SIMULATION OF A REDUCED RED BLOOD CELL SHELF LIFE

By

Matthew Hardy

Submitted in partial fulfilment of the requirements

for the degree of Master of Applied Science

at

Dalhousie University

Halifax, Nova Scotia

January 2015

Table of Contents

List of Tables

List of Figures

Abstract

Since the 1970’s, red blood cells (RBC) have had a rated shelf life of 42 days. Recently, questions have been raised regarding the safety of transfusing "older" red blood cells. Evidence suggests that mortality increases with blood use and believed that older RBC are less safe than younger ones. Shortening the shelf life of RBCs, however, may increase costs and lead to greater instances of outdates and shortages.

In this study, four simulation models are used to evaluate the impact of a reduced red blood cell shelf life on outdate, shortage, and emergency order rates. The first model is based on a network of two distribution centres, while the following three models are created using a reusable framework and are based on single distribution centres of varying sizes. Each model was evaluated under a 42, 28, 21, and 14 day maximum shelf life.

The results show that the size of both the overall distribution network, and the consumer sites impact the effects of a reduced shelf life. Small consumer sites and networks will experience shortage and outdate rates exceeding currently acceptable targets at even a 28 day shelf life, while the largest consumer sites and networks would be expected to meet current targets at a shelf life of 21 days or greater. At a 14 day shelf life almost all consumer sites and networks would experience outdate and shortage rates considered unacceptably high by current standards.

The effects of a reduced shelf life have a further impact based on blood type. The more universal types, such as O and A negative, would be expected to account for the greatest increases in both outdates and shortages under a reduced shelf life. The resulting requirements for suppliers to provide not only more blood to account for additional wastage, but to provide greater quantities of high demand blood types would create a significant challenge to the blood supply.

List of Abbreviations Used

CBS	Canadian Blood Services
FIFO	First In First Out
RBCs	Red Blood Cells
SMEs	Subject Matter Experts

Chapter 1 Introduction

Blood transfusions are required by individuals suffering from blood loss for a variety of reasons including blood loss due to surgery or accidental injury, as well as infections and illnesses causing anemia. Red Blood Cells (RBCs), the most commonly transfused blood component, perform the critical function of delivering oxygen to organs and tissues throughout the body.

Early transfusions required donors and patients to be side by side. However, by the end of World War I the use of anti-coagulants allowed for blood to be stored for several days prior to transfusion. By World War II, this was extended to 21 days, and by the 1970's the current standard of 42 for Red Blood Cells days was achievable (Blake, Hardy, Delage, & Myhal, 2013).

RBCs are collected either as whole blood which is then centrifuged to separate components including RBCs, plasma and blood platelets, or by apheresis where the RBCs are separated and the remaining components returned to the donors' body. Both the supply and demand for RBCs is stochastic and because of a lack of an artificial substitute for RBCs, suppliers must rely on voluntary donors to meet demand. The need for voluntary donors, combined with perishability, results in not only financial costs to collect, process, store, and distribute RBCs, but also donor goodwill costs associated with outdates if the blood supply is not well managed. Furthermore, aging demographics are expected to put pressure on the blood supply through higher demand and a smaller pool of donors (Greinacher et al. 2011).

1.1 Canadian Blood System

The blood supply in Canada is managed by two independent, not-for-profit organizations, Canadian Blood Services (CBS), and Héma-Québec. Héma-Québec manages the blood supply within the Province of Québec, while Canadian Blood Services provides services to the other provinces and the territories. Canadian Blood Services operates nine joint production/distribution centres and one distribution only centre. Héma-Québec operates two production/distribution centres within the Province of Québec. Both organizations manage the entire supply chain from collection, through production and testing, to distribution to healthcare facilities within their respective jurisdictions. While they operate independently, agreements exist for cooperation in times of shortage or crisis. The two organizations combined provide approximately 1.1 million units of RBCs to healthcare facilities across Canada annually, with Héma-Québec providing 240,000 units (Héma-Québec, 2014), and Canadian Blood Services 850,000 units (Canadian Blood Services, 2014). While both organizations directly serve different segments of the Canadian population, they share a common mandate to provide an adequate, safe, and, increasingly in today's economy, a cost efficient blood supply.

1.2 A Shorter RBC Shelf Life

Despite the current shelf life of 42 days, a body of literature exists that suggests poorer clinical outcome when older RBCs are used for transfusion. It is already common practice to transfuse younger RBCs to pediatric and some cardiovascular patients. Koch et al. (2008) through a retrospective study, found higher mortality rates and higher rates of complications for cardiac patients who were transfused with older RBC units. Frank et al. (2013) found that RBCs stored for longer than three weeks lose the flexibility required to pass through the capillaries in the body and thus lose their ability to deliver oxygen. Tinmouth et al. (2006) found that there is strong evidence that prolonged RBC storage may be harmful, but noted that there is a lack of clinical trials to confirm the impact. Further clinical trials are currently underway to evaluate the effects of transfusing young versus old RBCs in the general patient population (Steiner and Stowell, 2009).

A shorter shelf life is expected to impact the blood supply by reducing the availability of product, increasing the amount of wastage, reducing safety stocks and increasing order frequency. All would be expected to increase costs to a system already under pressure from provincial governments looking to reduce spending and to further strain a blood supply facing demographic shifts.

1.3 Problem Statement

The purpose of this research is to evaluate the impact of a reduced shelf life for Red Blood Cells on the Canadian blood supply. To achieve this, simulation models of the Québec blood system, as well as three regional centres of the Canadian Blood Services system were created and used to test the impact to wastage, service level, and ordering frequency under a 28, 21, and 14 day shelf life.

This thesis contains eight chapters. A literature review in Chapter Two details existing works which model the blood system and approaches specifically taken to evaluate a shorter shelf life. Chapters Three and Four describe the simulation model of the Québec blood system in detail along with the experiments and results. Chapters Five and Six describe the reusable model which was applied to the CBS blood system and detail the experiments conducted. Chapter Seven discusses the impact that a shorter shelf life would be expected to have.

Chapter 2 Literature Review[1]

Blood products are perishable and thus subject to natural wastage. Their management is complicated by the fact that, to avoid wastage and shortage, ordering decisions must be conditioned on the age of the stock on hand. If the shelf life of the product is one period, the perishable inventory problem reduces to the well-known news-vendor problem. When the shelf life is more than a single period, the problem becomes more complicated. The difficulty of solving the general *m*-period problem has led to the development of a rich body of literature describing approximate solutions (Telkin, Gurler, Berk, 2001).

Belien and Forcé (2012) review 97 papers relating to blood supply chain issues dating back to the 1960's and note peaks in publications in both the 1970's and 2000's. Much of this literature is specifically oriented towards red cell inventory and ordering. Pierskalla and Roach (1972) use a dynamic programming formulation to show that FIFO policies are optimal in perishable inventory problems, such as the red cell ordering problem. Fries (1975) also describes a dynamic programming approach to perishable inventory policies under the assumption of no backordering. Ordering policies in this case depend on the stock on hand, its age distribution, and the length of the planning horizon. Nahmias (1975) adopts a similar approach to Fries, but notes the extreme difficulty of computing optimal policies when *m* is greater than two days. In place of exact solutions, Nahmias (1978) argues for the use of heuristic solutions.

Literature on network planning for a blood supply chain is less well developed, perhaps as a result of the complexity of the network problem, which compounds an already difficult inventory management problem. Belien and Forcé (2012) note that while literature focusing on

---

[1] Sections of this literature review appear in Blake & Hardy (2014a) and Blake et al (2013)

the hospital and regional blood centre levels generally followed this trend in the literature, the overall supply chain was only briefly studied prior to 2000. Brodheim (1979) and his co-authors describe a number of studies to set inventory levels within a regional blood distribution network under the assumption of a 21-day shelf life for red cells. They note at that time (the mid-1970's) outdate rates for red cells in the Greater New York region were in the range of 20%. A statistical analysis of facilities within the New York region is used to derive a piece-wise linear relationship between target inventory level and mean daily demand for varying levels of availability (Brodheim, Hirsch, Prastacos, 1976). Inventory within hospitals is modelled as a Markov chain, under the assumption of daily deliveries of a fixed size *n* which allows the inventory state space to be represented in a compact form (Brodheim, Derman, Prastacos, 1975). Brodheim uses the model to set target inventories under the assumption of a fixed delivery schedule to investigate a dual use policy in which newer units are rotated between sites, while older units are retained.

Beliën and Forcé (2012) note that simulation is amongst the most common method for finding normative inventory policies for blood systems; its use has become more widespread as computational power has increased. They also suggest that it has seen an increase in popularity as work in the field has become more cross disciplinary, including health care managers and business economists with an interest in evaluating a number of scenarios rather than finding general policies. One of the earliest simulation studies appearing in the literature is due to Jennings (1973), who employs simulation to test red cell outdates and shortage rates in a customer network consisting of up to twenty facilities using a common inventory policy. Freidman, Abbott, and Williams (1982) describe the use of simulation to set inventory levels for red blood cells under the assumption of a 35-day shelf life. Hesse et al. (1997) describe an application of inventory management techniques to platelets in a system in which a centralized blood bank supplies 35 client hospitals. A periodic review model and (s, S) policies are developed for each of the client institutions using a simulation model as a test platform.

Sirelson & Brodheim (1991) use simulation to test platelet ordering policies for a blood bank, based on average demand and a fixed base stock level. They show that a base stock level derived from mean demand plus a multiple of the standard deviation of demand can be used to reduce current outdate and shortage rates. They also show that, on a regional level, low shortage and outdate rates can be readily obtained; within individual hospitals low outdate and shortage rates are more difficult to achieve.

Katsaliaki and Brailsford (2007) describe the use of a large scale simulation model to evaluate the function of a blood supply chain. A number of operational policies are tested via simulation. For a single producer and single consumer system, it is shown that the amount of inventory stored can be reduced if improved ordering and cross-matching policies are implemented. The model can be extended to cover a larger number of consumers, but the extended model must be executed via a distributed simulation environment (Brailsford, Katsiliaki, Mustafee, Taylor, 2006). Rytila and Spens (2006) describe the development and application of a simulation model to the blood system in Finland to test a series of scenarios for management of RBCs, plasma, and blood platelets. They provide an overview of the process required to develop, validate, and apply the simulation to the real world system and results from 11 policy scenarios.

Yegul (2007) describes the development of a custom model to evaluate inventory policies within a regional blood network. The model is applied to a region of Turkey with a two distribution centres and 49 hospital sites to identify policies that reduce shortages, outdates, and order frequency. Lang uses simulation based optimization to set inventory for a two-echelon system consisting of a single supplier and seven hospitals in which transshipment is allowed between hospitals as well as substitution of compatible blood types. The results show that the use of transshipments and substitution decrease shortage events but also result in higher usage of some rarer blood types (O-, A-, and B-). The study also found that implementing transshipment and substitution affected optimal order up to levels at hospitals differently by size, with smaller hospitals generally expected to increase holdings.

Blake et al. (2010) use simulation along with simple analytical rules to evaluate platelet inventory and ordering behavior for seven hospitals. Their model searches for the existence of ordering policies that jointly meet outdate and shortage targets. Results of the model show that better results are obtained when demand is larger and more regular. Their results suggest that larger tertiary centres have little difficulty managing inventory, but that smaller sites are necessarily subject to greater wastage rates. Fontaine et al. (2009) employ simulation of the platelet supply chain to recommend a series of practical policy improvements for a university-based blood supply chain. They note the importance of agility to respond to sudden changes in demand level and suggest a collection and testing regime that extends over seven days of the week and which is explicitly tied to expected demand.

There are four instances in the literature specifically addressing a potential overall shorter shelf life. Fontaine et al (2010) use a trace driven simulation to evaluate shorter shelf life in a single hospital / single supplier system. Shelf lives of 35, 28, 21, 14, and 7 days were simulated. Their model assumes a single hospital ordering daily from a single supplier according to an order up to policy. Shorter shelf life was simulated by reproducing historical orders and eliminating any items received after the maximum shelf life in the scenario. Demand was simulated as a stochastic process and both outdates and shortages were recorded. Results showed that shortages increased by 0.8% to 51% in the five scenarios over baseline, while outdates increase from 0.4% to 3.2% over baseline.

Atkinson et al. (2012) study the same single supplier/single hospital system as Fontaine et al., but assume stochastic arrival and demand for blood. They note that, in their baseline data, there are very few outdates and thus product arrivals are necessarily similar to demand. This results in queuing traffic intensities in excess of 0.99, suggesting a saturated queuing problem. To address this methodological difficulty, the authors evaluate shorter shelf life for RBC using a policy in which a varying number of units are imported from exogenous suppliers. They develop trade off curves in which the shelf life of RBC is compared against requirements for imported units under a variety of supply to demand ratios. The results of the study are used to illustrate the relationship between age at transfusion and import rate. Outdates are not reported, since they comprise an implicit input to the model, but the authors conclude that a regulated shelf life of 14 days would sufficiently reduce the age of RBC transfused to patients to be consistent with studies showing decreased mortality in cardiology patients when transfused with younger blood.

Simonetti et al. (2013) describe a simulation model of the US blood supply chain that is used to evaluate different distribution policies. Their model includes all suppliers and all distributors (hospitals) in the United States, but at an aggregate level only; all suppliers are incorporated into a single supplier object, while all distributors are aggregated into a single hospital. The supplier is assumed to collect blood and distribute it using a first-in-first-out (FIFO) policy. The study compares a hospital policy of distributing RBC with a 42 day shelf life using first-in-first-out (FIFO) policies against two non-FIFO policies: one of which skews to older units first, the other which skews to newer units first. Results suggest that the older unit first policy reduces the overall amount of inventory available in the US system by 6-8%, while a newest unit first policy reduces system inventory by 37%. Based on this result, Simonetti et al. conclude that a newer units first policy could lead to shortages, but could also sufficiently reduce the age of transfused RBCs to comply with a reduced shelf life.

