Universiteit Utrecht

Master Thesis Computing Science

Comprehensive Red Blood Cell Matching

Author:
R.H.G. van de Weem

Sanquin Supervisor:
Dr. ir. M.P. Janssen

Utrecht University Supervisor:
Dr. J.A. Hoogeveen

Utrecht University Second Supervisor:
Dr. ir. J.M. van den Akker

December 2020

Abstract

Advances in genotyping technology have opened up the possibility of typing donors and patients on many more antigens than just A, B and RhD, which are currently the standard. A future where all patients receive extensively matched Red Blood Cell (RBC) units is now foreseeable. Matching compatibly on all eleven clinically relevant minor antigens would eliminate almost all alloimmunization among patients and thereby make RBC transfusions safer and more effective. However, strictly compatible matches for all patients on all relevant antigens are unlikely due to an exponential increase in the number possible phenotypes. Large-scale extended matching should therefore neither decrease the availability nor increase the outdating of RBC units. We propose the MINRAR-Online Integer Linear Programming (ILP) approach which allows all AB-RhD compatible matches while minimizing the alloimmunization risk for the patients. This is achieved by allowing minor antigen mismatches at a cost based on the immunogenicity of the antigens. Furthermore, the ILP also contains terms in the objective to prevent shortages, outdating and alloimmunization in the long run. Simulations show that shortages and outdating can be prevented whilst reducing the alloimmunization risk per patient compared to previous work. Lastly, we investigated how the MINRAR-Online ILP can be extended to prioritize certain patient groups in the matching for which there is a larger incentive to prevent alloimmunization. Simulations of a single hospital show that this prioritization can be implemented effectively with patient group specific weights, while multi-hospital simulations show that the percentage of shortages for all considered patient groups can be kept below 1%.

Acknowledgements

Starting this thesis with only little background knowledge in biology meant that I had to familiarize myself quickly with the concepts of red blood cells, antigens, compatibility, immunogenicity and alloimmunization. This trend of familiarizing has continued throughout my thesis with visits to the OLVG Oost and UMCU hospitals. I have experienced a great willingness to explain how and why things are the way they are from everyone I got the opportunity to talk to, which was very helpful and motivating. I am especially grateful to Sanquin and in particular Mart Janssen for giving me the opportunity to write my thesis at the Transfusion Technology Assessment Group. Supervising my project together with Han Hoogeveen from the Department of Information and Computing Sciences at Utrecht University, I have both Han and Mart to thank for their excellent guidance, useful criticism and interesting discussions. I think that we were compatible from the beginning and I could not have written this thesis without the both of you. Furthermore, I would like to thank Masja de Haas and Jessie Luken for their expertise and enthusiasm to participate in brain storming sessions, which I have experienced as useful and insightful. I have also enjoyed the weekly meetings with everyone from the TTA group where we shared our progress, stories and sometimes frustrations. Finally, I would like to thank my friends and family for their moral support throughout this project and giving me the determination to make this thesis into what it is now.

– Ronald van de Weem, December 2020

Contents

1 Introduction

Blood transfusions are a vital part of modern health care. Various blood products exist which all serve specific purposes. One of the most common blood components are Red Blood Cells (RBC). Typically, these are transfused to improve the oxygen carrying capacity of the blood. The supply of blood products in the Netherlands comes from a large donor population which is managed by the Dutch blood bank Sanquin. Sanquin is responsible for recruiting and inviting donors, as well as processing and distributing the donated blood products among hospitals. In general, hospitals are responsible for matching blood to patients. Selecting suitable blood from a hospital inventory to patients who require a transfusion is a manual process. Matching donor RBC units to patients is complicated by the presence of antigens on the red blood cells. These antigens determine the blood group of every individual. Many different antigens are currently known, but antigens A, B and RhD are the most well known and clinically relevant. When a patient is transfused with blood that contains antigens which the patient lacks, the patient's immune system can form antibodies against the foreign antigen. Depending on which antigen is mismatched, this can have mild to severe consequences. When antibodies are formed against an antigen then every subsequent blood transfusion must not mismatch on this antigen, as otherwise it can lead to an acute or delayed hemolytic reaction (RBC degradation). Current matching policy is to match every patient on their major blood group (determined by antigens A, B and RhD). Because AB-RhD matching is adequate for the majority of patients, they are not typed for the remaining minor antigens. Only specific patient groups require more extensive antigen matching for their blood transfusion(s). This can be for various reasons, but the most common are that the patient needs transfusions on a regular basis or to prevent Hemolytic Disease of the Newborn (HDN) in future pregnancies for women within reproductive age.

Besides the type of matching a patient requires there is another important factor in the matching process, namely time. As most transfusions are elective, there is usually more than enough time to perform the necessary tests to determine the matching requirements for a patient. If no suitable unit is available from the hospital inventory a suitable unit can be ordered and added to the next delivery. Most hospitals receive daily deliveries that take much pressure off the hospital inventory capabilities as there is little to no gain in storing highly compatible units for these situations. Rather, these units are kept at the larger distribution centres such that when they are demanded they can easily be distributed to the hospital where they are needed. For acute transfusions extensive matching is often no option. Many hospital inventories have a small supply of highly compatible O^- blood with some of it also typed negative for antigens c, E and K as these are the most often requested negative units. This blood can be used in rare emergencies where an extensive match is required in an acute transfusion.

The presence of minor antigens in blood can be determined by serological tests, but it is too time consuming and costly to type all donors and patients on every antigen that is clinically relevant. This means that large-scale extensive matching is currently not possible and more importantly, not cost effective. However, this may change in the future. The availability of genotyping technology allows for near instant typing of blood on many antigens. As genotyping chips become more accurate and less costly, a future is foreseeable where large-scale genotyping is realized for all patients and donors. This brings along the possibility that in theory every patient can be matched on many more than the three major antigens. However, there are practical considerations which make such a large-scale extensive antigen matching approach not trivial. Firstly, a more extensive matching policy may lead to a decrease in the availability of highly compatible blood products for the patients who currently already require extensive matching. This is because new patients will be discovered with difficult to match phenotypes which

previously were only matched on antigens A, B and RhD. Another potential risk is that RBC units that express relatively many antigens are more prone to outdating because they become less compatible. RBC units have a maximum shelf life of 35 days, after which they expire. Large-scale extended antigen matching should not lead to an increase in RBC unit outdating. Outdating is to be avoided because it is wasteful and ethically wrong since it concerns donated products. Lastly, extensive matching is further complicated by the fact that not all antigens have equal clinical importance for matching. When identical blood group issuing is not possible, it means that some antigens must be mismatched in order to avoid shortages. As identical matching for all clinically relevant antigens is likely impossible, the problem arises of finding the best compromise which minimizes the number of clinically relevant mismatches. Together, these factors make it no longer obvious which units to allocate to which patients. A manual greedy FIFO issuing strategy, which is current practice in most hospitals, is likely no longer optimal. The goal of this research is therefore to quantify the quality of matchings such that they can be numerically optimized on a daily basis to minimize outdating and antigen mismatches without decreasing the availability of blood products to patients.

