain probability, which is a function of the system's current temperature.

This algorithm requires the following inputs: the maximum number of iterations to perform, the system's starting temperature and the problem data.

The algorithm will continue to create and evaluate neighbour solutions until on line 5 either the maximum number of iterations has been reached or the system has cooled to the desired temperature. At the end of each iteration, the system's current temperature is decreased by the cooling factor, see line 15. As a last step, the algorithm outputs the best solution, as seen on line 16.

Algorithm A.1 Simulated annealing

```
 1. Input: $Iterations$, $Temp_{starting}$ $Data$
 2. $Temp \leftarrow Temp_{starting}$
 3. $Sol_{best} \leftarrow \infty$
 4. $Sol_{current} \leftarrow$ CreateInitialSolution(Data)
 5. For($Iterations \,||\, Temp \leq 0$)
 6.   $Sol_{new} \leftarrow$ CreateNeighbour($Sol_{current}$)
 7.   If(Cost($Sol_{new}$) $\leq$ Cost($Sol_{current}$))
 8.     $Sol_{current} = Sol_{new}$
 9.     If(Cost($Sol_{new}$) $\leq$ Cost($Sol_{best}$))
10.       $Sol_{best} = Sol_{new}$
11.     End
12.   Else If(AcceptSolution($Sol_{new}$)
13.     $Sol_{current} = Sol_{new}$
14.   End
15.   $Temp$=NewTemp($Temp, iteration$)
16. End
17. Return $Sol_{best}$
```

As a simple extension, it is possible to alter the algorithm to include multiple restarts. This would enable the algorithm to perform local searches on multiple initial solutions which could yield better results, though at the cost of additional running time.

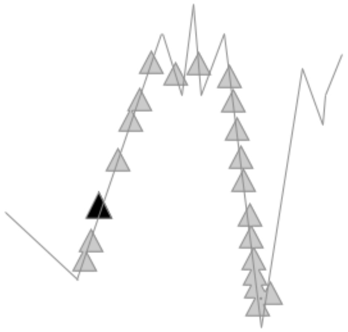

Figure A.1: Simulated annealing visualisation

To aide in visualising how the simulated annealing algorithm works consider Figure A.1. In this minimisation scenario the starting solution is indicated by the black triangle. In the early stages, when the temperature is high, the algorithm will frequently accept moves that are worse than the current solution. This ability to accept non-improving solutions however enables the simulated annealing heuristic to avoid getting stuck in the numerous local optimal in the solution space (the valleys), that would otherwise have trapped a simple hill decent algorithm. However, as the algorithm proceeds the likelihood of accepting a non-improving move becomes less and less likely with the heuristic eventually terminating at the optimal solution. A path that the heuristic could take is indicated by the grey triangles in Figure A.1.

Bibliography

Alfonso, E., Augusto, V., and Xie, X. (2015). Mathematical programming models for annual and weekly bloodmobile collection planning. *Automation Science and Engineering, IEEE Transactions on*, 12(1):96–105.

Alfonso, E., Xie, X., Augusto, V., and Garraud, O. (2012). Modeling and simulation of blood collection systems. *Health care management science*, 15(1):63–78.

Alfonso, V. A. and Xie, X. (2013). Tactical planning of bloodmobile collection systems. In *Automation Science and Engineering (CASE), 2013 IEEE International Conference on*, pages 26–31. IEEE.

Ali, A., Auvinen, M.-K., and Rautonen, J. (2010). Blood donors and blood collection: The aging population poses a global challenge for blood services. *Transfusion*, 50(3):584–588.

An, M.-W., Reich, N. G., Crawford, S. O., Brookmeyer, R., Louis, T. A., and Nelson, K. E. (2011). A stochastic simulator of a blood product donation environment with demand spikes and supply shocks. *PloS one*, 6(7):1–9.

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

Blake, J. and Hardy, M. (2013). Using simulation to evaluate a blood supply network in the canadian maritime provinces. *Journal of Enterprise Information Management*, 26(1/2):119–134.

Bosnes, V., Aldrin, M., and Heier, H. E. (2005). Predicting blood donor arrival. *Transfusion*, 45(2):162–170.

Brennan, J. E., Golden, B. L., and Rappoport, H. K. (1992). Go with the flow: improving red cross bloodmobiles using simulation analysis. *Interfaces*, 22(5):1–13.

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

Brodheim, E., Hirsch, R., and Prastacos, G. (1976). Setting inventory levels for hospital blood banks. *Transfusion*, 16(1):63–70.

Brodheim, E. and Prastacos, G. P. (1979). The long island blood distribution system as a prototype for regional blood management. *Interfaces*, 9(5):3–20.