Qinglin and Warren (2014) describe the application of a simulation optimization approach to a single distribution centre and a single consumer under a reduced shelf life for RBCs. The study applied a metaheuristic to minimize system wide outdates subject to a minimum service level under the assumptions of 21, 14, and 7 day shelf-lives. They report a 16% outdate rate at a 7 day shelf life, and outdate rates below 2% for 14 and 21 day shelf-lives. Qinglin and Warren note, however, that the search process carries a high computation cost, even in the case of a single centre and consumer.

Thus we conclude that while there is substantial literature on RBC inventory problems in general the existing literature, it is largely focused on single sites, such as a hospital or a distribution centre. Most of the literature addresses problems under either an increasing shelf life or the current 42 day shelf life and those studies that do address a shorter shelf life are limited to large US healthcare facilities or networks of aggregated demand. As a result, it is difficult to generalize the results of these studies, or indeed to apply them to a country such as Canada with networks of small, geographically distributed healthcare facilities.

In this study a simulation methodology is employed to test for the impact of shorter red cell shelf life within four broad, geographically distributed networks that includes multiple supply points and multiple demand points. Like Brodheim et al. (1979), the model does not attempt to suggest inventory ordering behaviour for either the producer or the consumer that minimizes cost. Our experience suggests that producers and consumers are interested in simple policies that perform well with respect to outdate and shortage targets rather than pure cost minimization. Like Fontaine et al. (2010), but unlike Brodheim, the simulation model is used to model inventory because of its flexibility to deal with a variety of potential shelf-lives. Furthermore, order rates from the supplier are used as a proxy for demand data, as was the case with Fontaine et al. The model differs from previous studies in the literature in terms of its scope and detail. It simulates the ordering behaviour of both suppliers and a broad base of consumers. Unlike Fontaine et al. it is assumed that each actor is an independent agent and places orders after evaluating the state of their inventory. Additionally, it is assumed that ordering policies and practices employed by agents in the model may change as red cell shelf life is reduced.

Chapter 3 Héma-Québec Methodology

A discrete event simulation of the Québec blood system was developed in Microsoft Visual Basic.Net with Microsoft Excel as a front end to manage simulation inputs and outputs. The modelled system includes the two regional distribution centres operated by Héma-Québec, located in Montreal and Québec City, as well as 97 consumer sites distributed throughout the province. The model includes processes covering the entire blood system, from the collection of RBCs by Héma-Québec, to distribution to consumer sites, and finally transfusion to the patient, or disposal in the case of outdating units. Since Héma-Québec imports and exports a minimal number of units to and from blood organizations outside of Québec the model is treated as a closed system where all units collected are either used within the network or discarded and where no additional product enters or leaves the system.

3.1 Héma-Québec Model Description

A two-level blood supply chain is employed by Héma-Québec to distribute blood products in the Province of Québec; the system consists of two supply centres, located in Montréal and Québec City, that provide blood and blood products directly to consumer sites located in cities and towns throughout the Province of Québec. There are no regional supplier depots, as is common in blood systems in other jurisdictions. However, there are a number of smaller consumers throughout the province that have informal arrangements with larger regional consumer sites to transship older product to the larger consumers to lower the overall system wide outdate rate. It should be noted that consumers in Québec, unlike the rest of Canada, are charged directly for the products that they order.

Blood Collections

The supplier collects blood on a daily basis based on a target collection volume that is determined several months in advance, but which is adjusted according to the current inventory state and in response to exogenous factors, such as the anticipated outbreak of pandemic flu. The amount collected each day is represented by a day-of-week-specific Poisson distribution with parameter μ_{Cj}. (Because of the large values of μ_{Cj}, a normal approximation to the Poisson distribution is used for computational efficiency). Of course, the blood type and Rh status of the donated unit is unknown until after the time of collection, thus in the model it is assumed that the blood agency "orders" a number of collections and is unable to directly influence the blood group and Rh status of the donated units. However, it is assumed that the distribution of blood types amongst donors is equivalent to that historically collected by Héma-Québec, which differs somewhat from the distribution of blood types amongst the general population. Thus, the number of units collected on a particular day is sampled from a normal approximation of the Poisson distributed daily value. Once the total collection value has been determined for a particular day, each unit is assigned a blood type (ABO and Rh status) using an empirical distribution.

The historical data from Héma-Québec indicates that inventory fluctuated significantly over the period 2008-2009. See Figure 1. At least three large increases and decreases in inventory over the period are evident, representing efforts by Héma-Québec to ramp up inventory ahead of expected peaks in pandemic flu or ahead of vacations periods when blood drives are not expected to perform well. This aggregate analysis suggests that inventory in the period covered by the historical data set (and perhaps all periods) is not truly in steady state; collections are clearly neither independent of inventory level nor exogenous factors, but rather are adjusted to either increase inventory in response to anticipated need or to deal with a surfeit of product. To represent the adjustments made by managers in response to fluctuations in inventory level and exogenous events, it is assumed a set of control bands for collections in the simulation. Accordingly, the blood agency is assumed to have a target inventory level equal to some number of days of average demand (i.e. seven days). When the total number of units on hand is within +/-1 days of the target level, normal collections are assumed and the daily collection is drawn directly from the normal approximation to the day-of-week-specific Poisson distribution. If on hand inventory exceeds target level by between 1 and 3 days, inventory is assumed to be reduced by decreasing collections by 10%. If on-hand inventory exceeds the target by more than 3 days of inventory, it is assumed that inventory is reduced by decreasing collections by 20%. In the event of an inventory shortage, it is assumed that collections are increased to return the inventory to the target level. In the model, if on-hand inventory is between 1 and 3 days below target, collections are assumed to be increased by 10%; if on-hand inventory is below 3 days of target, collections are assumed to increase by 20%. The percentage increases in collection are based on an analysis of the range of collection values recorded by Héma-Québec over the period from March 2008 through December 2009.

Figure 1: Héma-Quebec Aggregate Supplier Inventory Levels for March 2008- December 2009

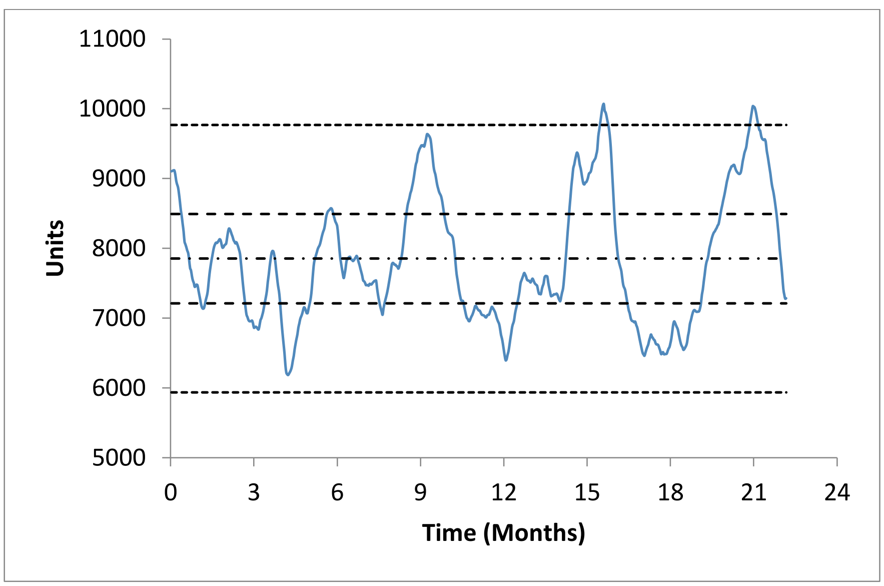

The bands indicate average demand, +/- 1 day, and +/- 3 days of average demand which are used in the inventory control policies.

Testing and Issuing

After collection, raw blood is tested for transmissible diseases and manufactured to produce a range of blood products. In general, production and testing take 1-2 days to complete. Thus, units that are collected on a Tuesday are typically produced and tested on Wednesday, end-labelled, and made available for distribution on Thursday morning. In this simulation model, it is assumed that testing takes one full day to complete and thus units are released into available inventory on the morning of the second day following collection. It is additionally assumed that there are no production, quality control, or testing losses; all units that are collected are assumed to be available for distribution at the end of the testing and production period.

Consumer Inventory Policies

All consumers are assumed to follow (R,S,s) inventory policies where *R* is the period between inventory checks, *S* is an order up to inventory quantity, and *s* is an inventory level order trigger point. The value of *R* for each consumer is assumed to be one day (i.e. inventory is checked at the beginning of each simulated day). Individual values are set for *S* and *s* for each consumer and each blood type. In the model, *S* is selected as average days demand multiplied by the number of days inventory desired by the consumer. The value of *s* is set as *S* less the average daily demand for the blood type multiplied by the average number of days observed between orders. When an order for any blood type is triggered (i.e. the current inventory level is below *s*) it is assumed that an order is issued for all blood types such that they are each returned inventory back to inventory level *S* even if the observed inventory level is greater than or equal to *s*.

Consumer Ordering

Consumer orders are assumed to be placed with the supplier at the beginning of each day before demand is observed. If an order is required, the consumer places the order with the supplier. Orders are assumed to be received by the supplier and are filled and then dispatched to the ordering consumer with no transport delay. In reality, consumer orders are a combination of standing orders for regular products, ad hoc orders to replenish inventory that has been depleted, and emergency orders to fill demand for which there is no available product on the shelf. Moreover, there is typically a delay to fill orders and, of course, a delay to transport units to the consumer site. However in practice it is usual for institutions to evaluate their inventory late in the day, after the bulk of demand has been experienced.   Orders are then typically transmitted to the supplier and are filled and dispatched overnight for early morning delivery, with the objective of having new product on hand prior to the start of peak demand. The assumption of instantaneous replenishment, therefore, while deviating from reality, is not entirely unreasonable. This assumption of instantaneous replenishment allows for simplification of the daily cycle within the model and accommodates the lack of data on when transfusions occur at the consumer sites.

Consumer orders are assumed to be processed as a batch, starting with the largest consumer, as measured by average daily demand, and proceeding to the smallest consumer, again as measured by average daily demand. Since inventory is issued by the supplier using a FIFO policy, this assumption results in slightly older blood being distributed to larger institutions with younger blood going to smaller ones. This mimics the common practice of ensuring that consumers with smaller demand receive newer product and thus have more available shelf life in which to turn inventory.

Demand is experienced by consumers throughout the day. In the model, total daily demand is assumed to be a random variable, drawn from a day-of-week specific Poisson distribution. Once the total daily demand is determined, each demand "item" is assigned a blood type and Rh status from an empirical distribution, based on the historical usage at each distribution centre.

It is assumed that demand is filled from inventory on hand using a FIFO ordering policy. If there are no units of a particular blood type on hand, it is assumed that consumers will substitute a compatible product where possible. Where no compatible unit exists to satisfy a demand, an emergency order is placed with the supplier to fill the demand. If a compatible unit is available, it is assumed to be dispatched (again instantaneously) and the demand item is counted as an emergency order. If, however, no compatible unit is available at the supplier, the demand item is counted as a shortage and is disposed. Due to the time sensitive nature of most transfusions, back ordering is not allowed and it is assumed that any shortages must be filled by emergency imports from outside the system. Demand items that are filled are counted and disposed.

Consumer Transshipments

While regulatory requirements prohibit product from being returned by consumers to Héma-Québec once issued, a number of consumers within the province have implemented transshipment agreements with other consumers. For the most part, this consists of an arrangement between a smaller consumer site and a larger consumer site to transship units when they reach a particular age. Within the Héma-Québec network, a number of such arrangements are known to exist, even though they are not part of the formal distribution network maintained by the supplier. In the simulation model it is assumed that a total of fifteen consumers transship approximately 2,377 units (1.0% of total system volume; 31.4% of the volume received at the transshipping consumer sites; 7.8% of the volume at the receiving consumer sites) per annum to four affiliated, regional consumer sites. Commonly, transshipment agreements allow for surplus units with 10 days or less of remaining shelf life to be transshipped. Within the model, it is assumed that transshipments occur at the end of the day when units reach the lesser of ten days of remaining shelf life or half of their original rated shelf life (i.e. if a unit has a rated shelf life of 14 days, it is assumed the unit will be transshipped with 7 days of remaining shelf life). As with order filling, transshipments are assumed to occur instantaneously and thus the units are available at the destination for the start of the next demand cycle.

Simulation Cycle

The simulation follows a repeating daily cycle of steps performed at both the supplier and consumers as shown in Figure 2. At the start of each day, stock at both the supplier and the consumers is aged by one day. Any stock that has zero remaining days left until expiry is outdated and counted. Similarly, any stock with one day of testing remaining becomes available inventory at the supplier.

After aging units, the supplier and consumers both observe their existing inventories. The supplier places an order for the total number of units to be collected. The supplier order is based on a normal approximation to a Poisson distribution with a mean value that is specific to the day of week. The mean value, however, may be adjusted for the supplier's inventory level relative to a target inventory level, as observed two days in the past. As noted previously, the two day horizon was identified through consultation with Héma-Québec subject matter experts and was based upon the minimum time to make adjustments to capacity at community blood clinics. The supplier then observes the day's collections, in terms of actual number of units and their blood type, and adds any newly collected units to the testing queue.

The consumer, conversely, after observing inventory level for each blood type and comparing these to target inventory levels, determines if an order is to be placed. If any blood group requires replenishment, the consumer issues an order to bring all blood groups up to their target levels. All consumers issue orders for products at the same time. Once issued, the supplier evaluates and fills consumer orders (if possible) from available inventory using a modified FIFO policy, where inventory is only issued with ten or more days of remaining shelf life, except in the event of a shortage. Consumer orders are aggregated by blood group and compared to available inventory. If available inventory is sufficient to meet all consumer requests, the supplier fills orders in decreasing order of consumer size as measured by average daily demand. If inventory is insufficient to meet all consumer requests, each individual order is prorated by (units available / units requested) before being filled by the supplier, again in decreasing order of average daily demand. Available units are then released from inventory and dispatched to consumers.