The remainder of this work is outlined as follows. In Chapter 2 we give a brief explanation of the relevant blood group systems, antigens, compatibility relations and the concept (relative) immunogenicity. Chapter 3 describes the current practice in the Netherlands concerning antigen matching and inventory management and alludes on previous work done on mathematical optimization of extended antigen matching. In Chapter 4 we provide an overview of existing literature on the blood supply chain and inventory management of blood products. In Chapter 5 we describe how the problem of assigning RBC units to patients can be mathematically modelled. We propose a new Integer Linear Programming (ILP) formulation called MINRAR which minimizes shortages and the risk of antibody formation. Chapter 6 discusses how to extend the MINRAR formulation for use in hospitals or distribution centres. The performance of the resulting MINRAR-Online ILP formulation is compared to previous work using simulation experiments. In Chapter 7 we introduce another ILP formulation to compute the optimal allocation of RBC units to patients over an entire simulation. We then compare these results to the performance of the MINRAR-Online ILP on the same simulation to estimate how close to the theoretical optimal performance the MINRAR-Online ILP is. Finally, in Chapter 8 we investigate how the MINRAR-Online ILP formulation can be used in combination with specific weights for different patient groups to perform large-scale extensive matching while prioritizing certain special patient groups.

2 Antigens and Blood Compatibility

Human blood consists of red blood cells, white blood cells, platelets and plasma. Red blood cells are the most numerous cells in the blood and their main function is to facilitate the oxygen carrying capacity of the blood. White blood cells are larger, less abundant, cells which are part of the immune system. Platelets are the smallest blood components and their principal function is to prevent bleeding. Lastly, there is the plasma, which accounts for about 55% of blood volume. Besides transporting the components mentioned above to all parts of the body, plasma facilitates the transport of nutrients, hormones and proteins to the parts of the body that need it. Furthermore, cells can dispose their waste products by putting them in the plasma. On the surface of the red blood cells are molecules called *antigens*. These are proteins or polysaccharides which serve a variety of functions within the cell membrane. Specific genes in the human DNA encode for the presence or absence of each of these molecules. The exact genetic encoding for all antigens together is called the *genotype* of an individual. The actual presence or absence of the antigens on the red blood cell is called the *phenotype*.

The International Society of Blood Transfusion (ISBT) is a collaboration of transfusion medicine professionals from over 100 countries. Their goal is to improve transfusion safety and this includes creating a consensus on recognized blood group systems and standardizing them. Currently, the ISBT recognizes more than 300 different antigens [1]. The theoretical number of unique blood types is therefore 2^{300} as each antigen is either present or absent. The probability of occurrence of antigens is not strictly independent. Antigens often come in groups called *blood group systems*. The 300 recognized antigens reside in 38 systems. The occurrence of antigens within a system is not independent but between systems there is no dependency.

2.1 Compatibility

The compatibility of blood of a donor and patient is determined by the presence or absence of antigens in the donor's and patient's blood respectively. When an antigen is present in the blood we call this blood *positive* for the antigen. Similarly, if the antigen is absent then the blood is *negative* for the antigen. When a patient receives blood from a given donor there can be four possibilities per antigen as shown in Table 1.

	Patient Negative	Patient Positive
Donor Negative	OK	OK (Substitution)
Donor Positive	Mismatch	OK

Table 1: Donor-Patient compatibility for a certain antigen.

When a patient who is negative for a certain antigen is transfused with donated blood which is positive for that antigen then we call this a *mismatch*. A mismatch may lead to *alloimmunization*, which is when the immune system starts forming antibodies against the supplied antigen, as it is considered foreign. Whether this occurs depends on multiple factors. Some individuals will never form antibodies whereas others will. It is not yet understood why this is the case. Furthermore, the alertness of the immune system also influences whether antibodies will be formed. Lastly, antibody formation is also determined by the *immunogenicity* of the antigen. This is the likelihood that alloimmunization will occur given a mismatch and it is different for each antigen. Usually, alloimmunization does not have direct consequences. However, the immune system has recognized the antigen as a threat and marked it as such. This means that on a possible subsequent transfusion the same mismatch may lead to a *transfusion reaction*.

In case of a transfusion reaction the immune system starts breaking down the transfused blood and this can lead to severe illness of the patient. Due to the high risk of illness, patients with antibodies should always receive blood which lacks the antigen where they have formed antibodies against. Patients alloimmunized against multiple antigens are therefore hard to transfuse as they require very specific RBC units with the absence of the correct antigens. Thus, there is an incentive to prevent alloimmunization, especially for those patients who receive multiple or periodic transfusions.

The upper right term in Table 1 corresponds to a patient who is positive for an antigen receiving a unit which is negative for the antigen. This is called a *substitution*. Substitutions are generally to be avoided, as the patient is not at risk of a mismatch and therefore can safely be transfused a positive unit. On the other hand, a patient who is negative for a particular antigen must be transfused with a unit which is also negative for the antigen to eliminate alloimmunization risk. All combinations in Table 1 which are marked as "OK" do not expose the patient's immune system to a foreign antibody. This means that these three combinations are safe for transfusion and therefore called *compatible*.

2.2 ABO System

The ABO blood group system is by far the most important for determining blood compatibility. It contains antigens A and B, which are the only antigens with the property that if the antigen is absent the immune system will naturally form antibodies against it. This means that every mismatch on either A or B is almost guaranteed to cause a severe transfusion reaction. Since both antigens can be either present or absent there are $2^2 = 4$ different phenotypes: O, A, B and AB. Here O denotes the absence of both antigens. Because mismatches only occur when blood positive for an antigen is transfused to a patient who is negative for that antigen it is obvious that O blood is the most usable, being compatible with all ABO blood types. Similarly, AB is the least compatible blood as it can only be issued to patients who also have the AB phenotype. Patients with the AB phenotype can however receive from all other phenotypes. Thus O is the universal compatible donor while AB is the universal receiver.

2.3 Rhesus D

The next most important antigen is RhD which is part of the Rhesus (Rh) system. The Rhesus system contains five antigens in total but RhD is the most immunogenic. It is not strictly as important as the A and B antigens but important enough that it should be matched if possible for all transfusions. Because correct matching on antigens A and B as well as RhD is sufficient to prevent direct transfusion reactions they together form the major blood group of an individual. The presence of the RhD antigen is often indicated with a + symbol and the absence is denoted with - and the blood is accordingly called Rhesus-positive or Rhesus-negative. Thus, the $2^3 = 8$ major blood groups are: $O^-, O^+, A^-, A^+, B^-, B^+, AB^-, AB^+$. These eight major blood groups and the corresponding antigens are shown in Table 2. The prevalence of each blood group shows the percentage of occurrence of this blood group among Caucasians.

Name	Antigen A	Antigen B	RhD	Prevalence
O^-	-	-	-	7.2%
O^+	-	-	+	35.8%
A^-	+	-	-	7.4%
A^+	+	-	+	36.6%
B^-	-	+	-	1.5%
B^+	-	+	+	7.5%
AB^-	+	+	-	0.7%
AB^+	+	+	+	3.3%

Table 2: Eight major blood types with their prevalence among Caucasians

RhD negative blood (-) can be transfused to both rhesus negative and positive patients. Rhesus positive blood can only be transfused to positive patients. The compatibility relations can be visualized with a *compatibility matrix*. For each donor-patient major blood group combination the matrix specifies whether the combination is compatible. In Table 3 the compatibility relations are shown for the eight major blood groups.

Patient→ Donor↓	O^-	O^+	A^-	A^+	B^-	B^+	AB^-	AB^+
O^-	✓	✓	✓	✓	✓	✓	✓	✓
O^+		✓		✓		✓		✓
A^-			✓	✓			✓	✓
A^+				✓				✓
B^-					✓	✓	✓	✓
B^+						✓		✓
AB^-							✓	✓
AB^+								✓

Table 3: Donor-Patient ABOD-compatibility for the eight major blood types.