Cerveny, R. P. (1980). An application of warehouse location techniques to bloodmobile operations. *Interfaces*, 10(6):88–96.

Chapman, J., Milkins, C., and Voak, D. (2000). The computer crossmatch:

a safe alternative to the serological crossmatch. *Transfusion Medicine*, 10(4):251–256.

Chazan, D. and Gal, S. (1977). A markovian model for a perishable product inventory. *Management Science*, 23(5):512–521.

Chiavetta, J., Herst, R., Freedman, J., Axcell, T., Wall, A., and Rooy, S. v. (1996). A survey of red cell use in 45 hospitals in central ontario, canada. *Transfusion*, 36(8):699–706.

Cohen, Pierskalla, and Sassetti (1981). Regional blood inventory control and distribution. *Proceedings of the 1980 Conference on the Management and Logistics of Blood Banking*, 5:21–88.

Cohen, M., Pierskalla, W., and Sassetti, R. (1983). The impact of adenine and inventory utilization decisions on blood inventory management. *Transfusion*, 23(1):54–58.

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

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

Cumming, P. D., Kendall, K. E., Pegels, C. C., Seagle, J. P., and Shubsda, J. F. (1976). A collections planning model for regional blood suppliers: description and validation. *Management Science*, 22(9):962–971.

Currie, C., Patel, T., McEwan, P., and Dixon, S. (2004). Evaluation of the future supply and demand for blood products in the united kingdom national health service. *Transfusion Medicine*, 14(1):19–24.

De Angelis, V., Felici, G., and Impelluso, P. (2003). Integrating simulation and optimisation in health care centre management. *European Journal of Operational Research*, 150(1):101–114.

Denesiuk, L., Richardson, T., Nahirniak, S., and Clarke, G. (2006). Implementation of a redistribution system for near-outdate red blood cell units. *Archives of pathology & laboratory medicine*, 130(8):1178–1183.

Deuermeyer, B. L., Pierskalla, W. P., Sassetti, R. J., and Director, B. B. (1976). Methods for reducing outdating without altering physician ordering patterns. Technical report, Working Paper, Department of Industrial Engineering, Northwestern University.

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

Dumas, M. B. and Rabinowitz, M. (1977). Policies for reducing blood wastage in hospital blood banks. *Management Science*, 23(10):1124–1132.

Fontaine, M. J., Chung, Y. T., Erhun, F., and Goodnough, L. T. (2010). Age of blood as a limitation for transfusion: potential impact on blood inventory and availability. *Transfusion*, 50(10):2233–2239.

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

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

Georgsen, J. and Kristensen, T. (1998). From serological to computer cross-matching in nine hospitals. *Vox sanguinis*, 74(S2):419–425.

Glynn, S. A., Busch, M. P., Schreiber, G. B., Murphy, E. L., Wright, D. J., Tu, Y., Kleinman, S. H., et al. (2003). Effect of a national disaster on blood supply and safety. *JAMA: the journal of the American Medical Association*, 289(17):2246–2253.

Goyal, S. and Giri, B. C. (2001). Recent trends in modeling of deteriorating inventory. *European Journal of operational research*, 134(1):1–16.

Gunpinar, S. (2013). *Supply Chain Optimization of Blood Products*. PhD thesis, University of South Florida.

Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109.

Heddle, N. M., Liu, Y., Barty, R., Webert, K. E., Whittaker, S., Gagliardi, K., Lauzon, D., and Owens, W. (2009). Factors affecting the frequency of red blood cell outdates: an approach to establish benchmarking targets. *Transfusion*, 49(2):219–226.

Hoyt, P., Hutchins, J., and Lewis, D. (1996). Improvement of the us navy mobile blood bank through simulation analysis and forecasting. Technical report, DTIC Document.

Jabbarzadeh, A., Fahimnia, B., and Seuring, S. (2014). Dynamic supply chain network design for the supply of blood in disasters: A robust model with real world application. *Transportation Research Part E: Logistics and Transportation Review*, 70:225–244.

Jacobs, D. A., Silan, M. N., and Clemson, B. A. (1996). An analysis of alternative locations and service areas of american red cross blood facilities. *Interfaces*, 26(3):40–50.

Jagannathan, R. and Sen, T. (1991). Storing crossmatched blood: a perishable inventory model with prior allocation. *Management science*, 37(3):251–266.

James, R. and Matthews, D. (1996). Analysis of blood donor return behaviour using survival regression methods. *Transfusion Medicine*, 6(1):21–30.

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

Kendall, K. E. (1980). Multiple objective planning for regional blood centers. *Long Range Planning*, 13(4):98–104.

Kendall, K. E. (1984). Planning for innovation in health care: a breakthrough in blood banking. *Long Range Planning*, 17(6):127–131.