Consumers receive their inventory from the supplier instantaneously upon issue. Demand for the day is then experienced. Daily demand at consumer sites is assumed to be Poisson distributed with a mean value that is specific to the day of the week. Once total demand is determined, each demand item is assigned a blood group. The consumer then attempts to satisfy demand exactly. If an exact match cannot be found for a particular demand, the consumer will fill demand using a compatible product. If no compatible unit can be found to meet a particular demand locally, the consumer will issue an emergency order to the supplier. Emergency orders are, like regular orders, assumed to arrive at the supplier, be filled, dispatched and returned to the consumer instantaneously, if an exact or compatible unit is available. If no suitable unit can be found at the supplier, the demand is assumed to be lost.

Finally, all consumers with a transshipment agreement in place evaluate their stock. Any unit which has the lesser of 10 days of remaining shelf life or half of the original rated shelf life is removed from the originating consumer's inventory and entered into the destination consumer's inventory. The day then ends and the cycle repeats.

Figure 2: Simulation Cycle Diagram

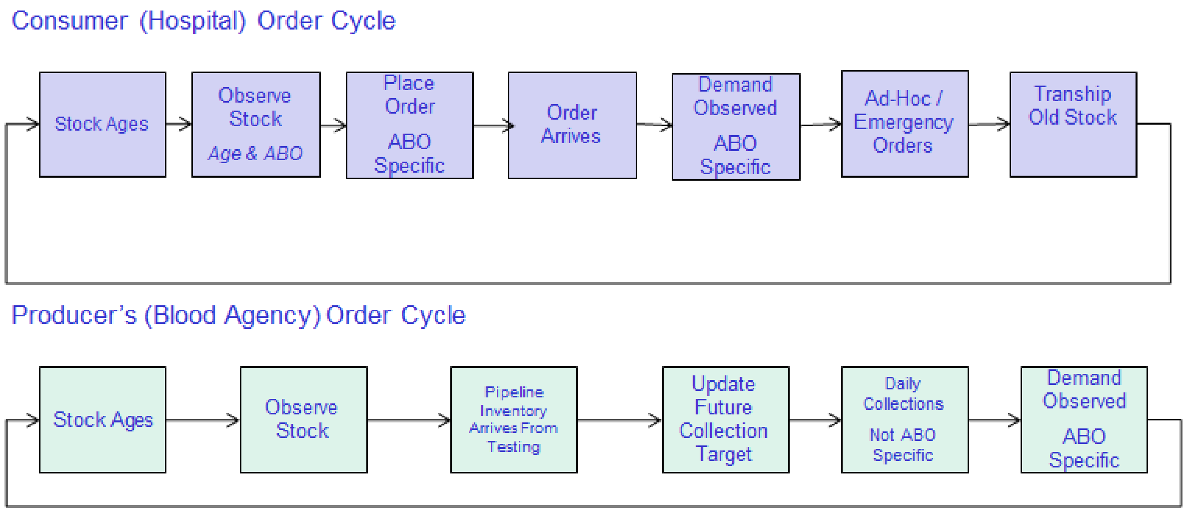

The daily process cycles for both the supplier and consumers. The observed demand and the supplier consist of the sum of consumer orders.

3.2 Héma-Québec Data

Data for the study was provided by Héma-Québec from their operational data collection systems. The data consists primarily of transaction level records for all RBC units collected and shipped to consumers or outdated by Héma-Québec during the period January 2008-July 2010. In total, there are 597,982 records in the data set. Some paring of the data set was necessary, due to incomplete records for 2010. Accordingly, transaction records including blood collections and distribution to consumers, for the years 2008 and 2009 were used as the primary data source for this project. There are a total of 487,046 records for this time period covering all 185 customers listed in the Héma-Québec database. However a total of 88 customers were excluded from the analysis. Six consumers were eliminated because their total annual volume of RBC received was less than 1.5 units per week. In addition, 39 customers were eliminated because they are not hospitals, but blood agencies outside of Québec. Finally, 43 consumers were eliminated for which no data on outdates was available and therefore actual demand could not be estimated. The remaining dataset includes 97 consumers with a total of 470,482 transactions over the 2008-2009 period.

The average annual demand for RBCs, as determined from the transaction level data in the reduced dataset, is approximately 225,000 units or roughly 620 units per day. The ABO distribution of units shipped by Héma-Québec roughly resembles that of the Canadian population by ABO type; However there is a 60% higher demand for Rh negative blood in the units shipped by Héma-Québec than would be expected in the general population, due to the effects of type matching at transfusion. Demand, as measured by shipments from Héma-Québec distribution centres, is skewed towards weekdays with weekend demand (Saturday and Sunday combined) being about one fourth that of an average weekday (857 units Monday through Friday versus 230 units on Saturday and Sunday combined). Monthly demand is relatively consistent, averaging between 18,500 and 22,000 units ($\bar{x} = 20{,}032;\ \sigma = 958$).

Both consumer demand and supplier blood collections were tested for seasonality using Analysis of Variance (ANOVA). The number of units collected was separated by day of week and a one-way ANOVA was run to test the hypothesis of equality of means for collections by day of week. The analysis showed a strong day of week effect ($p < 0.01$). As with the collections, aggregate demand was separated by day of week and a one-way ANOVA was run to test the hypothesis of equality of means for demand by day of week. This analysis also indicated a strong day of week effect for demand ($p < 0.01$). Tests were similarly conducted to evaluate month-to-month seasonality for both collections and aggregate demand. The resulting p-values for these tests were 0.582, and 0.381, respectively, suggesting that there is no evidence to support the hypothesis of monthly trend in seasonality for either collections or aggregate demand. Finally, an ANOVA was repeated for demand at the five largest consumers. Again there was found to be a strong day of week effect but no evidence to support a hypothesis of month-to-month variation in the data set.

In addition to transaction level data for individual red cells, Héma-Québec also provided the annual outdate rates for some of the consumers for the period of 2004-2008. (Note: consumers for which there were no outdate records were excluded from the analysis as described above.) The summary data consisted of the total units shipped to each consumer, the total units outdated by the consumer, and the outdate rate for the year. The rates provided for the year 2008 were used to estimate demand and to validate the simulation model because they were the most recent available and also overlapped with the available shipment data. Because consumer demand was not directly available, it was necessary to estimate demand from shipping data. Since units shipped to a particular consumer site include both RBCs that are transfused as well as those that eventually outdate, it was necessary to adjust the annual demand rate at each consumer site by prorating for the number of units shipped to each consumer by (1-the outdate rate) as follows:

$$D_i^{Annual} = H_i^{Annual}\left(1 - O_i^{Annual}\right)$$

Where:

D_i^{Annual} is the annual demand at consumer i
H_i^{Annual} is the annual shipment from Héma-Québec to consumer i
O_i^{Annual} is the 2008 outdate rate (in %) at consumer i

Consumers that reported outdate rates less than 0.5% of total volume received per year were assumed to transship older units to larger institutions. Consumers that transship product are expected to send all older RBC units (those with the lesser of 10 days of remaining shelf life of ½ of their maximum rated shelf life) to the largest affiliated consumer within in their regional network. Demand at transshipping consumer sites was estimated in the same way as demand at consumer sites that are not assumed to transship: demand less reported outdates. Of course, since these consumer sites are assumed to transship, the demand estimate is only approximate and likely over-estimates the actual demand at these consumer sites, since a portion of products shipped are neither transfused nor outdated, but rather shipped to a larger consumer. However, the magnitude of error associated with this simplification is modest; the total volume of products transshipped, based on model estimates is less than 1.0% of total system volume.

Data on target inventory levels were available for a number of larger consumers. In total, inventory targets were available for 21 institutions, comprising 53.8 % of the demand in the reduced dataset. The consumers in this data set were grouped according to the distribution centre (Montréal or Québec City) that supplies the majority of their product. This data suggested that consumers supplied from Montréal have average target inventories of approximately 6.83 days demand, while consumers supplied from Québec City hold about 7.74 days of average demand in inventory. Target inventories for consumers which were not know *a priori*, were estimated by extrapolating from their daily demand under the assumption that they would hold the same target days inventory as the larger consumers supplied by the same distribution centre, for which target inventory was available:

$$\lceil Inv_{i,j} \rceil = D_i^{Daily} * P_{i,j} * n_i$$

Where:

$\lceil Inv_{i,j} \rceil$ is the quantity of blood type j to hold at consumer I, rounded up to the nearest integer value

D_i^{Daily} is the total average daily demand for all blood types at consumer i

$P_{i,j}$ is the proportion of total demand of blood type j at consumer i

n_i is the average number of days demand held in inventory held by consumers supplied from the same distribution centre as consumer i

3.3 Héma-Québec Validation

Validation of the simulation model took part in stages, using an iterative process as suggested by Law (2006). Subject matter experts (SME) were consulted regularly, over the course of a year, to comment on model assumptions and preliminary designs. A modular design process was used to build the actual simulation and each module was tested and verified before going into production.

Verification efforts focused on ensuring that both collections and demand matched the historical records. Standard statistical tests were used to compare the simulated output against values observed in the historical record for both collections and orders. The mean daily collections and mean daily demand were determined for both the historical record and the simulation model and compared at a confidence level of 95%. In both instances, there was no evidence to suggest that the simulation output differed from the input data derived from the historical records. We do note that while the means are accurately replicated the simulation has lower variability than the real world system. This was due to both the longer run times of the simulation to achieve a steady state relative to the historical data sets, as well as the absence of exogenous events in the simulation.

Table 1: Verification of Collections and Demand Rates

	Simulation		Historical Data*	
	Mean	**Std Dev**	**Mean**	**Std Dev**
Daily Collections	642.30	3.67	642.33	35.46
Daily Demand	624.40	0.75	624.40	

* Because daily data is not available on daily demand data, we cannot calculate the variance of demand at consumers. Expected demand can be calculated from aggregate data, however.

Validation efforts focused on ensuring that inventory levels and outdate rates observed in the simulation matched the historical record. This was not straightforward. Validation typically involves a statistical comparison of historical data to simulated data. For such comparisons to be meaningful, an assumption of independent and identically distributed values is required for both the simulation and the historical dataset. The Héma-Québec data, however, shows clearly that inventory levels and, by extension, collections and outdates, are influenced by exogenous events. Three distinct instances were observed in the 2008-2009 data in which inventory values increased substantially and were then allowed to decline. Not surprisingly, validating the simulation model, which assumes a steady state, against the real-world system that responds dynamically to exogenous events, proved challenging.

For all validation runs, metrics from the historical data were divided into 22 periods of one month each (March 2008 to December 2009). The simulation was then run for 30 replications of 770 days and each replication was likewise divided into 22 periods for measurement purposes. This allowed for a statistical comparison of the two sets of metrics with each having a sample size of 22.

Initial validation efforts assumed collections were independent of inventory levels and that inventory would reach steady state. Under this assumption, results from the model indicated collections and demand could be validated against the historical data, but inventory levels and outdate rates could not. Furthermore, the length of time for the simulation to reach a steady state inventory level was exceedingly long, representing several years of real time. See Figure 3. The model did produce an outdate rate similar to that observed by Héma-Québec (15.7 per day in the simulation versus 16.5 in the real system); however, outdates in the simulation occurred almost exclusively at consumer sites. See Tables 2 and 3. Based on this analysis it was concluded that collections are not independent of inventory state and that managers use control policies to limit inventory levels and to respond to exogenous events. This assumption was validated by the SMEs.

Figure 3: Simulated Supplier Aggregate Inventory with No Inventory Control Policy

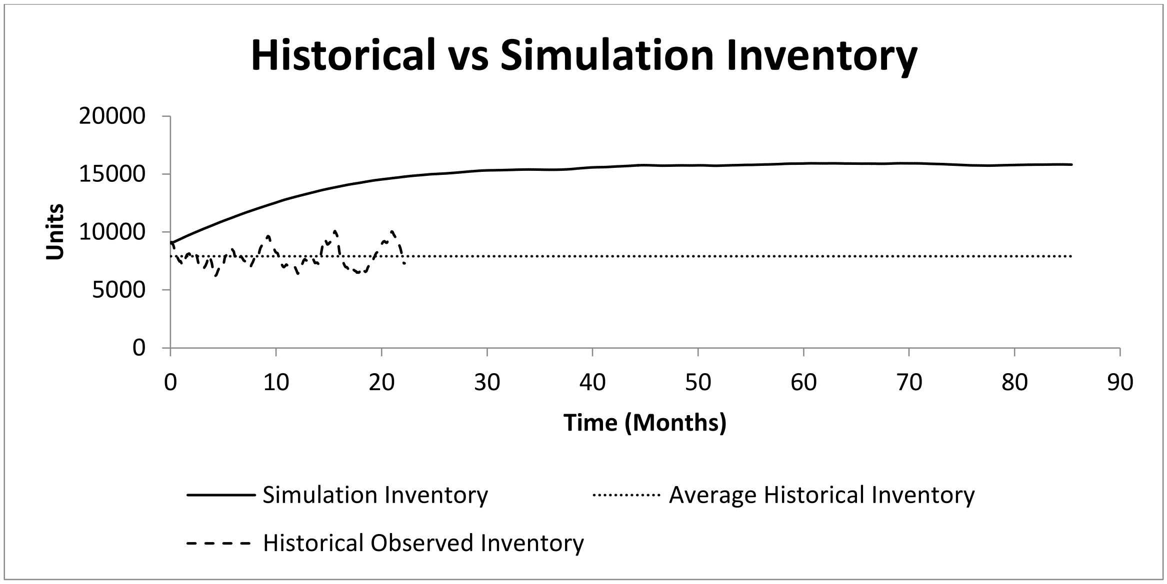

The historical inventory state for the Héma-Quebec distribution centres in aggregate and the simulation inventory state without an inventory control policy. The steady state inventory is significantly higher than in the real world system.