In practice all patients who require a transfusion are matched with donor blood that is of a compatible major blood type. We call this *AB-RhD compatible* matching. Only in very rare instances or emergencies a non-AB-RhD compatible transfusion can be allowed. This is however beyond the scope of this research and therefore we will only consider AB-RhD compatible matches.

2.4 Other Systems

The major blood type of a patient is determined by the *A* and *B* antigens from the *ABO* system and the *RhD* antigen from the Rhesus system. There are however many other antigens present in human blood. Under normal circumstances they are not significant enough that a mismatch on one of these makes an otherwise compatible match incompatible. However, each mismatch on any of these antigens has an associated risk. As mentioned earlier there are more than 300 antigens currently recognized by the ISBT. We will limit our view to only those antigens that are considered clinically relevant according to the CBO transfusion guidelines [2]. Apart from the ABO and Rhesus systems this means we will include Kell, MNS, Duffy and Kidd. An overview is shown in Table 4 below.

Blood Group System	ABO		Rhesus					Kell		Duffy		Kidd		MNS			
Antigen	A	B	D	C	c	E	e	K	k	Fy(a)	Fy(b)	Jk(a)	Jk(b)	M	N	S	s
General Reference	Major		Minor														

Table 4: Blood group systems considered and their antigens.

As mentioned before, there is no dependency between systems in terms of antigen expression. For example, there is no correlation between expressing antigen C from the Rhesus system and Antigen M from then MNS system. However, within a system there is dependency. Some combinations within a system occur frequently while others do not exist. These probabilities also vary between individuals from different ethnicity. The frequency of occurrence of the different antigen combinations within each blood group are shown in Appendix A.

2.5 Relative Immunogenicity

Not all antigens have the same urgency for matching. A large factor in determining the matching priorities between the antigens is their *relative immunogenicity*. The immunogenicity of an antigen expresses the likelihood that a mismatch leads to alloimmunization. As the immunogenicity of most antigens is low, it is more interesting to look at the relative immunogenicity of antigens. This relative immunogenicity was estimated in a 2016 study by Evers *et al.* [3]. Transfusion data was gathered from six hospitals totalling a group of 54347 transfused patients. These were all patients who had not previously been transfused. Patients who received extensively matched products (more than just AB-RhD compatible) were excluded, because in order to estimate the risk of antibody formation mismatches should be possible such that antibody formation could occur. Furthermore, all patients were excluded who received no follow up transfusion(s), because for this group there was no data on whether they formed antibodies. Finally, patients were excluded for which the reason of antibody formation could have come from a different cause than the transfusion. The result is a group of 21512 patients who received one or more transfusions and were screened for antibody formation afterwards. The authors investigated antibody formation against antigens in the following blood systems: Rhesus, Kell, Duffy, Kidd, MNS, Lewis and Lutheran. Because of the lower clinical relevance of the latter two blood group systems, we will omit their results. Similarly, the authors did not include any results for antigen k, as nearly the whole Caucasian population (99.8%) is positive for k, making mismatches extremely unlikely and antibody formation even more so.

The authors first estimated an *antigen negative cohort* per antigen. This is the theoretical subgroup of patients which is expected to lack the antigen in their phenotype. As the patients had not been extensively phenotyped before transfusion, it was not known which patients expressed which antigens. Therefore, the size of the antigen negative cohorts is not known and must be estimated based on the prevalence of the corresponding antigens. However, due to the large number of participants these estimates are likely to correspond to the true number of antigen negative individuals in the patient group. Furthermore, because all patients were not matched on any antigens other than A, B and RhD, it is unlikely that there is a bias for patients with negative antigens to receive more RBC units which were positive for those antigens which would result in skewed estimates of the relative immunogenicity. Therefore, dividing the number of patients who were alloimmunized for a certain antigen by the size of negative cohort for that antigen is a unbiased indicator of the relative immunogenicity. The size of the antigen

negative cohort per antigen is shown in Table 5.

In total 474 times antibodies were formed against any of these antigens: C, c, E, e, K, $Fy(a)$, $Fy(b)$, $Jk(a)$, $Jk(b)$, M, S, s. As all units transfused were extensively typed, the authors were able to study the relation between the number of units transfused and the antibody formation. First they calculated the *cumulative alloimmunization incidence* which specifies per antigen the percentage of the antigen negative cohort that formed antibodies when transfused with a specific number of antigen positive units. This measure gives an indication of the immunogenicity of the antigen because it is not dependent on the antigen prevalence in the population.

The authors compute the cumulative alloimmunization incidence per antigen after transfusion with 1, 2, 5, 10, 15 and 20 antigen positive units. We will use the incidence levels for transfusion with two units to estimate the relative immunogenicity. This is done for the following reasons: Firstly, transfusion with two units is the most common and therefore the most representative of the alloimmunization risk after one transfusion episode. Secondly, the cumulative incidences of all antigens increase when the cumulative number of mismatching units transfused is increased. However, the ratio of these values between antigens remains relatively constant, implying that the immunogenicity does not heavily depend on the number of mismatching units transfused. Lastly, we have chosen to use the incidence values after two antigen negative units transfused instead of after just one unit. The reason for this is that after transfusion with only a single antigen negative unit there was no alloimmunization observed for several antigens (e, $Fy(b)$ and $Jk(b)$), implying that the relative immunogenicity of these antigens would be zero. It is not clear if these antigens require more than one unit of exposure to lead to alloimmunization or there were just no cases in the studied patient cohort because the absolute probability of antibody formation is low.

All in all, we think that the alloimmunization incidence values for exposure to two units is most representative for the immunogenicity values of the corresponding antigens. Whether the immune system forms antibodies when exposed to a foreign antigen is a complex process which is not yet completely understood. We think that using the values as mentioned is a simple yet effective way to estimate the relative risk of antibody formation given a mismatch.

As mentioned earlier, the absolute risk of alloimmunization for these antigens is low. When computing the relative immunogenicity we will therefore normalize these values to sum up to one. Besides being more user friendly, it will make optimization slightly easier as (very) small values are numerically difficult to handle. The alloimmunization incidence levels and corresponding relative immunogenicity values are shown in Table 5. The percentages of patients from the antigen negative cohorts which formed antibodies were directly copied from [3]. We also show the size of the antigen negative cohort to give an indication of the size of the group of patients at risk of antibody formation. Evers *et al.* [3] do not present the actual number of patients who were transfused with two antigen positive RBC units per antigen, and therefore this data is not shown.

Antigen	Size of Antigen Negative Cohort (Total = 21512 Patients)	Percentage of Negative Cohort that formed antibodies after transfusion with 2 antigen-positive RBC units	Estimated Relative Immunogenicity
C	6758	0.21%	3.45%
c	4247	0.43%	7.06%
E	15122	1.46%	23.97%
e	433	0.51%	8.37%
K	19274	2.34%	38.42%
Fya	7181	0.27%	4.43%
Fyb	3583	0.08%	1.31%
Jka	4890	0.51%	8.37%
Jkb	5482	0.02%	0.33%
M	4648	0.18%	2.96%
N	5895	0%	0.00%
S	9479	0.08%	1.31%
s	2317	0%	0.00%

Table 5: Relative Immunogenicity per antigen. This measure is estimated by Evers *et al.* [3] by comparing the percentage of patients in the antigen negative cohorts that formed antibodies when transfused with two antigen positive RBC units. As alloimmunization incidence rates are low, these percentages are normalized to sum up to 100% to estimate the relative immunogenicity per antigen.