Kendall, K. E. and Lee, S. M. (1980). Formulating blood rotation policies with multiple objectives. *Management Science*, 26(11):1145–1157.

Kirkpatrick, S., Gelatt, C. D., Vecchi, M. P., et al. (1983). Optimization by simulated annealing. *science*, 220(4598):671–680.

Lowalekar, H. and Ravichandran, N. (2010). Model for blood collections management. *Transfusion*, 50(12pt2):2778–2784.

Mathoulin-Pélissier, S., Salmi, L. R., Verret, C., and Demoures, B. (2000). Blood transfusion in a random sample of hospitals in france. *Transfusion*, 40(9):1140–1146.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092.

Millard, D. (1960). Industrial inventory models as applied to the problem of inventorying whole blood, ohio state university engineering experiment station. *Bulletin*, (180).

Moore, C., Sambrook, J., Walker, M., Tolkien, Z., Kaptoge, S., Allen, D., Mehenny, S., Mant, J., Di Angelantonio, E., Thompson, S. G., et al. (2014). The interval trial to determine whether intervals between blood donations can be safely and acceptably decreased to optimise blood supply: study protocol for a randomised controlled trial. *Trials*, 15(1):363.

Nahmias, S. (1975a). A comparison of alternative approximations for ordering perishable inventory. *Infor*, 13(2):175–184.

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

Nahmias, S. (1982). Perishable inventory theory: A review. *Operations Research*, 30(4):680–708.

Nahmias, S. (2011). *Perishable inventory systems*, volume 160. Springer.

Nahmias, S. and 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.

Nahmias, S. and Pierskalla, W. P. (1976). A two-product perishable/nonperishable inventory problem. *SIAM Journal on Applied Mathematics*, 30(3):483–500.

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

Osorio, A. F., Brailsford, S. C., and 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*, (ahead-of-print):1–22.

Ownby, H., Kong, F., Watanabe, K., Tu, Y., and Nass, C. C. (1999). Analysis of donor return behavior. *Transfusion*, 39(10):1128–1135.

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

Pegels, C. C., Seagle, J., Cumming, P., Kendall, K., and Shubsda, J. (1977). An analysis of selected blood service policy changes. *Medical care*, 15(2):147–157.

Pegels, C. C., Seagle, J. P., Cumming, P. D., and Kendall, K. E. (1975). A computer-based interactive planning system for scheduling blood collections. *Transfusion*, 15(4):381–386.

Pereira, A. (2011). Modeling the effect of reducing maximum shelf life on red blood cell availability. *Transfusion*, 51(3):662–662.

Perera, G., Hyam, C., Taylor, C., and Chapman, J. (2009). Hospital blood inventory practice: the factors affecting stock level and wastage. *Transfusion Medicine*, 19(2):99–104.

Pierskalla, W. P. (2005). Supply chain management of blood banks. In *Operations research and health care*, pages 103–145. Springer.

Pierskalla, W. P. and Roach, C. D. (1972). Optimal issuing policies for perishable inventory. *Management Science*, 18(11):603–614.

Pitocco, C. and Sexton, T. R. (2005). Alleviating blood shortages in a resource-constrained environment. *Transfusion*, 45(7):1118–1126.

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

Price, W. L. and Turcotte, M. (1986). Locating a blood bank. *Interfaces*, 16(5):17–26.

Raafat, F. (1991). Survey of literature on continuously deteriorating inventory models. *Journal of the Operational Research Society*, pages 27–37.

Rabinowitz, M. (1973). Blood bank inventory policies: a computer simulation. *Health services research*, 8(4):271.

Reesink, H., Davis, K., Wong, J., Schwartz, D., Mayr, W., Devine, D., Georgsen, J., Chiaroni, J., Ferrera, V., Roubinet, F., et al. (2013). The use of the electronic (computer) cross-match. *Vox sanguinis*, 104:350–364.

Şahin, G., Süral, H., and Meral, S. (2007). Locational analysis for regionalization of turkish red crescent blood services. *Computers & operations research*, 34(3):692–704.

Seifried, E., Klueter, H., Weidmann, C., Staudenmaier, T., Schrezenmeier, H., Henschler, R., Greinacher, A., and Mueller, M. (2011). How much blood is needed? *Vox sanguinis*, 100(1):10–21.

Services, U. B. (2011). Blood stocks management scheme.

Sönmezoglu, M., Kocak, N., Öncul, O., Özbayburtlu, S., Hepgul, Z., Kosan, E., Aksu, Y., and Bayik, M. (2005). Effects of a major earthquake on blood donor types and infectious diseases marker rates. *Transfusion Medicine*, 15(2):93–97.