Table 2: Model Metrics with No Inventory Control Policy

	Simulation		Historical Data	
	Mean	**Std Dev**	**Mean**	**Std Dev**
Collections	642.30	3.67	642.33	35.46
Supplier Outdates	0.02	0.20	8.08	3.13
Inventory	14644.96	413.70	7895.83	911.32

Table 3: Model Metrics with No Inventory Control Policy

	Simulation		Historical Data	Prediction Interval
	Mean	**Std Dev**	**Mean**	
Demand	624.40	0.75	624.40	(622.99, 625.90)
Consumer Outdates	15.50	11.00	8.50	(-6.54,37.54)

Two control policies were then evaluated in the process of validating the simulation model against historical records. The first policy assumes that collections are increased during periods of shortage and decreased during periods of surplus. In this scenario it was assumed that if inventory is within +/- 1 day of target, managers will not adjust collections. If inventory levels are greater than +/- 1 day and less than +/- 3 days of target, collections are assumed to be decreased or increased by 10%. If inventory levels exceed +/- 3 days of target, collections are assumed to decrease or increase by 20%. Implicitly, this transforms collections from a pure stochastic process to a feedback control system in which managers monitor inventory and adjust collections according to an error signal created by the difference between target and observed inventory. Because feedback control loops are not instantaneous, a two-day lag between managers' ability to observe the state of the system and their ability to affect change in the collections process is assumed. The choice of a two day lag is based on discussions with the SMEs. Thus, in the simulation, managers are assumed to "observe" the state of the system on a Monday and implement changes (i.e. allowing collections to increase or decrease) on the following Wednesday. The policy is assumed to be applied to the mean value of the collections distribution before a sample is drawn on a particular day; it assumed that the variance of the collections distributions are unaffected by the implementation of the control policies. To test this policy the simulation model was run for a total of 30 replications of 770 days and the results were compared to the historical Héma-Québec data. It was observed that the simulation estimate of consumer demand and centre inventory levels are statistically indistinguishable from the 2008-2009 actuals, but that collections (630 per day vs. 638 per day) and thus, by extension, outdates (7.1 per day versus 16.5) are lower in the simulation than in the historical data. Furthermore, it was observed that almost all of the outdates in the simulation continue to occur at consumer sites.

Table 4: Model Metrics with Controlled Collection Rates

	Simulation		Historical Data			
	Mean	Std Dev	Mean	Std Dev	N	p-value
Collections	629.22	2.47	642.33	35.46	22	0.098
Supplier Outdates	0.13	0.15	8.08	3.13	22	0.0
Inventory	7858.30	34.50	7895.83	911.32	22	0.849

Table 5: Model Metrics with Controlled Collection Rates

	Simulation		Historical Data	Prediction Interval
	Mean	Std Dev	Mean	
Demand	624.40	0.75	624.40	(622.90,625.90)
Consumer Outdates	4.85	1.04	8.50	(2.77,6.93)

As Tables 4 and 5 show, outdates are nearly zero at the supplier in the simulation model. This result suggests, not an error in the function of the simulation, per se, but rather that a pure FIFO issuing policy is not invariably observed in practice. There are a number of good reasons why FIFO might be violated – for instance rare blood types or phenotyped units may be set aside for contingency. In addition, in instances of over collection, the system operator and customers might prefer that newer stock be issued. An analysis of periods in 2008-2009 in which there were substantial outdates in the historical data shows that FIFO violations do tend to precede instances of large outdates. To illustrate this effect, five weeks with large numbers of outdates were extracted from the database. The date of outdate was noted and the issues for the preceding seven days were extracted from the database. The age of the issued product was compared to the age of the products that outdated, and the numbers of newer products issued ahead of older products (i.e. in violation of FIFO) were noted. These results are shown in Table 6 and indicate that FIFO is not necessarily followed when there is a surfeit of stock.

Table 6: Observed Historical FIFO Violations

Outdate Date	Units outdated	Number of shipped units in violation of FIFO
26/03/2008	38	3901
18/07/2008	34	4175
24/12/2008	48	4849
16/07/2009	62	5292
23/07/2009	48	4524
Totals	**230**	**22741**

Based on this analysis, a second control policy was evaluated for validation purposes. Reducing collections when overstocked discourages donors and thus potentially reduces the base for future collections. Thus, an alternative policy might maintain collections, but later allow newer units to be issued ahead of older units in violation of FIFO and implicitly allow some units to outdate. This policy was implemented in the simulation model by tagging a fraction of units collected when above target levels as surplus. The model assumes that when inventory exceeds target levels by more than 1 day, but less than 3 days, that 1% (the approximate outdate rate at the supplier) of collected units are tagged as surplus. If inventory is more than 3 days above target, 5% of collected units are tagged as surplus. Surplus units are assumed to be set aside and not issued to consumers unless there is a subsequent shortage of stock. As before, it is assumed that if inventory is below target levels that mangers affect change to increase collections as outlined above. To test this second policy the simulation was again run for 30 replications of 770 days and the results compared to historical records. These results show that the model reproduces historical inventory levels, outdates, collections, and demand rates at the centres. See Tables 7 and 8.

Table 7: Model Metrics under a 'Collect to Outdate' Control Policy

	Simulation		Historical Data		n	p-value
	Mean	Std Dev	Mean	Std Dev		
Collections	642.30	3.67	642.46	35.46	22	0.983
Centre Outdates	10.39	3.25	8.08	3.13	22	0.021
Inventory	7925.01	96.54	7895.83	911.32	22	0.883

Table 8: Model Metrics under a 'Collect to Outdate' Control Policy

	Simulation		Historical Data	Prediction Interval
	Mean	Std Dev	Mean	
Demand	624.40	0.75	624.40	(622.99, 625.90)
Consumer Outdates	6.80	0.54	8.50	(5.72,7.88)

Under this assumption, the model, while faithfully replicating performance indicators for the supplier, is not fully able to replicate outdating behaviour at consumer sites. The model reproduces consumer demand but under-estimates orders and outdates by approximately 2 units per day.

By expanding the idea of allowing controlled outdates in violation of FIFO to sometimes allow newer units to be used to the consumer sites the model was able to achieve both the expected outdates and orders at the consumer sites along with the expected outdate levels at the supplier (see Tables 9 and 10). This was implemented in the model by randomly tagging a portion of units at consumer sites to be passed over for younger units, effectively allowing them to outdate in violation of FIFO, similar to at the supplier. This approximation attempts to capture FIFO violations which occur for a number of reasons including;

- high inventory levels at the supplier,
- specific requests from consumers for fresh blood,
- the supplier holding specific units, such as rare phenotypes, in reserve.

While FIFO violations are known to occur, the circumstances are not captured in available data and thus the specific policies are difficult to model. In total, 19 consumers with lower than expected outdates in the simulation were selected for this process. Each of these 19 consumers was given a daily probability, equal to the difference between their simulated outdates per day and their historical reported outdates per day, of tagging a unit which would then be allowed to outdate in violation of FIFO. The probabilities of tagging a unit to outdate in violation of FIFO ranged from 2% to 33% across the 19 consumers and allowed for 0.2% of total system volume to outdate annually. While there is no direct data available from consumer sites to confirm this assumption, it is it is not entirely unreasonable to assume that violations of FIFO do occur in practice.

Table 9: Model Metrics Under a 'Collect to Outdate' Control Policy with Controlled Consumer Outdates

	Simulation		**Historical Data**		**N**	**p-value**
	Mean	**Std Dev**	**Mean**	**Std Dev**		
Collections	642.30	3.67	642.46	35.46	22	0.983
Centre Outdates	8.24	3.26	8.08	3.13	22	0.869
Inventory	7926.95	115.23	7895.83	911.32	22	0.875

Table 10: Model Metrics Under a 'Collect to Outdate' Control Policy with Controlled Consumer Outdates

	Simulation		**Historical Data**	**Prediction Interval**
	Mean	**Std Dev**	**Mean**	
Demand	624.40	0.75	624.40	(622.99, 625.90)
Consumer Outdates	8.62	0.50	8.50	(7.62,9.62)

The validation process demonstrated two important aspects of modelling the blood supply. Firstly, the blood supply is not well modelled as a pure queue, but requires a control policy to manage inventory levels. As the results of the initial validation show, without a control policy, the inventory levels will either grow well beyond those of the real world system, or will deplete resulting in high levels of shortages. In practice, blood suppliers will use a variety of tools to raise public awareness and encourage donations in the event of low inventories, as well as to recruit and maintain new donors, in order to maintain the blood supply (Canadian Blood Services, 2013). The second aspect of modelling the blood supply revealed through validation are the occurrences of FIFO violations in distribution. Allowing FIFO violations allows the supplier to reduce inventory in the event of surplus product, while maintaining a stable average age of the stock distributed to the consumer sites. Incorporating both of these aspects of the real world blood supply was critical to developing a working model which accurately captured the real world system.

Chapter 4 Héma-Québec Experiments

The model was initially run under the assumption of a 42 day shelf life for RBCs to create a baseline scenario. The baseline scenario was run for 100 replications of one year under the method of replication/deletion. The selection of 100 replications was made to provide an absolute error of no more than 1% on the key metrics of outdates, shortages, and emergency orders. Warm-up for the model is a two-stage process. For the first 18 days of the run, collections occur, but consumers do not place orders. This allows the total inventory at the supplier to rise to the target level of approximately 7800 units. After building an initial inventory at the supplier, consumers start to order and the simulation is run for an additional 42 days (60 days in total) before statistics are collected. The sufficiency of the 42 day warm up period was confirmed using the Welch technique outlined in Law (2006) with the age of outgoing stock as the metric of interest. For the baseline scenario, it is assumed that collections are scaled in the event of a shortage or surplus, rather than being forced to outdate when stock levels are high. The baseline model further assumes a modified FIFO policy in which units are not issued from the centre with 10 or fewer days of remaining shelf life, unless there are no other units available. For remaining stock, FIFO policies are assumed to be strictly followed and units are not selected to be allowed to outdate as in validation. As noted above, there is some deviation between the baseline scenario and the historical record, which arises because of the real world system responds to exogenous events not included in the simulation (i.e. pandemic flu). Accordingly, the baseline scenario is used as the benchmark for evaluating shorter shelf policies, rather than history, since the model based comparison allows for analysis using a common methodology and common set of assumptions.

Three shorter shelf life scenarios were created corresponding to 28, 21, and 14 day shelf life. Each of these scenarios was run for 100 replications of one year with up to a 60 day warm up.

As the amount of inventory held in the system is reduced, the 18 day supplier warmup period is also reduced as less initial inventory is needed to bring the simulation to steady state and thus the total warmup decreases. The 42 day second stage of the warmup was again used and was confirmed using the Welch technique.

The primary purpose of the experiments was to determine how shorter shelf life affected RBC availability and outdate rates across the entire network. However, the model was also used to investigate policy options related to how much total inventory should be held in the system and where that inventory should be held (i.e. at the supplier or the producer) for each tested shelf life. To test total system inventory at each rated shelf life, the ordering behaviour of the supplier and consumers was adjusted such that the total inventory held in the system, as measured by days' supply, equaled 40%, 50%, or 60% of the assumed RBC shelf life. Thus, in the scenario in which RBCs are assumed to have a 28 day shelf life, total system inventory level was set such that 12 days (40%), 14 days (50%) or 16 days (60%) total supply was held. The inventory was split evenly with half held at the supplier and half held at the consumer sites. The results of the reduced shelf life scenarios were compared against the baseline scenario at a 42 day shelf life, as well as against the currently accepted target thresholds of 5% or less outdates and a minimum 99% service level.

4.1 Héma-Québec Results

Baseline

In the baseline scenario an overall outdate rate of 0.34% was observed across the entire network. Of the observed outdates, over 99% occurred at the consumer sites. A shortage rate of 0.17% was observed with all shortages occurring in Rh negative units. Consumer sites placed emergency orders (orders placed by a consumer outside of the regular order cycle in response to a demand that could not be filled by stock on the shelf) at a rate equal to 0.34% of demand, which they were unable to fill with on-hand inventory.

28 Day Shelf Life

The results for the 28 day shelf life, in terms of outdate rate, shortage rate, and emergency order rate are presented in Table 11. For the case of a 28 day shelf life, it can be seen that policies which simultaneously result in low outdate rates, low shortage rates, and low emergency order rates can be easily found. Outdate rates, which ranged from 0.63% to 0.68%, while representing statistically significant increases, are still modest and in fact below currently observed rates. Holding 12 or 14 days' worth of inventory increases both shortages (0.44% and 0.27%) and emergency orders (0.8% and 0.48%). At 16 days of inventory, a slight decrease in emergency orders to 0.28% is observed, while shortages remain virtually unchanged at 0.16%. It is worth noting that holding 16 days of inventory would be a slight increase from current levels.

It was also observed that the effects of reduced shelf life do not affect all consumers equally. As shown in Table 12, larger consumers, as measured by transfusion volume, see almost no effect in terms of outdates (0.36%) or emergency orders (0.03%) if the shelf life is reduced to 28 days. Smaller consumers however, are disproportionately affected and thus will see increases in outdates and emergency orders.