3 Problem Description

In this chapter we will describe the current blood supply chain in the Netherlands. Furthermore, we will elaborate on current antigen matching practices and the priorities in matching RBC units to patients. We will also discuss how the availability of genotyping opens up future possibilities of large-scale extensive matching and which problems come with it. We summarize the previous work on mathematical optimization of large-scale extensive matching and the shortcomings it has. We then identify which priorities should be adhered to in a mathematical optimization approach to increase practical relevance.

3.1 Current Situation

3.1.1 Supply Chain

The Dutch blood bank Sanquin has a donor population of more than 300,000 donors and processes more than 400,000 whole blood donations per year [4]. These donors receive regular invitations for donation. Because blood donation is voluntarily, the donors are not obliged to respond to an invitation. It is the responsibility of Sanquin to ensure that enough donors are invited to maintain a sufficiently large stock to be able to cope with variations in both supply and demand. Blood is donated in either one of the fifty fixed donation centres or in one of the eighty mobile donation units. This blood is then brought to one of the two processing facilities where the blood is checked for any irregularities and processed into 300ml units for further distribution. From these two facilities, one in Amsterdam and the other in Nijmegen, the blood units are distributed over seven distribution centres in the Netherlands which in turn supply the 100 Dutch hospitals. To ensure that variations in supply or demand can be dealt with, the blood bank maintains a buffer of on average five days of demand. The more infrequent blood types such as the AB groups are likely to be somewhat longer than five days in the distribution centres whereas the higher throughput groups usually spent less days before distribution to the hospitals.

All Dutch blood donors are typed for their AB-RhD major blood group. Furthermore, Sanquin has also performed more extensive typing on a smaller number of donors. For the 22 antigens most often requested, the blood bank has set up a procedure to perform large-scale typing. For each of these antigens a target level is set which states the fraction of O and A donors who should be typed for this antigen to meet the demand for typed units [5].

Sanquin has arrangements with each hospital on order-up-to levels for each of the eight major blood groups. These specify minimum and maximum inventory levels. The minimum levels indicate when to order and the maximum levels are used to determine how much of each blood group should be ordered. These maximum levels are also known as order-up-to levels. The order-up-to levels are based on the size of the hospital, the patient cohort and the interval at which the hospital receives deliveries. Almost all hospitals have at least one daily delivery, with larger hospitals having the option to receive multiple deliveries per day. Besides an amount for each major blood group, hospitals can also order more specific products from Sanquin. This includes specific requests for certain blood units that lack certain (combinations of) antigens.

When a hospital receives a delivery, all units are scanned into an inventory management system. All new units are also checked again for their AB-RhD type, as a way to make sure that the typing on the bag is correct. Next, the units are stored in a cooling cabinet that is subdivided into a free inventory and a reserved inventory. Units that have entered the hospital inventory will stay there until they are transfused or expire. Alternatively, they can also be transported to be used for surgery but remain unused. Then they are returned to the unassigned inventory and available for re-use.

3.1.2 Antigen Matching

Currently the matching of RBC units to patients in the Netherlands is done according to the CBO blood transfusion guidelines [2]. These prescribe which patients should receive additional matching on top of the regular AB-RhD compatible issuing. The first step in the matching process is the initial *Type and Screen* procedure. A Type and Screen procedure is intended to obtain information about the patient's blood. It consists of two steps: typing and screening. In the typing procedure the major blood group of the patient is determined. The screening refers to a combined series of tests in which the blood of the patient is screened for the presence of irregular antibodies, that is, antibodies other than against A, B or RhD. Most likely these antibodies have been formed at a prior transfusion where a mismatch occurred on a minor antigen. When the screening test result is positive, it means that one or more irregular antibodies are found. The blood is then sent to be examined at the Sanquin lab where the exact antibody is determined. When an antibody is present, the patient should receive blood that lacks the corresponding antigen, because if not a transfusion reaction is very likely. The first Type and Screen test is usually performed far ahead of the actual transfusion. This can range from days to weeks or months.

A second Type and Screen test is done shortly before the transfusion. This test is required to validate the results of the first test. After this second screening the definitive requirements are known for the patient and thus a unit can be selected from inventory. CBO guidelines state that the results of a Type and Screen test are valid for at most 72 hours. This is a precautionary measure to minimize the risk of antibodies being formed after the Type and Screen procedure. After the definitive Type and Screen test, patients can roughly be classified into two groups. On the one hand there is the majority of patients with a negative antibody screening and who therefore do not require additional matching. They do not have irregular antibodies and there is no direct incentive to perform extended matching, although it would still be beneficial as additional antibody formation would be prevented. The remainder of patients have some additional matching requirement. Apart from the presence of antibodies there can be various other reasons for a patient to require additional matching. The most common reasons are listed below.

- There is an incentive to prevent antibody formation for this patient. This can have various reasons but the most common one is that it concerns a patient who requires multiple or chronic transfusions. Every time antibodies are formed subsequent transfusions are complicated with an additional matching requirement. Common patient groups who require extensive matching are patients with sickle cell disease, thalassemia, myelodysplastic syndromes or autoimmune hemolytic anemia.

- The patient requires additional matching because of another reason. For example, women within reproductive age (< 45) require additional matching because of possible complications during future pregnancies.

Current matching policy prescribes that every patient must receive AB-RhD compatible blood. Normal patients that are not expected to receive multiple transfusions are not matched on any additional antigens. This is mainly done for three reasons: 1) There is a very low risk that mismatching minor antigens in the transfusion will lead to any complications for the patient. 2) Comprehensively typing these patients is labour intensive and therefore costly. 3) Only a fraction of all RBC-units is comprehensively typed and therefore it is better to save these units for patients who actually need a comprehensively matched product. For this patient group there is usually a set of antigens that are considered a 'must', for instance the antigens for which the patient has known

antibodies in their blood. The remaining antigens do not have to be correctly matched but this can be recommended.

3.1.3 Multiple Units

All RBC units have a fixed volume of 300ml. It may be obvious that patients who require a transfusion do not have a demand limited to only 300ml. Many patients require more than one unit. Most common demand is one or two units. Demand for more units does occur but is relatively uncommon. An empirical distribution of demand is shown below.

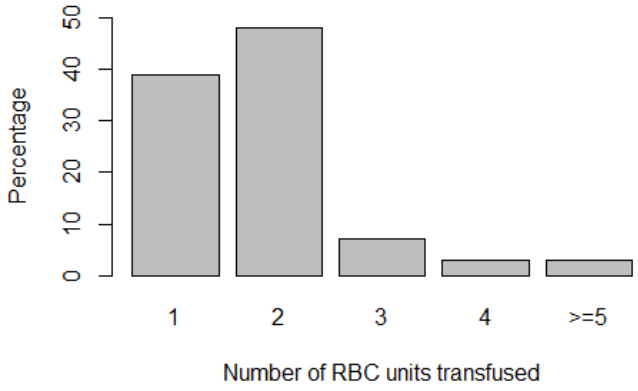

Figure 1: Distribution of number of units transfused per transfusion episode.

All transfusion episodes (=one a period of one or more transfusions) in which a patient has been transfused with five or more units have been aggregated in one column. This is done because it is of little relevance to distinguish demand for more than four units. The reason for this is that the dataset from which this data is aggregated contains transfusion episodes with a very large number of transfused units. Some cases have transfusion episodes in which more than 100 RBC units are transfused to a single patient. These high numbers do not concern planned transfusions but (extremely) severe traumas and are therefore not within the scope of this study. We have consulted with transfusion lab staff which demand to exclude from this study and concluded that nearly all non-emergency demand is limited to at most four units. Therefore we will include only this demand in our study.

3.2 Future

Current matching regulations take great care in ensuring compatible extensive matching for patients for whom extensive matching is essential. The hospitals cooperate with Sanquin to supply every patient with suitable blood whilst making sure that overall the fulfillment rate, that is the percentage of requests that can be met with blood that satisfies the requirements, is kept close to 100%. Although current practice is already of high quality, it does not mean that it cannot be improved. Advances in technology open up possibilities previously thought impossible. One of these new technologies is large-scale genotyping.

3.2.1 Genotyping

While a portion of the donor population is more extensively typed than AB-RhD, this is not the case for ordinary patients. Donors regularly give blood, so performing extensive typing is an investment with a lasting benefit. Once a donor has been typed for an antigen, every RBC-unit donated by this donor is also marked with the presence or absence of this antigen. For patients the typing is not nearly as cost effective. Apart from periodically transfused patients, the majority of patients receives only one transfusion. Serological typing, that is the determination of the presence of every antigen with an individual test, is a time consuming and costly process.

Recently, genotyping technology has made it much easier to determine the full patient's phenotype. Using molecular DNA sequencing the presence of dozens of antigens can nearly instantly be determined. Note that this is different from the screening procedure in the Type & Screen test in which the blood is screened for irregular antibodies. The presence of antibodies cannot be obtained from the DNA as antibodies are a result of an immune system response. Genotyping therefore is not a replacement for a Type and Screen procedure, but instead a replacement for many individual serological tests.

3.2.2 Implications

When both donors and patients can be extensively typed on a large-scale it may be possible to perform better antigen matching for many more patients. However, the matching process does become more difficult. Where most blood is now typed on three antigens with $2^3 = 8$ different blood groups, a modest amount of 17 antigens already produces $2^{17} = 131072$ theoretical blood groups. Such a large number of different products requires a more intelligent way of inventory management and product issuing.

3.3 Mathematical Optimization

In case large-scale genotyping would be introduced it is not trivial anymore to select the most suitable products for patients while ensuring that the fulfilment rate does not decrease. Therefore, this problem calls for a more mathematical approach.

3.3.1 Previous Work

Because large-scale genotyping for both donors and patients is not yet implemented there has not been much research on effective and efficient matching procedures. Initial work on this subject was done by Van Sambeeck *et al.* [6], [7]. A mathematical framework was created to model the daily matching of RBC-units to patients, assuming known phenotypes of both RBC-units and patients. Using historical transfusion data and antigen prevalence, the supply of donor blood and patient requests are modelled. The daily assignment problem that arises is solved by transforming it to a Minimum Cost Maximum Flow (MCMF) problem which can be efficiently solved. The compatible combinations of RBC unit and patient are determined by first limiting the view to a (sub)set of antigens and then only allowing matches that are compatible on all those antigens. To compute the best possible matching a weight is assigned to each allowed match which is composed of two factors: *remaining shelf life* and *relative opportunity loss*. The remaining shelf life factor penalizes the issuing of relatively fresh units, making sure to not let older units outdate. The relative opportunity loss quantifies how well an RBC-unit matches the patient's requests. Because only compatible matches are allowed, the RBC-unit always expresses the same negative antigens as those requested. If the RBC-unit has more negative antigens, it is more compatible than necessary. This therefore implies a higher relative opportunity loss, as the unit might be better used to

satisfy a more extensive request. The constructed MCMF model, called FIFO/MROL after the two terms that make up the objective function, minimizes both outdating and relative opportunity loss.

A simulation study is performed to test the performance of the issuing strategy over a longer period of time. The results show that the outdating and shortage percentages can be kept low (0.8% and 1.3% respectively) with compatible matching on 14 minor antigens when maintaining one week of national demand worth of inventory (average daily demand = 1184 RBC units).

3.3.2 Our Contribution

While previously mentioned work is the first to investigate the availability of RBC-units under large-scale extensive matching, there are some aspects that are not taken into account. Our work aims to continue the same mathematical approach while adding some modifications to the model to be able to get a better estimate of the possibilities of extended matching. To get an adequate understanding of the important factors that play a role in the matching process, we discussed various issues with the Sanquin supply chain manager, a hospital's clinical chemist and immunology experts. Each of these individuals appreciates the matching process from a different perspective. From these conversations we have learned the following:

- Extensive matching for patients who require extensively typed blood products (in one way or another) is far more important than it is for patients who do not.
- Large-scale extensive matching should therefore not decrease the availability of extensively typed antigen-negative blood for those patients who actually need it.
- The majority of patient requests are for more than one unit which complicates the matching process.
- There is little pressure on the hospital inventories due to the high supply frequency and extensive lifespan of RBCs.
- Consequently, shortages rarely occur and a good model should therefore also have near-zero shortages.

These factors shed a different light on the priorities in a large-scale extensive matching effort. Besides these non-technical factors, we noted some shortcomings of the matching model used in the previous work. This matching model needs a predetermined set of antigens to consider and will subsequently compute a matching that only accepts strictly compatible matches on these antigens. By enforcing strict compatibility, the number of possible matches is unnecessarily reduced. When unable to cope with all requests, one has to remove one antigen from the entire matching and try again. Instead, we propose a different approach where all AB-RhD compatible matches are allowed and the compatibility of the remaining antigens is converted into a cost based on relative immunogenicity. In this way, we do not eliminate any solutions but instead allow solutions which the previous model does not consider. This model also explicitly allows for requests for multiple units, both when minimizing shortages and alloimmunization risk. We also show that our model can be used to prioritize patient groups which currently already receive extensively matched RBC units, which makes it more relevant in practice.

4 Related Literature

The blood supply chain is a complex system with different processes and decisions involved from donation to transfusion. We will give an overview of the literature addressing optimization problems which occur within the supply chain. We will focus mostly on literature concerning inventory management problems, as they are most closely related to the problem at hand.

We will look at hospital inventory management and categorize the literature on the following variants:

- Order-up-to-levels
- Age Categorization in Demand
- Inventory Age Distribution
- Compatible Matching
- Reducing Shelf Life
- Centralized Inventory Management

Most optimization work on inventory management combines one or more of these aspects in their approach. Although all this work concerns inventory management, only few studies have investigated inventory management under large-scale extensive antigen matching. This means that in almost all previous work antigen compatibility concerned only antigens A, B and RhD and therefore the methods used in most previous work cannot directly be used to solve the more combinatorially challenging problem of extended antigen matching. Lastly, we briefly summarize some work on the different genotyping methods available, as the possibility of large-scale genotyping is the main assumption in this work.