Table 11: Outdate, Shortage, and Emergency Ordering Rates for 28 Day Shelf Life

Inventory Held	Outdate Rate	Shortage Rate	Emergency Order Rate
16 Days	0.68%	0.16%	0.28%
14 Days	065%	0.27%	0.48%
12 Days	0.63%	0.44%	0.8%

Table 12: Average Consumer Outdate and Emergency Order Increases by Consumer Size for a 28 Day Shelf Life Holding 14 days Inventory Split 50%/50%

Consumer by Size	Outdate Rate	Emergency Order Rate
Largest 1/3	0.36%	0.03%
Middle 1/3	1.45%	0.81%
Smallest 1/3	2.96%	5.97%

21 Day Shelf Life

When shelf life is reduced to 21 days it is still possible to find policies that result in modest increases in outdate rates, shortage rates, and emergency ordering rates over the baseline. However the trade-off between shortage and outdate rates becomes more apparent (see Table 13). When 8 days of inventory are held in the system, an outdate rate of 1.07% was observed, along with a shortage rate of 1.2% and emergency orders equal to 2.18% of demand. Conversely, holding 12 days of inventory yields lower increases in shortage and emergency order rates (0.45% and 0.81% respectively), but higher outdates (1.3%). Once again, as shown in Table 14, it may be observed that the effects of a reduced shelf life are not equally distributed amongst all consumers; smaller consumers, as measured by transfusion volume, can be expected to see the larger increases in outdates and emergency ordering than larger consumers.

Table 13: Outdate, Shortage, and Emergency Ordering Rates for 21 Day Shelf Life

Inventory Held	Outdate Rate	Shortage Rate	Emergency Order Rate
12 Days	1.3%	0.45%	0.81%
10 Days	1.14%	0.72%	1.33%
8 Days	1.07%	1.2%	2.18%

Table 14: Average Consumer Outdate and Emergency Order Increases by Consumer Size for a 21 Day Shelf Life Holding 10 Days Inventory Split 50%/50%

Consumer by Size	Outdate Rate	Emergency Order Rate
Largest 1/3	0.62%	0.26%
Middle 1/3	2.53%	2.84%
Smallest 1/3	5.53%	11.82%

14 Day Shelf Life

When the rated shelf life is lowered to 14 days, the simulation suggests that managing the blood supply chain becomes more difficult. From Table 15 it can be seen that none of the tested policies were able to maintain shortage rates below the current target of 1%. As expected, holding 6 days of inventory in the system produces the smallest increases in outdate rates. This, however, is accompanied by large increases in both supplier shortages and consumer emergency ordering (2.44% and 4.21% respectively). Holding 8 days of inventory results in increased outdates over 6 days; however, shortage and emergency ordering rates drop sharply to approximately half the levels seen at 6 days of inventory held. As Table 12 shows, the effects of a 14-day shelf life are experienced by all consumers. However, the largest effects are again seen at smaller consumers, where increases in outdate rates are expected to exceed 13% on average.

Table 15: Outdate, Shortage, and Emergency Ordering Rates for 14 Day Shelf Life

Inventory Held	Outdate Rate	Shortage Rate	Emergency Order Rate
8 Days	3.07%	1.18%	2.17%
7 Days	2.37%	1.46%	2.36%
6 Days	2.63%	2.44%	4.21%

Table 16: Average Consumer Outdate and Emergency Order Increases by Consumer Size for a 14 Day Shelf Life Holding 5 Days Inventory Split 50%/50%

Consumer by Size	Outdate Rate	Emergency Order Rate
Largest 1/3	2.07%	1.99%
Middle 1/3	6.28%	8.46%
Smallest 1/3	13.83%	16.12%

Summary

Inventory policies which meet current targets of 5% or less outdates and 1% or less shortages are easy to find at a shelf life of 28 or 21 days. Once the shelf life drops to 14 days however shortages present a challenge and reaching a 99% service level may require holding a large percentage of the shelf life worth of demand in inventory. This would put pressure on more universal blood types (O and Rh negative in particular) which are typically already over collected relative to their ratio of occurrence in the Canadian population. The large impact on the smallest consumers may also worsen the outdate rates as administrators and clinicians would likely attempt to increase safety stocks at these locations which tend to be located further from distribution centres than larger consumers to mitigate shortages.

Chapter 5 Reusable Model

Following the Héma-Québec model, a reusable simulation framework, based on the Canadian Blood Services network, was developed. A reusable simulation framework was created which could represent any of the 10 regional distribution centres in the Canadian Blood Services network and instances of the model were created to represent each distribution site. To confirm the results of the Héma-Québec study, this regional distribution centre model was extended to evaluate the impact of a shorter shelf life for red blood cells within the CBS network. To generalize the results, the experiments were conducted at three regional distribution sites in the Canadian Blood Services network representing a small, medium, and large volume distribution centre and their associated networks of consumer sites.

As with the Québec model, the r0eusable version was built in Microsoft Visual Basic.net, but with a framework including two databases to manage inputs and outputs. The first database stores historical, transaction level data detailing the progress of blood from collection through disposition at the distribution centre and the second contains parameters related to execution of the simulation model and records output results. Unlike Héma-Québec, Canadian Blood Services distribution centres each service a unique geographic region. These regions form catchment areas for blood collections, production, and distribution; consumers within the area receive RBCs exclusively from the single distribution centre responsible for the region. Canadian Blood Services routinely moves product between regional distribution centres. As such, while each region is modeled individually, they are not closed systems since additional product enter and exit the system via imports and exports with other distribution centres.

The ability to develop a model of an additional CBS distribution centre by adding an additional transaction database allows the model to be quickly reused by CBS research staff.

The models themselves are useful not only for additional experiments related to a shorter RBC shelf life, but other research into inventory and logistics policies for CBS. While the initial development time for the reusable framework and the initial model was lengthy (approximately 210 days), subsequent models were able to be created and validated much more quickly due to the development of a standardized process for preparing the raw data and performing validation. As additional models were created with the reusable framework development time was reduced to between 7 and 14 days per additional distribution network modelled. In addition to this project, the reusable framework has also been applied to an analysis of current CBS Inventory policies for RBCs at the current 42 day shelf life (Blake & Hardy 2014a).

Transaction Database

Data for the modelling framework was derived from transaction level records extracted from Canadian Blood Services' production database. Information regarding all RBC units collected, distributed, or disposed of at any site in Canada between the periods of 18 Feb 2011, and 12 May 2012, (i.e. fiscal 2011-2012 with a 42 day buffer at the beginning and the end of the period) were obtained. In total, the data included slightly more than 1.3 million records. Each record tracks the journey of one unit through the supplier's portion of the supply chain: collection date, release date, expiry date, disposition date, and disposition status. In addition, since approximately 10% of all units collected in Canada are transferred between centres prior to distribution, a list of unit transfers for the period 18 Feb 2011 and 12 May 2012 was also obtained.

To prepare the data for use in the simulation models, a series of queries were written in MS-Access. The queries performed three functions – separating the data into region specific databases, formatting the data for use in the simulation models and building input distributions, and summarizing the data for verification and validation purposes. See Figure 4 for a summary flow chart of the steps required to prepare data for use in the framework.

Figure 4: Schematic representation of transaction level database and queries required to prepare data tables for the simulation framework

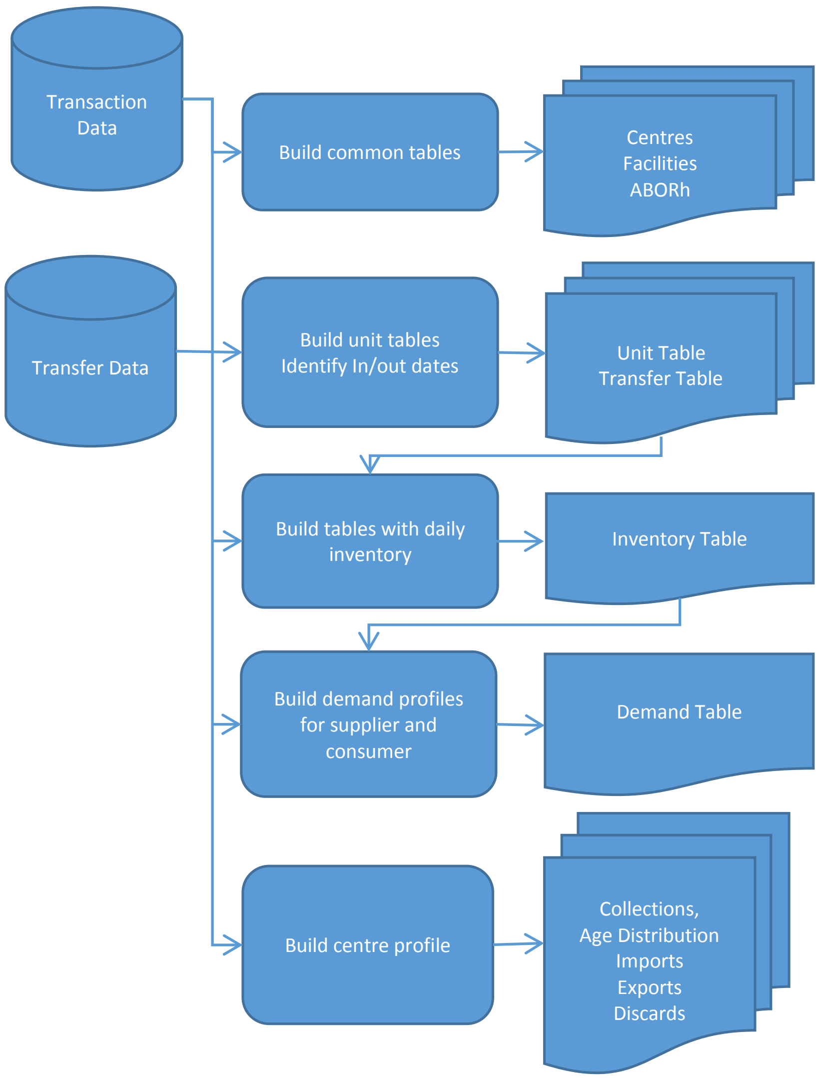

This process is repeated for the transaction and transfer data from each distribution centre to prepare the data for use with the simulation model.

Application Database

The application database stores data to control the framework and serves as a repository for output information from the simulation model. It stores information about user specifiable model parameters. Also retained in the application database is data related to simulation control parameters, such as length of runs, number of replications, and warm-up period. In instances where the framework is used to execute a series of experiments, the application database also stores the experiment parameters. When model runs are complete, results are output to the application database for review.

Simulation Framework

The simulation framework is built on a paradigm in which it is assumed that there is a single distribution centre of a blood agency (the "supplier") that collects red blood and distributes it to a set of consumers ("consumers") within its catchment area. The distribution centre is assumed to exist within a network of other centres, which may exchange red blood cells (RBC) with the modelled centre as either imported or exported units. Consumers are assumed to manage their own local supply of RBC; local inventory, local demand, and local ordering behaviour is, however, incorporated into the framework paradigm to allow for a simulated end-to-end supply chain.

Producers and consumers are modelled as separate object classes with the framework. Each object has a series of properties that define the state of the object at any given time, and a series of methods that change or update the system's state.

The Supplier Object

The supplier object represents the blood producer in the model. It is assumed that there is only one supplier in each region and that all consumers receive products from that sole supplier. The supplier object has a number of properties that are either informational in nature, define parameters used by the simulation to update the state of on hand inventory, bring in new inventory from collections, import units from or export units to another region, and fill consumer orders. For a complete list of methods in the supplier object see Appendix A.

Blood Collections

As with the Québec model, blood collections occur on a daily basis based on target volumes determined several months in advance, with some adjustment for the current inventory state. Overall daily collections are sampled from a day of week specific Poisson (a Normal approximation is used in the cases where $\lambda > 200$) distribution and units are then assigned ABO/Rh status from an empirical distribution based on historical collections. To maintain supplier inventory levels, a control policy is used similar to that used in the Quebec model. A seven day moving average of the supplier inventory is tracked and when the moving average of the supplier inventory deviates from the target level by greater than +/- one day collection levels are adjusted in an effort to bring the inventory back to target levels. Instead of the flat adjustments used in the Quebec model, collections are adjusted according to the inverse ratio of the inventory level to the target inventory. The level of adjustment to collections is subject to a cap of +/- 20%, preventing unreasonable collection levels. This allows for a more gradual adjustment than the Quebec control policy of flat adjustments according to inventory levels.

Testing and Issuing

Unlike with the Québec model, units are not held for a constant, one-day period after collection to account for testing and production, before being released into the general inventory. New units are instead assigned an age from a day of week specific empirical distribution representing the historical age of inventory when it becomes ready to ship. This allows for variations in the time required for testing and production procedures, as well as accommodating the reduced shelf life of irradiated units. It is assumed that there are no losses from production processes, Quality Control, or testing.

Imports and Exports

Unlike the Québec model, which operated as a closed system, the reusable version allows for imports of units from and exports of units to centres outside of the region serviced by the distribution centre. It is assumed that import (or export) capacity exists elsewhere in the network and thus, whenever a centre requests units for import, these units are always made available. To keep the supplier from importing (or exporting) unrealistic numbers of product, a 7-day moving average of imports and exports is kept for each blood group and type and compared to historical averages. If the average number of units imported or exported over the past 7 days does not exceed the historical average, an import or an export is permitted. Imported units arrive in a manner similar to collected units. However, since they are requested directly from a distribution centre elsewhere in the network, the ABO and Rh status are known. As with collections, the imported unit is assigned an age from an empirical distribution based on the arrival age of historical imports. Exported units are removed from the distribution centre's inventory in a FIFO manner.

The Consumer Object

The consumer object represents consumers in the simulation and is similar to the supplier object. The consumer object also encapsulates properties that update the state of the consumer's inventory, place orders for and receive new units from the supplier, and generate and fill demand. For a complete list of methods in the consumer object see Appendix B.

Consumer Inventory Policies

As with the Québec model, consumers are assumed to follow an ABO/Rh specific (R,S,s) inventory policy. The order up to levels (S) were determined from consumer target inventory levels provided by CBS. Appropriate order trigger points (s) were selected using the order up to level and the consumer order frequency. In cases where target inventory levels were not available for a given consumer site, an estimate was made by assuming the average number of days demand held by other consumers in the same region. It is also again assumed that when any new units are ordered from the distribution centre that all ABO/Rh groups are replenished back to the level of 'S'.