4.1 Hospital Inventory

Inventory management is the most widely studied aspect of the blood supply chain. It considers the daily decisions that have to be taken in a hospital blood bank inventory both in terms of ordering and allocation of blood to requests. Generally, two objectives rule these decisions. Firstly, shortages should be minimized. A shortage is the failure to satisfy a request for blood and can have severe consequences. An emergency order might have to be placed or an operation might have to be cancelled. In the most extreme case a shortage might lead to the death of a patient. It may be obvious that these consequences are hardly expressible as costs. Therefore, shortages are generally always to be prevented by having a large enough stock. Maintaining a large stock can possibly lead to the outdating of some blood products. Depending on the type of product the maximum shelf life differs, but for red blood cells the legal shelf life in the Netherlands is 35 days. If a unit is not used within this period it cannot be used for transfusion anymore and is destroyed or sent back to the production facility for possible extraction of useful parts or research. Since blood donations are voluntarily, having a high outdate percentage is ethically unacceptable. It is easy to see that maintaining a larger inventory will decrease shortages but increase outdates, whereas smaller inventories will prevent outdating but increase shortages. The conflicting nature of these objectives is part of what makes inventory management difficult. The real difficulty however, lies in the stochasticity of the problem. The exact demand is generally not known beforehand. Some transfusions are scheduled far in advance, but the majority is not known until a few days before or, in case of emergency, not until the day of transfusion. This makes it complex to maintain

inventory levels that can supply all requests whilst still preventing outdating. Supply is also stochastic since donors are invited to donate in a span of a few weeks but when they come in for donation is completely up to them. Furthermore, the donors are not obliged to respond to an invitation. Although this stochastic characteristic in supply is certainly challenging for the whole blood supply chain, it is not of direct relevance for hospitals. For the scope of hospital inventory management, we can assume that placed orders at AB-RhD level are always fulfilled.

4.1.1 Order-up-to Levels

Inventory management of perishable products has been widely researched. Both red blood cells and platelets in blood banks have been studied, but also perishables in supermarkets such as fruits have similar problems. Most research dealing with the inventory management of perishable products focuses on determining suitable *order-up-to levels* [8]–[16]. An order-up-to level specifies the target size of the stock at the beginning of a certain time interval, called the *review period*. Typically for blood products the review period is one day. In that case, order-up-to levels indicate the size of stock per type of product at the start of each day. At the end of each review period an order can be placed for each type of product to fill the stock back up to the order-up-to level. The time between the placement of an order and its delivery is called the *lead time*. Depending on the definition, in systems dealing with blood products the lead time is one or zero days. Adequate order-up-to levels should minimize the probability of a shortage occurring during the review period and lead time while also not maintaining such a large stock that causes items to outdate. Order-up-to levels are widely used in practice because of their simplicity and effectiveness. They do not involve complicated decisions for the inventory manager. In the simplest case there is one static order-up-to level per product, but alternatively a different level can be computed for each day of the week if demand is more periodic.

Order-up-to levels are widely studied in the recent blood supply chain literature. Different methods have been used to determine order-up-to levels such as Markov Decision Process (MDP) [8], [9], [11], [12], (Stochastic) Mixed Integer Linear Programming [13], [15]–[20], Meta Heuristics [12], [14] and Simulation [8], [10], [21]–[24]. First work on analysing blood blank inventory management was done in the 70s and 80s of the previous century. Much of this early research focuses on analytical determination of suitable inventory levels. Prastacos [25] provides a good review. In more recent years, research has shifted to more sophisticated models to capture more of the complexity of the inventory management process with regard to blood products. Most of this work considers either red blood cells or platelets. Although these blood products have different uses and maximum shelf lives, we will not separate the literature regarding these products since many techniques are applicable to both platelets and RBCs. In perishable product inventory management, a first in first out (FIFO) policy is the standard. It has a natural tendency to prevent outdating. Most of the literature indeed does follow this allocation policy, but deviations are also investigated. Atkinson *et al.* [21] propose using a LIFO policy for all RBC units older than a specified number of days d while issuing younger units as FIFO. They note that a hospital usually receives slightly more blood than is demanded, which makes some outdating inevitable. But the age of transfused blood is not as fixed as the outdating percentage. Some research suggests that transfusing younger blood might lead to lowered mortality rates [26]–[28]. Atkinson *et al.* [21] find that they can lower the mean age of transfused blood with 10 to 20 days with a corresponding shortage increases of at most 0.5% when using a LIFO policy for old units. Simonetti *et al.* [22] compare FIFO to two other policies: *likely newest* and *likely oldest*. Both methods use a negative binomial distribution to assign probabilities

to the issuing of inventory RBC units of different ages, skewed to newer or older units respectively. Their results are in line with those of Atkinson *et al.* [21] as they find that the number of outdates is more or less invariant to the allocation policy. This is to be expected when both supply and demand are relatively constant and shortages are rare. Both non-FIFO policies do cause the steady state inventory to decrease which is not necessarily positive as a smaller inventory is not prepared to high spikes in demand in for example emergency situations. Abdulwahab and Wahab [11] vary between FIFO, LIFO and a Circular allocation policy for a platelet bank. A FIFO policy is shown to perform the best. Their model also allows deviating from the policy by assigning younger compatible units if this is beneficial in terms of costs.

4.1.2 Age Categorization in Demand

The demand for blood products cannot always be satisfied by a standard FIFO issuing policy. This is because blood products are used for many different treatments and some of these require the freshest blood products. Haijema *et al.* [8] explicitly model this in their MDP by considering two types of demand for platelets: *young* and *any*. Platelets have a shelf life of approximately one week. Demand for young platelets is for units of at most two days old. Demand for any is for treatments that do not explicitly need young platelets or for emergency use. Because of the two types of demand the order-up-to rules are also extended for both types of demand and therefore called "2d-order-up-to levels". Since the shelf life of platelets equals a week, the authors find 5-day periodic order-up-to levels. For each day of the week these 2D order-up-to levels specify the ideal number of young units and the ideal number of total units in stock. Furthermore, the model is extended for irregularities in production during breaks such as Christmas and Easter [9]. Civelek *et al.* [12] extend the trend of differentiating between demand for products of different ages to three categories: *young*, *mature* and *old*. Their model allows substitution of requests for platelets in different categories. Due to this possibility the number of young units in inventory is endangered because young units can be substituted for any type of request. Therefore, a heuristic allocation policy is used with protection levels for this category, thus preventing disproportional substitution of young platelets. For RBC units the distinction between requests of different ages is also investigated. In their stochastic integer programming model Gunpinar and Centeno [18] model requests for types *young* and *any* for both platelets and RBC units. Their model has explicit capacity constraints on both types of products and both types of ages. Especially the capacity for young products heavily influences the quality of the solution. Increasing this capacity reduces the shortage rate but increases overall costs.

4.1.3 Inventory Age Distribution

Categorizing inventory units based on their shelf life is not only useful for satisfying categorized requests. Regular order-up-to levels or even 2D order up-to-levels [8] ignore the age-distribution of the current stock, which is not always optimal [29]. Duan and Liao [10] pose a solution to this problem as they look at platelet inventory management with special attention for the age distribution of the stock. They make a distinction between *normal* units and *old* units in inventory when placing an order. With these they compute the *Old Inventory Ratio (OIR)*, which is the fraction of old units in the total inventory. If this fraction exceeds parameter δ then additional platelets are ordered to anticipate possible future shortages. Pauls-Worm *et al.* [16] consider the planning problem for a single perishable product under erratic demand. Their Mixed Integer Linear Programming Model computes suitable review periods and order-up-to levels which account for the expected number of outdates given the age distribution

of the current stock. The model is made deterministic by using chance constraints to satisfy a certain target service level.

4.1.4 Compatible Matching

When optimizing RBC inventory management, the substitution relationships between the major blood groups cannot be ignored. Allowing compatible substitutes when satisfying requests can help to mitigate shortages. Saving more applicable blood like O^- whenever possible gives more future flexibility. Allowing a model to substitute compatible blood for requests adds extra options which might help the performance under variation in demand. Modelling these substitutions also reflects common practice in hospitals. Duan and Liao [14] study the performance of a single-hospital single-blood centre system with no AB-RhD substitution, substitution only at the hospital and system-wide substitution. The model uses the OIR [10] for determining order-up-to levels. The model is solved using a meta heuristic based on Tabu Search and Simulated Annealing. Results show that allowing AB-RhD substitution can reduce system-wide outdates with 16% when the maximum shelf life is one week. Dillon *et al.* [15] have created a stochastic integer programming model for determining the optimal order-up-to ordering policy. The authors test their model with and without AB-RhD substitution with interesting results. They present the order-up-to levels for each major blood group with and without substitution. With substitution the target inventory levels for all blood groups decrease, except for O^-. Furthermore, substitution lowers the average age of transfused blood. Najafi *et al.* [17] also allow AB-RhD substitution in the matching. As mentioned before their model differentiates between young and any age demand and supply. A table is presented which shows the occurrence of different substitutions between blood groups. Abdulwahab and Wahab [11] consider the substitutions among platelets when satisfying demand. Their FIFO model assigns fixed costs to different kinds of substitution. Identical matches have the highest reward, compatible non-O^- blood slightly less and the substitution of O^- blood is valued the least. Furthermore, a penalty linear in the number of days is used to prevent substituting young blood. The authors find that substitution improves performance, but the rewards assigned to the substitutions are more or less arbitrary. The authors also investigate the ideal percentage of O^- blood in inventory. Graphs show that both shortage and outdates reduce sharply when increased to 25% and then slightly until this percentage reaches 40%. More theoretical work has also been done on calculating steady state levels given the possible substitutions between blood groups based on differential equations [30]. Unfortunately this work does not consider the Rhesus factor, rendering it not particularly useful. Several studies have approached the problem of assigning compatible RBC units as a multiple knapsack problem [31]–[34]. The demand for each blood type is modelled as an individual knapsack. All inventory units have size one and can have a value equal to their receivability factor. This means that O^- blood has low value and is therefore not preferable to satisfy demand. The authors then compute an optimal allocation of RBC units while allowing for extra ordering of units, albeit at a cost. Results show that the MKP formulation reduces the overall number RBC units imported compared to identical matching. However, the number of O^- RBC units imported does increase because where an identical matching procedure would import if a type is unavailable, the compatible matching procedure will first use compatible available units first. Since O^- is always compatible, it will itself inevitably run short. All implementations use a FIFO policy for minimizing outdated units, which proves to be effective.

4.1.5 Reducing Shelf Life

Some studies suggest that transfusions with blood with increased shelf life is correlated with possible complications for the patient, as well as reduced short-term and long-term survival [35]. Therefore, Fontaine *et al.* [23] conduct a simulation study using actual transfusion data in which they analyze different scenarios with varying maximum shelf lives under a FIFO issuing policy. Results show that shortages increase with 51%, 20%, 10%, 4%, and 1% for maximum shelf lives of 7, 14, 21, 28, and 35 days respectively. Outdates are much less affected. Grasas *et al.* [24] also investigated the reduction and concluded that indeed the maximum shelf life can be reduced to 28 days without serious complications. They also tested the robustness of this by simulating supply shocks. They found that for maximum shelf lives of 21 days and less the inventory capability to compensate these shocks is severely reduced. Furthermore, the Rhesus negative blood stock is more affected due to its lower prevalence. Using their meta heuristic, Duan and Liao [14] show that AB-RhD substitution can help reduce outdating when the maximum shelf life for RBC units is limited. In their study they test with a maximum shelf life of one, two and three weeks. Outdates heavily reduce from one to two weeks while extending it to three weeks only has a very small improvement. This means that two weeks of effective maximum shelf life is essentially needed for RBC units to have both shortage and outdate rates low, while extensions do not significantly improve the quality. Therefore, they conclude that average issuing ages larger than two weeks can and should be prevented.

4.1.6 Centralized Inventory Management

Most literature only considers one hospital blood bank or a single-hospital single-blood bank system. However, system-wide shortages and outdates can be mitigated by co-operation between hospitals [10], [17], [36], [37]. Kendall and Lee [37] create a model that performs weekly rotation of units between regional hospitals. They make use of a goal programming approach which allows an administrator to set different performance goals with different priorities. Results show that allowing rotation leads to less outdating, however no costs were incurred for the systematic rotation, leaving the extra costs of such a system unclear. Sapountzis [20], [36] instead looks at a central inventory for a group of hospitals. Their model assumes that for each individual hospital and blood group the probability is known of an i days old unit eventually expiring. These probabilities form a curve per blood type per hospital. The model then distributes units of different groups and age optimally among the hospitals as to minimize the overall probability of expiration. This is calculated by an ILP which is proven to be unimodular. The expiration probability curves are estimated by known data points consisting of a known number of units of remaining shelf life j and the number that eventually expire. A curve is fitted that has the property that the probability of a unit being used is independent of its shelf life, except for the final days. Duan and Liao [10] compare platelet inventory management for a central inventory versus a decentralized inventory. Results show that system-wide outdates can be reduced from 19.6% to 1.04% using a centralized approach. Maintaining centralized control over multiple inventories can be beneficial in terms of shortages and outdates. In practice however, implementing such a system may prove difficult.

4.2 Genotyping and Antigen Matching

4.2.1 Genotyping Methods

4.2.1.1 Molecular Background Since the discovery of the human genome researchers have uncovered much of the molecular basis that is responsible for the antigens present on human red blood cells. The presence of many antigens is determined by a single nucleotide polymorphism (SNP) or one allele, allowing a relatively straightforward phenotype prediction from the DNA. Some antigen systems however are more complex. Most notably are the ABO and Rhesus systems for which the determination of the phenotype is challenging due to the large genetic complexity and many variant alleles [38]. Mistyping a donor or patient on their ABO blood group can have severe consequences. Therefore, it is unlikely that genotyping will replace the serological ABO blood group typing since the serological methods are fast, cheap and reliable. For most populations the molecular typing of the Rhesus antigens (D, C, c, E, e) is not too difficult. The genetic structure of the Rhesus system is however prone to many variations which lead to more complex and hybrid variants. These variants are more likely to incorrect serological typing and consequent alloimmunization. Rhesus variation is uncommon in Caucasian population but more frequent with individuals of African descent which is also a group more prone to Sickle Cell Disease (SCD) and thus to receiving more blood transfusions. Rhesus variants have been shown to contribute to alloimmunization among SCD patients [39]. The application of genotyping should facilitate a more exact blood type determination. A combined use with reliable serological tests is likely the most promising approach as serological tests can give results for blood systems which are more difficult to genotype. On the other hand, genotyping is much more effective in typing many antigens whose presence is generally hard to show serologically.

4.2.1.2 Commercial Methods Several commercial systems have been developed based on DNA microarray technology to determine a wide range of antigens and variants. Microarrays are used to perform simultaneous detection of many thousands of genes. With the use of specific software this information is converted into a patient's phenotype which expresses the presence of a predetermined number of antigens and/or mutations. Examples are the HEA BeadChip (Immucor, Warren, New Jersey,USA) [40] and BloodChip (Progenika Biopharma, Derio, Spain) [41].