Consumer Ordering

Consumers place orders at the beginning of the day, prior to knowing the days' demand. Orders are assumed to be instantaneously replenished. Consumer orders are filled by the supplier according to a FIFO policy and a consumer sort order which was determined during the validation procedure.

Overall daily demand is drawn from a consumer specific, day of week specific Poisson distribution with ABO/Rh status assigned from an empirical distribution based on historical transfusions. In addition, a Zero Inflated Poisson distribution is used for to account for consumers with large numbers of days with no demand. The Zero Inflated Poisson distribution combines a Bernoulli distribution, which generates additional zeroes, with a Poisson process, resulting in an overall demand pattern that follows a Poisson distribution when demand occurs but has significantly more zero occurrences then a Poisson distribution alone (Lambert 1992).

Demand is filled from consumer inventory according to a FIFO policy. If insufficient inventory is available to fill demand, substitution of a compatible unit may occur at the consumer, or, if necessary an emergency order will be placed to the distribution centre.

Consumer transshipments do not occur in the reusable model due in part to the small number of units affected by the practice, as well as a lack of data to support confirmed transshipment arrangements between consumers.

Simulation Cycle

The reusable model follows a repeating daily cycle, much like the Québec model. On each simulated day, the model executes a sequence of calls to the supplier and consumer objects. Each day begins with a call to advance the supplier inventory. This ages the stock on hand at the supplier by one day and causes any stock with zero days of shelf life remaining to be outdated and to leave the system. A call is then made to the supplier object to have inventory arrive. After receiving units from its own collections process, the supplier object may import additional units from other suppliers. Once all incoming inventory is in place at the supplier, the simulation loops through each of the consumer objects and makes a call to advance the inventory. As was the case for the supplier, advancing the inventory at the consumer causes the stock on hand to age by one day. Any units with zero days of shelf life remaining are counted as outdated units and exit the system. Each consumer object then determines if an order is required. The consumer object evaluates, by blood group and type, its inventory position and compares it to a threshold level. If the current inventory level is less than the threshold level, the consumer issues an order for additional stock to return the inventory to a target level.

To simulate demand for product, a call is made to each consumer object. Demand is filled FIFO from available units on the shelf. If no units are available, compatible units may be substituted. If no exact or compatible unit is available, the consumer site issues a demand for additional units, on an emergency basis, from the supplier. If available, emergency units are transferred with zero delay to the consumer object, using the same logic as regular demand. The supplier may substitute a compatible unit, if no exact match is available. If no compatible unit is available, the demand is considered to be lost.

Once all demand from consumer sites has been met, the supplier may export units to other distribution centres, using logic similar to that for imports. The simulated day then ends and statistics are collected about inventory, outdates, arrivals, demand, and shortages. The statistical counters are reset and the daily cycle begins again.

At the end of each replication, a statistics collection routine is executed to summarize the daily simulation results. When all replications are complete, a final collection routine is executed to summarize results across all replications and to output the results to the application summary database. See Figure 5.

Figure 5: Simulation framework flow diagram

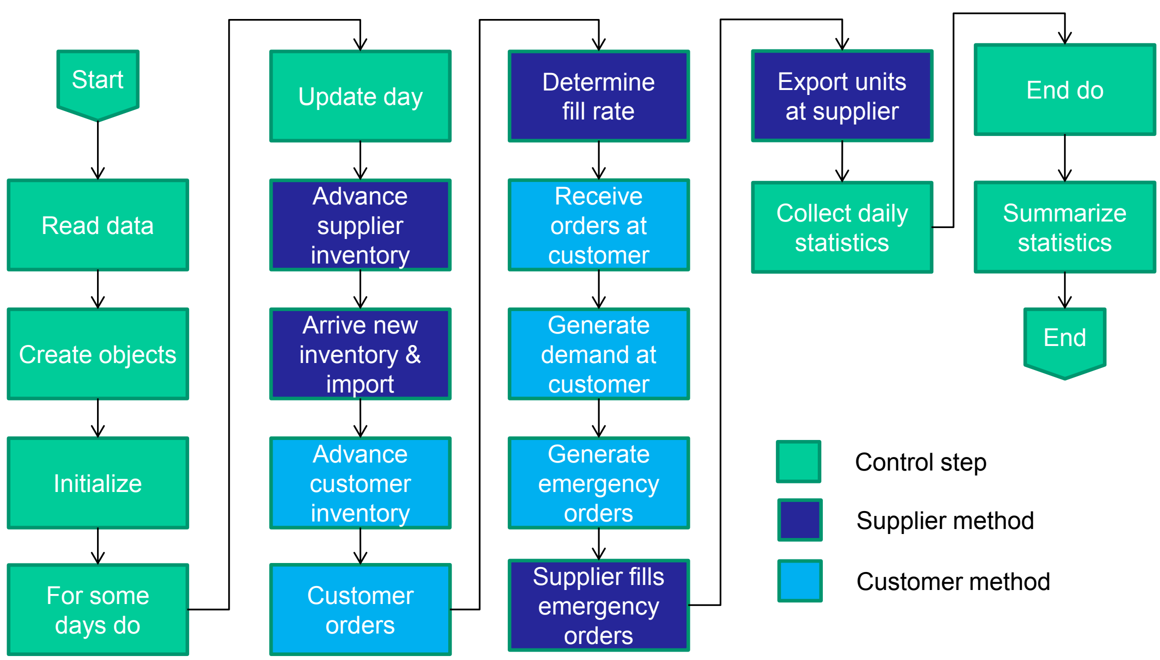

The process diagram for the simulation model. The steps from ‘update day’ to ‘collect daily statistics’ are repeated as a daily cycle.

5.1 Reusable Model Validation

To test the validity of the framework, three metrics, which were not inputs to the model, but rather were derived from its function, were compared against historical results for each site: the average inventory level, the average age of outgoing stock, and the average age of arriving product at each of the consumer sites in the catchment area.

A modified FIFO distribution policy was used to allow for practices, such as shipping preferentially younger units to certain consumers, as well as allowing older units to outdate in the case of high supplier inventory levels. In this policy, each consumer is assigned an age range from which they receive stock of a preferential age from the supplier. A consumer with an assigned range of (0, 41) for instance would receive stock in a true FIFO manner, whereas a consumer which receives primarily fresh blood might have a designated age range of (35, 41). The supplier fills consumer orders according to a modified FIFO policy where stock is issued FIFO, but only within the designated age range for the ordering consumer. Only if there is insufficient stock available within the target age range to fill an order will the supplier issue stock from outside the age range.

Both the age ranges and consumer sort order were determined using an exhaustive search process. Consumers were initially ordered from oldest average age of received stock to youngest and assigned age ranges of (0, 41). The age of received stock was then compared to the historical value with a prediction interval for each consumer. The age range for the first consumer in the sort order where the historical age of stock received did not fall within the prediction interval would then be narrowed, with the upper limit lowered in the case where the modeled age was too young, or the lower limit raised if it was too old. If an appropriate age range could not be found the consumer would be advanced in the sort order and the process repeated.

The three distribution centres selected for shorter shelf life experiments were validated under the current 42 day shelf life. The distribution centres represent larger (Centre X), a medium (Centre Y), and smaller (Centre Z) sized distribution centres within the CBS network. Each model was run for a total of 12 replications of one year, with a 42-day warm-up period. Models were run under the assumption that import and export practices mirrored the historical record. The comparisons suggested that there was no reason to assume that the model results differed from the historically observed values. As with the validation of the Québec model we note that the simulation has lower variability than the real world system due to longer simulation runs and the absence of exogenous events in the model. See Tables 17-19.

Table 17: Model validation – Average daily inventory at distribution centre

	Historical			Model			
Centre	$\bar{x}$	σ	n	$\bar{x}$	σ	n	p-value
X	8197.5	1204.2	12	8175.0	433.1	12	0.97
Y	1621.3	308.4	12	1628.2	9.8	12	0.94
Z	471.8	78.8	12	471.2	14.5	12	0.75

Table 18: Model validation – Average daily consumer demand

	Historical			Model			
Centre	$\bar{x}$	σ	n	$\bar{x}$	σ	n	p-value
X	879.9	32.4	12	880.5	9.4	12	0.86
Y	165.7	7.2	12	166.1	1.2	12	0.63
Z	47.9	3.0	12	47.7	4.9	12	0.99

Table 19: Model validation – Average age at receipt at consumer

	Historical			Model			
Centre	$\bar{x}$	σ	n	$\bar{x}$	σ	n	p-value
X	30.1	1.5	12	29.9	0.8	12	0.48
Y	31.3	1.0	12	30.8	0.1	12	0.18
Z	30.1	1.5	12	29.9	0.8	12	0.48

Chapter 6 Reusable Model Experiments

Each model was then run under the assumption of a 42 day shelf life for RBC to create a baseline scenario. The baseline scenario was run for 100 replications of one year under the method of replication/deletion, with a warm up period of 42 days. As with the Québec model, the replications were selected to provide an absolute error of no more than 1% on metrics of interest, and the warmup period was determined using the Welch method as described in Law (2006). The baseline scenario was used to establish the benchmark for evaluating shorter shelf policies. The results of the baseline runs appear in Table 20. It may be noted that wastage and shortage rates are inversely related to the volume of demand experienced at the distribution centre; larger centres have fewer outdates and shortages. The average age of RBCs on transfusion under the 42 days baseline was 18 days.

Table 20: Baseline wastage and shortage rates for 42 day shelf life.

Centre	Wastage	Shortage	Emergency Orders
X	1.56%	0.04%	2.03%
Y	4.22%	0.07%	3.56%
Z	6.44%	0.82%	6.41%
Overall	2.24%	0.08%	2.45%

Shorter shelf-lives of 14, 21, and 28 days were evaluated using three inventory scenarios. These three values of shelf life were selected to match the values tested with the blood distribution network model for the Province of Québec. Each of the simulations was run for 100 replications of one year with a 42 day warm up. To test the impact of system inventory at each rated shelf life, the ordering behaviour of the supplier and consumers was adjusted such that the total inventory held in the system, as measured by days' supply, equaled 40%, 50%, or 60% of the assumed RBC shelf life. As with the Québec model, the results were compared against both the baseline values, and the currently accepted thresholds of 5% or less wastage and a minimum 95% service level.

28 Day Shelf Life

In Table 21 the results for the three distribution centres in the model under the assumption of a 28 day shelf life are presented. In the table, wastage and shortage rates are shown for each of the three networks, under a scenario in which total inventory is either 40% (12 days), 50% (14 days), or 60% (16 days) of the maximum shelf life. The average age of RBCs on transfusion under a 28 day shelf life was 15 days, a reduction of 3 days from the baseline.

Table 21: Network results for a 28 day shelf life.

Inventory Level	Centre	Wastage	Shortage	Emergency Orders
40% of Shelf life	X	1.31%	0.73%	2.14%
	Y	4.10%	0.47%	3.20%
	Z	7.84%	5.04%	9.15%
	Overall	2.11%	0.88%	2.61%
50% of Shelf life	X	1.53%	0.36%	1.42%
	Y	4.33%	0.28%	2.52%
	Z	8.63%	3.98%	7.29%
	Overall	2.35%	0.51%	1.85%
60% of Shelf life	X	1.80%	0.18%	1.09%
	Y	4.73%	0.18%	2.22%
	Z	9.30%	3.66%	6.90%
	Overall	2.67%	0.33%	1.52%

As can be seen, system wide shortage and outdate rates are inversely related to the volume of product; Centre X, which is the largest network observes the smallest outdate rate and emergency order rate, although it is slightly outperformed in some cases by Centre Y in terms of shortages, Centre Z, the smallest network, experiences the largest outdate, shortage, and emergency order rates. It may also be seen that, as expected, shortage rates decrease while wastage rates increase, as more inventory is held. For comparison, the shortage and outdate rates at 40% and 60% of shelf life were compared, on a site by site basis, with the rates observed when inventory equals 50% of total shelf life. In all instances, the differences in outdate rates were statistically significant. All results were statistically significant at the 5% level, except the shortage rates for Centres X and Y at 16 days inventory held. Overall, at a 28 day shelf life it can be seen that outdates and shortages rates would be considered acceptable for large (X) and medium (Y) centres based on the currently accepted targets. The smaller centre (Z) had difficulty identifying an inventory policy that would result in acceptable levels of outdates and shortages, even at this relatively long shelf life.

21 Day Shelf Life

Continuing with the same form of analysis, in Table 22 the model results for a 21 day shelf life are presented. The average age of RBCs on transfusion under a 21 day shelf life was 11 days, a reduction of 7 days from the baseline.

Table 22: Network results for a 21 day shelf life.

Inventory Level	Centre	Wastage	Shortage	Emergency Orders
40% of Shelf life	X Y Z Overall	1.78% 5.61% 11.18% 2.90%	1.35% 0.99% 6.31% 1.51%	3.67% 5.44% 11.79% 4.29%
50% of Shelf life	X Y Z Overall	2.06% 6.00% 11.80% 3.22%	0.78% 0.54% 5.41% 0.94%	2.35% 3.57% 9.70% 2.85%
60% of Shelf life	X Y Z Overall	2.53% 6.62% 12.53% 3.74%	0.44% 0.34% 4.86% 0.62%	1.63% 2.74% 8.68% 2.10%

At a 21 day shelf life it is evident that inventory management becomes more difficult for small (Z) and medium size (Y) centers. Table 22 shows that wastage rates for Centres Y and Z exceed the target threshold of 5%, under all three inventory scenarios. When only 8 days of inventory are held, shortage rates at both the medium and large centres also exceed the target of 1%. The small centre was observed to experience large outdate and shortage rates at a 21-day shelf life under all three inventory scenarios.

14 Day Shelf Life

The results for a 14 day shelf life are presented in Table 23. The average age of RBCs on transfusion under a 28 day shelf life was 8 days, a reduction of 10 days from the baseline.

Table 23: Network results for a 14 day shelf life.

Inventory Level	Centre	Wastage	Shortage	Emergency Orders
40% of Shelf life	X	3.22%	2.49%	8.06%
	Y	10.50%	1.89%	10.78%
	Z	18.54%	11.09%	20.02%
	Overall	5.29%	2.78%	9.00%
50% of Shelf life	X	3.83%	1.51%	5.07%
	Y	11.08%	1.31%	7.36%
	Z	18.89%	9.08%	15.85%
	Overall	5.89%	1.81%	5.89%
60% of Shelf life	X	4.57%	1.14%	3.54%
	Y	11.69%	0.91%	5.41%
	Z	19.22%	8.33%	13.37%
	Overall	6.59%	1.42%	4.25%

The model results show that all centres, regardless of size, struggle with inventory management when RBC have a rated shelf life of 14 days. Extremely large wastage and shortage rates were observed for the small network (Z), while the medium sized network (Y) was observed to experience large wastage rates and generally unacceptable levels of product shortage. Large networks (X) continue to be better able to manage wastage at a 14 day shelf life. However, wastage rates were still consistently above the 1% threshold.

Consumer Results

The effect of size on the ability to manage inventory applies to consumers as well as the overall distribution networks. To evaluate the impact of consumer size, outdate and shortage rates for the consumers within the three simulated networks was recorded. The results were aggregated on the basis of the size of the consumer's transfusion program. For the purposes of this analysis, all consumers receiving blood and blood products from Canadian Blood Services were ranked by annual volume and the resulting list was divided into thirds, consisting of small consumer sites (0-400 units/year), medium consumer sites (401-1500 units/year) and large consumer sites (1500+). Of the 168 consumer sites in the three regional networks, 56 were rated as small, 46 as medium-sized, and 66 as large.

In baseline runs, with a shelf life of 42 days for RBCs, using existing inventory policies, simulated wastage rates varied from 0.6% at large consumers to 30.03% at small consumers. Shortage rates varied between 0.05% and 0.45%, depending on consumer size.

Table 24: Baseline wastage and shortage rates for consumers assuming a 42 day shelf life.

Consumer Size	Wastage	Shortage	Emergency Orders
Large	0.60%	0.05%	1.33%
Medium	2.74%	0.19%	8.67%
Small	30.03%	0.45%	25.52%
Overall	1.36%	0.07%	2.45%

Tables 25-27 show the impact of a shorter shelf life for RBC at consumers. In general, as shelf life is reduced, wastage rates and shortage rates increase, with the effect of reduced shelf life becoming more noted in smaller consumers. Again, it may be seen that wastage rates decrease with lower levels of system inventory, while shortages decrease as more inventory is held in the system. Please note that wastage and shortage rates are calculated as a fraction of blood transfused.

Table25: Consumer results for a 28 day shelf life

Inventory Level	Consumer Size	Wastage	Shortage	Emergency Orders
40% of Shelf life	Large Medium Small Overall	1.23% 2.57% 30.21% 1.98%	0.70% 2.14% 3.19% 0.88%	1.31% 10.60% 25.57% 2.61%
50% of Shelf life	Large Medium Small Overall	1.45% 2.63% 30.02% 2.25%	0.37% 1.34% 2.74% 0.51%	0.68% 7.86% 25.16% 1.85%
60% of Shelf life	Large Medium Small Overall	1.66% 4.37% 31.66% 2.58%	0.22% 1.02% 2.01% 0.33%	0.43% 7.29% 25.58% 1.52%

Table 26: Consumer results for a 21 day shelf life.

Inventory Level	Consumer Size	Wastage	Shortage	Emergency Orders
40% of Shelf life	Large Medium Small Overall	1.36% 5.75% 40.66% 2.75%	1.27% 3.35% 3.85% 1.51%	2.70% 15.50% 25.45% 4.29%
50% of Shelf life	Large Medium Small Overall	1.58% 6.64% 40.60% 3.07%	0.75% 2.33% 3.19% 0.94%	1.48% 11.39% 26.08% 2.85%
60% of Shelf life	Large Medium Small Overall	1.92% 8.22% 41.46% 3.56%	0.47% 1.60% 2.78% 0.62%	0.90% 8.81% 26.28% 2.11%

Table 27: Consumer results for a 14 day shelf life.

Inventory Level	Consumer Size	Wastage	Shortage	Emergency Orders
40% of Shelf life	Large Medium Small Overall	1.71% 15.46% 59.56% 5.22%	2.46% 5.51% 4.52% 2.77%	6.94% 25.41% 26.12% 8.99%
50% of Shelf life	Large Medium Small Overall	2.19% 16.27% 60.77% 5.77%	1.54% 3.84% 4.61% 1.81%	4.06% 19.05% 28.11% 5.89%
60% of Shelf life	Large Medium Small Overall	2.74% 17.13% 60.44% 6.37%	1.19% 3.21% 3.64% 1.42%	2.65% 15.27% 26.23% 4.25%

From the results of Tables 25-27, it can be seen that small consumers will have difficulty managing inventory as RBC shelf life is reduced. Even at a 28 day shelf life, wastage is substantial and shortages are evident for small consumers. When shelf life is 14 days, wastage rates at smaller consumers are prohibitive at over 12 times the currently accepted target levels. Large consumers, conversely, are able to maintain acceptable wastage and shortage rates even as RBC shelf life decreases to as low as 21 days. At a 14 day shelf life, large consumers would still be able to find inventory policies that maintain acceptable wastage rates. However, even large consumers would encounter difficulties with shortage rates. Medium sized consumers fall somewhere in the middle – a 28 day shelf life would be compatible with accepted thresholds, but a 21 or 14 day shelf life would result in substantially increased outdates and shortages.

Results by Blood Type

The impact of a shorter shelf life is not uniformly distributed across the various blood types. Even under the current 42 day shelf life, blood is not collected in proportion to the ratios at which it is naturally found in the general population, rather specific blood types are targeted and efforts are made to recruit and maintain donors of those types. The O- blood type, for instance, accounts for roughly 11% of collected units in the historical data, but is found in only about 6% of the Canadian population. As the shelf life for RBCs is reduced shortages occur in the more universal blood types (those that can be used as substitutes for a large portion of demand), particularly O- and A-, while outdates occur more heavily in the remaining types, generally spread across the Rh positive types. Shortages and outdates by ABO/Rh for each shelf life where 50% of the shelf life is held as inventory are shown in Table 28.

Table 28: Percentage of Outdates and Shortages by ABO/Rh

Centre	ABO/Rh	28 Day		21 Day		14 Day	
		Outdates	Shortages	Outdates	Shortages	Outdates	Shortages
X	A-	9.56%	36.93%	11.00%	35.58%	12.47%	31.80%
	A+	17.08%	0.92%	20.07%	2.90%	23.43%	9.55%
	AB-	1.85%	0.82%	1.45%	0.61%	1.15%	0.56%
	AB+	19.87%	0.00%	14.46%	0.00%	9.27%	0.00%
	B-	0.91%	1.29%	0.55%	2.07%	1.08%	0.98%
	B+	29.66%	0.00%	23.16%	0.00%	15.39%	0.00%
	O-	5.69%	59.75%	9.01%	58.22%	13.56%	48.76%
	O+	15.38%	0.29%	20.30%	0.62%	23.65%	8.34%
Y	A-	3.94%	15.12%	5.21%	15.31%	6.18%	18.03%
	A+	16.23%	0.07%	20.15%	0.22%	24.66%	0.63%
	AB-	3.61%	0.04%	2.83%	0.07%	1.87%	0.05%
	AB+	21.82%	0.00%	16.47%	0.00%	9.72%	0.00%
	B-	1.04%	28.98%	1.17%	23.97%	1.22%	14.77%
	B+	33.05%	0.00%	26.24%	0.00%	17.68%	0.00%
	O-	5.92%	55.77%	8.03%	60.37%	11.21%	65.27%
	O+	14.39%	0.02%	19.91%	0.06%	27.47%	1.26%
Z	A-	2.00%	13.65%	2.70%	16.11%	3.74%	19.89%
	A+	8.48%	0.00%	14.79%	0.00%	25.44%	0.00%
	AB-	1.94%	1.18%	1.44%	1.66%	1.11%	1.82%
	AB+	61.25%	0.00%	46.96%	0.00%	26.54%	0.00%
	B-	0.78%	20.67%	0.98%	19.67%	1.05%	13.58%
	B+	14.15%	0.00%	15.83%	0.00%	13.89%	0.00%
	O-	2.89%	64.50%	4.96%	62.52%	8.39%	64.41%
	O+	8.51%	0.00%	12.35%	0.04%	19.85%	0.30%

As Table 28 shows, types O- and A- consistently account for large portions of shortages, with O- accounting for nearly 50% or more of all shortages at all three centres under each shelf life. The medium and small distribution centres also see a large portion of their shortages accounted for by B- blood (between 14% and 29%), although this is not observed at the largest centre. By contrast, Rh + blood types make up a small portion of the observed shortages with only O+ and A+ types accounting for more than 1% each of overall shortages, and then only at a 14 day shelf life.

Outdates are more distributed than shortages, occurring in all blood types at some level. Rh positive blood types, however, consistently account for a larger portion of outdates than Rh negative blood types. Interestingly, as the shelf life shortens O- and A- begin to account for increasing portions of outdates, even while accounting for the majority of shortages and thus seemingly being in demand. This is largely explained by the reliance of small consumers on these two types to serve all of their demand through compatibility matching. These small consumers thus hold disproportionate levels of O- and A- blood while also suffering from the highest outdate rates of all consumers, due to their small size and typically unreliable demand.

These results suggest that under a shorter shelf life, simply collecting more blood in general will not address the problem of shortages, and may in fact result in additional unwanted wastage of less desirable blood types. A shorter shelf life, would in fact require blood suppliers to collect additional Rh-, particularly O- and A-, blood to maintain an adequate supply. As these types are already over collected, relative to their occurrences in the population, it is unclear how much further collection ratios can be pushed in favour of these types. This presents a significant challenge to the suppliers.

Chapter 7 Conclusion

In completing this project, a discrete event simulation of the Héma-Québec blood system was implemented in Microsoft Visual Basic .net. Additionally, a reusable modelling framework based on the Canadian Blood Services system was created, also in Visual Basic .net, and was applied to three CBS distribution centres of varying size. The models were used to test three sets of inventory policies under the assumption of a reduced shelf-lives of 28, 21, and 14 days for RBCs.

The results show that under the currently accepted target metrics of 5% or fewer outdates and a service level of 99% or greater, some reduction in shelf life is feasible. At a 28 day shelf life the Héma-Québec system, as well as the large and medium CBS distribution centres were able to meet both target metrics. The small CBS distribution centre, however, was unable to meet the target metrics and would likely need to import additional product from larger distribution centres to meet demand. At a 21 day shelf life, both the Héma-Québec system and the large CBS distribution centre continue to be able to meet the target metrics. However, the medium CBS distribution centre was no longer able to meet the 5% outdate target. At a 14 day shelf life, none of the distribution networks were able to meet both target metrics. In particular, shortage rates were consistently above the 1% maximum target in all cases, except one of the medium CBS distribution centres. While these targets are sensible in the current blood system with a 42 day shelf life, they may ultimately not be achievable targets at a reduced shelf life, particularly for smaller consumer sites which have difficulty meeting them at even a 28 day shelf life.

It should be noted that shortages do not typically occur uniformly across all blood types, but are observed at higher rates in the more universal, and thus more desirable, types such as type O and Rh negative variants. As the overall level of blood available in the distribution networks decreases, the demand for these more universal types, particularly O-, increases due to compatibility matching from shortages at consumer sites and because smaller consumers cease to hold less common blood. Donors with these blood types are already highly desired and thus it is unclear if the collection ratio of these types could be further increased in response to the additional demand that a reduced shelf life would entail. Any increases in outdate rates would also require additional collections to replenish inventory depleted by wasted units. This may prove difficult in practice as a 1% increase in outdates nationally would require a corresponding increase of over 11,000 donations to maintain inventory level.

These results demonstrate that the effects of a reduced shelf life for RBC are influenced by both the size of the consumer site and the size of the overall distribution network. As noted in Chapter 2, previous publications on shorter RBC shelf life have focused either on a single consumer of comparable size to a tertiary healthcare facility in a major Canadian city, or have treated demand in aggregate. Thus the existing literature has not, and indeed could not, detect the interaction between consumer and network size and shorter shelf life as has demonstrated in this study.

While results suggest great challenges associated with managing the supply of RBCs under a reduced shelf life, there are a number of practical options available to improve both outdate and shortage rates. Firstly, the “one size fits all” policy approach assumed in this study would almost certainly not be optimal. Smaller, more remote consumer sites might be forced to hold higher inventories than larger consumers to mitigate shortages and thus higher outdate targets for smaller consumers may be necessary to accommodate a shorter RBC shelf life. The costs of the increased outdates at smaller consumers may, however, be mitigated by lower ordering costs associated with the higher inventories as suggested by Blake & Hardy (2014a).

Smaller consumer sites may also reduce the impact of a shorter shelf life by entering into transshipment agreements with larger consumer sites and/or joining together with other consumers to form regional "hub" blood banks.

Four models applicable to the blood supply in Canada were developed as part of this project. In addition to evaluating the impact of shorter shelf life, the reusable modelling framework has also been used to evaluate CBS inventory policies under the existing 42 day shelf life (see Blake & Hardy, 2014a). The framework also provides the conceptual basis for two further papers (Blake & Hardy, 2014b), (Blake et al, 2013) appearing in the literature.

The models developed for this project have significant potential for further development. With the addition of a dynamic import and export algorithm for distribution centres, the simulation can be run as a network model, potentially encompassing an entire national blood supply chain; this could allow for improvements in the efficiency of inventory rebalancing between distribution centres. Such models will be particularly important if shorter shelf life is adopted, since it is not apparent that blood can be as easily collected and transferred between sites as has been assumed in this study. We note, in particular, that small, geographically isolated supplier sites may find it difficult to import sufficient products to maintain service levels when shelf life decreases. Finally, there is excellent potential for integrating exact or heuristic methods with the simulation framework to search for optimal inventory policies.

Our results clearly show that outdates, rather than shortages, are the primary outcome of a shorter shelf life. Thus, as shelf life decreases and outdates increase additional blood collections will be necessary. An important area for future research is the study of the effectiveness of blood suppliers to influence both the number and type of blood donations. In all of the models in this study, an inventory control policy based on changes to collections was vital to maintaining a stable inventory. While model control policies were based on both historical collections data, the true effectiveness of suppliers to encourage more (and more targeted) donations, is not well known. Both the ability of the supplier to influence collection levels, and the cost effectiveness of the tools used to achieve this influence would be crucial to determining if a reduced RBC shelf life would be truly sustainable in any blood system.

Bibliography

Atkinson, M. P., Fontaine, M. J., Goodnough, L. T., & Wein, L. M. (2012). A novel allocation strategy for blood transfusions: investigating the tradeoff between the age and availability of transfused blood. *Transfusion, 52*(1), 108-117.

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

Blake, J., & Hardy, M. (2014a). A generic modelling framework to evaluate network blood management policies: The Canadian Blood Services Experience. *Operations Research for Health Care*, pp. 116-128.

Blake, J., & Hardy, M. (2014b). Evaluating the impact of a shorter shelf life for red blood cells with a generic simulation. *Industrial and Systems Engineering Research Conference.* Montreal.

Blake, J., Hardy, M., Delage, G., & Myhal, G. (2013). Déjà-vu all over again: using simulation to evaluate the impact of shorter shelf life for red blood cells at Héma-Québec. *Transfusion*, pp. 1544-1558.

Blake, J., Heddle, N., Hardy, M., & Barty, R. (2010). Simplified platelet ordering using shortage and outdate targets. *International Journal of Health Care Management, 1*(2), 144-156.

Brodheim, E., & Prastacos, G. (1979). The Long Island blood distribution system as a prototype for regional blood management. *Interfaces, 9*(5), 3-20.

Brodheim, E., Derman, C., & 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., & Prastacos, G. (1976). Setting inventory levels for hospital blood banks. *Transfusion, 16*(1), 63-70.

Canadian Blood Services. (2013). *A Report to Canadians 2012/2013.* Ottawa: Canadian Blood Services.

Canadian Blood Services. (2014). *Annual Report 2013-2014.* Ottawa: Canadian Blood Services.

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

Fontaine, M., Chung, Y., Rogers, W., Sussman, H., Quach, P., Galel, S., . . . Erhun, F. (2009). Improving platelet supply chains through collaborations between blood centres and transfusion centres. *Transfusion, 49*(10), 2040-2047.

Frank, S. M., Abazyan, B., Ono, M., Hogue, C. W., Cohen, D. B., Berkowitz, D. E., . . . Barodka, V. M. (2013). Decreased Erythrocyte Deformability After Transfusion and the Effects of Erythrocyte Storage Duration. *Anesthesia & Analgesia*, 975-981.

Friedman, B., Abbott, R., & Williams, G. (1982). A blood ordering strategy for hospital blood banks derived from a computer simulation. *American Journal of Clinical Pathology, 78*(2), 154-160.

Greinacher, A., Fendrich, K., Brzenska, R., Kiefel, V., & Hoffmann, W. (2011). Implications of demographics on future blood supply:a population-based cross-sectional study. *Transfusion*, 702-709.

Héma-Québec. (2014). *2013-2014 Annual Report.* Montreal: Héma-Québec.

Hesse, S., Coullard, C., Daskin, M., & Hurter, A. (1997). A case study in platelet inventory management. In G. Curry, B. Bidanda, & S. Jagdale (Ed.), *Sixth Industrial Engineering Research Conference Proceedings* (pp. 801-806). Norcross, GA: IIE.

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

Katsaliaki, K., & Brailsford, S. (2007). Using simulation to improve the blood supply chain. *Journal of the Operational Research Society, 58*(2), 219-227.

Koch, C., Liang, L., Sessler, D., Figueroa, P., Hoeltge, G., Mihaljevic, T., & Blackstone, E. (2008). Duration of red-cell storage and complications after cardiac surgery. *New England Journal of Medicine, 358*(12), 1229-1239.

Lambert, D. (1992). Zero-Inflated Poisson Regression, with an Application to Defects in Manufacturing. *Technometrics*, 1-14.

Law, A. (2006). *Simulation Modeling and Analysis, 4th Edition.* New York: McGraw-Hill.

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

Nahmias, S. (1978). The fixed charge perishable inventory problem. *Operations Research, 26*(3), 464-481.

Pierskalla, W., & Roach, C. (1972). Optimal issuing policies for perishable inventory. *Management Science, 18*(11), 603-614.

Qinglin, D. T., & Warren, L. (2014). Optimization of blood supply chain with shortened shelf lives and ABO compatibility . *International Journal of Production Economics*, 113-129.

Rytilä, J. S., & Spens, K. M. (2006). Using simulation to increase efficiency in blood supply chains. *Management Research News*, 801-819.

Simonetti, A., Forhsee, R., Anderson, S., & Walderhaug, M. (2013, August 29). A stock-and-flow model of the US blood supply. *Transfusion*. doi:10.1111/trf.12392

Simonetti, A., Forshee, R. A., Anderson, S. A., & Walderhaug, M. (2014). A stock-and-flow simulation model of the US blood supply. *Transfusion*, 828-838.

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

Steiner, M. E., & Stowell, C. (2009). Does red blood cell storage affect clinical outcome? When in doubt, do the experiment. *Transfusion*, 1286-1290.

Telkin, E., Gurler, U., & Berk, E. (2001). Age-based versus stock level control policies for a perishable inventory system. *European Journal of Operational Research, 134*(2), 309-329.

Tinmouth, A., Fergusson, D., Yee, I., & Hebert, P. (2006). Clinical consequences of red cell storage in the critically ill. *Transfusion, 46*(11), 2014-2027.

Yegul, M. (2007). *Simulation analysis of the blood supply chain and a case study.* Master's Thesis, Middle East Technical University, Industrial Engineering.

Appendix A: Supplier Object Properties and Methods

Table A1: Supplier Properties

Property	Parameters	Description
MyName		Text name of object instance
Wastage	{intType}	The number of units wasted today, or the number of units wasted in blood group *intType*, if specified
Shortage	{intType}	The number of units demand not filled today, or the number of units not filled in blood group *intType* if specified
Inventory	intAge, {intType}	The total number of units with age *intAge* days remaining shelf life in inventory, or the total number of units with age *intAge* remaining of blood type *intType* remaining if specified
Arrivals	{intType}	The total number of units arriving today, or the total number of units of type *intType* if specified
Demand	{intType}	The total number of units demanded today, or the total units demanded of type *intType* if specified
AgeUsed	{intType}	The average age of outgoing stock, or the average age of outgoing stock of type *intType* if specified
OrderSize	{intType}	The number of units received today, or the total number of units received of type *intType* if specified
Imports	{intType}	The number of units imported today, or the total number of units imported of type *intType* if specified
Exports	{intType}	The number of units exported today, or the total number of units exported of type *intType* if specified
Filled	{intType}	The total number of units shipped to consumer sites today, or the total units shipped of type *intType* if specified
EmergencyFilled	{intType}	The total number of units to fill emergency demand, or the units required to fill emergency demand of type *intType* if specified
ArrivalTarget	intDay	The average expected units to arrive from testing on *intDay* of the week (supplier) or the average expected demand on intDay of week
ArrivalProb	intDay	The probability that any units/demand will arrive on *intDay* of the week
ArrivalAge	intDay, intAge	The probability that a unit arriving from testing on *intDay* of the week has a remaining shelf life of *intAge* days
ImportTarget	intDay	The average number of units to import on day *intDay*

Property	Parameters	Description
ImportType	intType	The probability a unit imported is of type *intType*
ImportProb	intDay	The probability that at least one unit is imported on day of week *intDay*
ImportAge	intDay, intAge	The probability that a unit imported on day of the week *intDay* is of age *intAge*
ExportTarget	intDay	The average number of units to be exported on day of week *intDay*
ExportType	intType	The probability a unit exported is of type *intType*
ExportProb	intDay	The probability that at least one unit is exported on day of the week *intDay*
TypeMatch	intPType, intDType	The number of emergency order units filled with compatible units of type *intDType* in response to demand of type *intPType*

Table A2: Supplier Methods

Method	Description
InitializeInventory	Causes the inventory at the supplier to be initialized at the start of the simulation run. The routine sums the total expected demand from all consumer sites and brings this number of units into inventory (units are assumed to have a randomly assigned remaining shelf life). Inventory is initialized so that all consumer sites start the simulation with their desired inventory level.
AdvanceInventory	Causes stock on hand to age by one day. Stock with remaining shelf life of 0 days is outdated.
ArriveInventory	Causes units to arrive at the supplier site from testing. Arrivals are assumed to be Poisson distributed with a mean arrival rate specific to day of week. The mean arrivals may, however, be adjusted up or down if a 7-day moving average is below or above a target inventory. Using the adjusted mean arrival rate, a random number of units is generated. These units are assigned a blood type from an empirical distribution and a remaining shelf life, again drawn from an empirical distribution. The arriving inventory is then entered into inventory.
Import	Causes units to be imported from other supplier sites. The number of units to import is drawn from a Poisson distribution, specific to day of week. The units are assigned a blood group and type and are assigned a remaining shelf life (both values are drawn from empirical distributions) and are entered into inventory.
Export	Causes units to be exported to other distribution centres. The number of units of demand for exports is drawn from a Poisson distribution, specific to day of week. The demand units are assigned a blood group and type, drawn from an empirical distribution, and if units are available are removed into inventory.
FillOrder	Calculates the total demand for units of each blood type and determines if there is sufficient inventory to fill all demand. If there is not, FillOrder returns a scaling factor used to adjust consumer demand to meet available inventory.
TransferInventory	Removes units from inventory and enters them into a temporary array to be passed to the consumer in response to daily demand.
EmergencyOrder	Causes units to be withdrawn from inventory and transferred to the consumer in response to emergency orders.
CheckTypeMatch	Checks for compatible units in the event that insufficient stock is available to fill emergency orders.
DoTypeMatch	Causes compatible units identified by *CheckTypeMatch* to be withdrawn from inventory and transferred to the consumer.
CollectStatistics	Collates statistics for today and then clears the statistical counters.

Appendix B: Consumer Object Properties and Methods

Table B1: Consumer Properties

Property	Parameters	Description
MyName		Text name of object instance
Wastage	{intType}	The number of units wasted today, or the number of units wasted in blood group *intType*, if specified
Shortage	{intType}	The number of units demand not filled today, or the number of units not filled in blood group *intType* if specified
Inventory	intAge, {intType}	The total number of units with age *intAge* days remaining shelf life in inventory, or the total number of units with age *intAge* remaining of blood type *intType* remaining if specified
Demand	{intType}	The total number of units demanded today, or the total units demanded of type *intType* if specified
AgeUsed	{intType}	The average age of transfused stock, or the average age of transfused stock of type *intType* if specified
OrderSize	{intType}	The number of units received today, or the total number of units received of type *intType* if specified
ArrivalTarget	intDay	The average expected demand on intDay of week
ArrivalProb	intDay	The probability that any demand will arrive on *intDay* of the week
EmergDemand	{intType}	The number of units ordered on an emergency basis today, or the number of units ordered on an emergency basis of type *intType* if specified
EmergencyFilled	{intType}	The number of units received today, or the number of units received of type *intType* if specified
Transfused	{intType}	The number of units transfused to patients today, of the number of units transfused of type *intType* if specified
TypeMatch	intPType, intDType	The number of units of demand for type *intPType* which are filled with units of the compatible type *intDType*

Table B2: Consumer Methods

Method	Description
AdvanceInventory	Causes stock on hand to age by one day. Stock with remaining shelf life of 0 days is outdated.
OrderInventory	Causes a request for inventory replenishment to be issued to the supplier. Consumer sites observe their inventory on a blood-type by blood-type basis. The current inventory is compared against a consumer site specific inventory trigger. If inventory is below this level, an order is generated to return inventory to a consumer site specific inventory target. To minimize orders, it is assumed that if any blood type requires replenishment, orders are issued to return all blood types to their order up to level.
ReceiveInventory	Causes units to be withdrawn from the supplier and entered into the consumer's inventory. Units are entered into the consumer's inventory according to blood type and remaining shelf life.
GenerateDemand	Causes demand for product to be generated. A random variable is drawn to determine if any demand will be observed on this day of the week. If so, a second random variate is generated from a Poisson distribution specific to day of week to determine the total number of units to be transfused. The blood type is then determined for each demand request by sampling from a hospital specific empirical distribution. Requests are then filled, FIFO, from available inventory. Type matching for compatible units may or may not be allowed (this is an optional parameter in the model). Any demand that cannot be satisfied with stock available on the shelf generates an emergency order to the supplier.
ReceiveEmergOrder	Counts units received from the supplier as part of an emergency order
CheckTypeMatch	Checks for compatible units in the event that insufficient stock is available to fill demand.
DoTypeMatch	Causes compatible units identified by *CheckTypeMatch* to be withdrawn from inventory and transfused.
CollectStatistics	Collates statistics for today and then clears the statistical counters.