4.2.2 Genotyping Applications

Genotyping may be performed for multiple reasons. According to Sapatnekar and Figueroa [42] the following patients are appropriate candidates: Patients requiring extensively matched transfusions, patients with autoantibodies or other serologic reactivity that impedes the exclusion of clinically signicant antibodies or patients with suspected antibody against an antigen for which typing sera are not available.

4.2.2.1 Sickle Cell Disease and Thalassemia Genotyping of Sickle Cell Disease or Thalassemia patients is not uncommon as these patient groups often require chronic transfusions as treatment. The genotyping has many benefits. First of all, the exact phenotype determination is important as it can be used as a baseline for investigating the cause of potential transfusion reactions or to distinguish between allo and auto (self-induced) antibodies. An additional benefit is that many genotyping methods can detect the GATA site mutation in the DARC gene, which is common in individuals of African descent. Individuals with this mutation are not at risk of Fy(b) alloimmunization and can safely receive Fy(b+) units, leading to the preservation of the scarce Fy(b-) units [43]. Rhesus variants are another cause of alloimmunization among SCD patients. A

frequent variant is a partial C antigen. These individuals are serologically typed as C+ but when exposed to conventional C+ antigens they form anti-C. The detection of this partial antigen and the subsequent allocation of C- blood can prevent alloimmunization for these individuals [44]. Not all SCD patients are prone to alloimmunization. This cannot be determined a priori, and thus all SCD patients should receive antigen matched blood if the risk of alloimmunization is to be reduced. Such a preventive strategy can have an impact on the supply of antigen negative blood. Wilkinson *et al.* [45] explore the availability of extensively matched RBC units for SCD patients. This was investigated by querying a three-day blood bank inventory from primarily Caucasian donors with SCD patient requests. Three levels of matching were investigated: basic level (Rh and K), medium level (Rh, K and Duffy) and high level (Rh, K, Duffy, Kidd, S, s). The availability was determined by counting the number of compatible RBC units from the three-day inventory for each of the 70 SCD patients. A large drawback of this approach is that the units are only counted but not allocated and thus highly compatible units are counted multiple times. In reality a unit can only be issued once, so the availability is very likely to be overestimated. Castilho and Dinardo [46] perform a similar study among Brazilian SCD patients. They also use three levels of matching: basic (Rh and K), extended (Rh, Kell, Duffy, Kidd, MNS, Diego) and extended including RH allele variants. Compatible products for SCD patients are usually available at the basic and extended level. Matching the Rh allele variants is much more challenging, largely due to racial discrepancies between the donor and patient populations.

4.2.2.2 Donors Large-scale genotyping of donors can reveal donors with rare blood types or infrequent negative antigen combinations. Several studies have already demonstrated the effect of such a system. Perreault *et al.* [47] maintained a genotyped donor database of 10555 donors. These donors are selected if they donate at least three times per year so that the genotyping stays cost effective. Flegel *et al.* [48] have created and studied a genotyped donor database consisting of 43066 donors. All donors who self-identified as African American, Hispanic, Asian or Native American were eligible for genotyping as were all donors who donated more than three times in the past year and who did not have the AB phenotype. Both databases report very low discrepancies between the genotype and observed phenotype.

5 Mathematical Formulation

In this chapter we will describe how the problem of matching RBC units to patients can be mathematically modelled. We will discuss how to appropriately model shortages and antigen mismatches for requests for multiple units. Furthermore, we will show how an optimal assignment of RBC units to patients can be computed using Integer Linear Programming. Lastly, we will discuss the computational complexity of resulting matching problem.

5.1 Graph Representation

In essence, matching RBC units to patients is an *assignment problem*. An assignment problem is always characterized by the following structure: there is a supply set with items that can be issued and a demand set with requests that need to be satisfied with items. Furthermore, supply constraints limit how many times an item can be assigned to a request and demand constraints state how many items should or can be assigned to requests. In the case of assigning blood to patients the supply and demand constraints are as follows:

- Every RBC unit can only be issued once.
- Every patient's request must be supplied with the requested number of RBC units.

Assignment problems are easy to visualize using a graph representation. This requires transforming the requests and items and their relations to nodes and edges in a graph. To do this, we create a node for every RBC unit and a node for every patient. Edges between these nodes express the possibility of transfusing a particular patient with a particular unit. There are no edges possible between two patient nodes or two RBC unit nodes. This means that the nodes in the graph can be divided into two groups where every edge connects two nodes of different groups. A graph with this property is called *bipartite*. As not every RBC unit can be assigned to each patient because of AB-RhD compatibility, we cannot create an edge for every patient-unit combination. Instead, only those edges are included for which the RBC unit is AB-RhD compatible with the patient. An example of a bipartite graph representing RBC units and patients is shown in Figure 2. RBC Unit nodes are shown as circles while patient nodes are shown as squares which also contain a number representing the number of units they require.

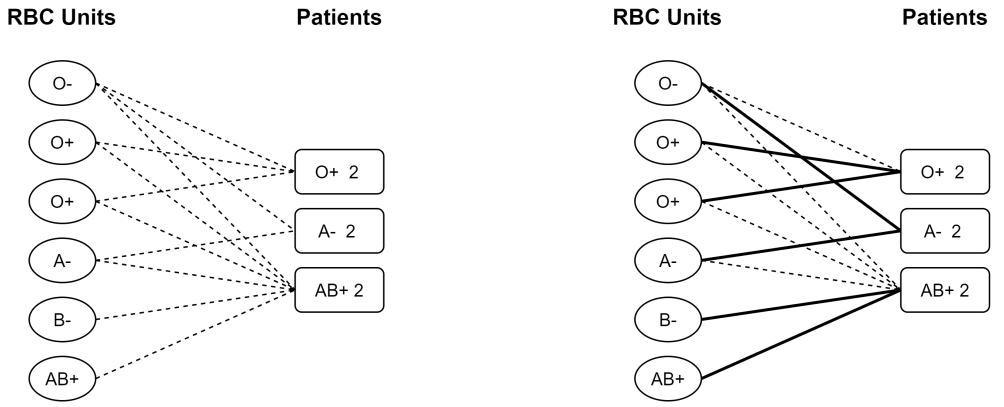

Figure 2: Left: Assignment problem shown as bipartite graph. Edges show compatible unit-patient combinations. Right: The unique valid assignment with the minimal number of shortages.

5.1.1 Shortages

Now that we have represented the assignment problem as a graph, we can construct an algorithm to find an optimal assignment of RBC units to patients. The first priority when computing an assignment is to minimize shortages. We define a shortage as the failure to satisfy a patient with the requested number of RBC units, irrespective of the size of the request. Note that this is different from the definition used by Van Sambeeck *et al.* [6] where one shortage corresponds to one missing RBC unit. That these definitions are not equal is made clear in the following example:

> **Shortage Example** Consider two patients, both with a demand for two RBC units of type O^-. Suppose only two O^- RBC units are available. Therefore there are two options for matching. Either one patient receives both units and the other one receives none or both patients receive one unit each. This situation is shown as a graph in Figure 3.

Figure 3: Left: Both patients are partially satisfied. Right: One patient receives both units so that at least one request is satisfied fully.

> We argue that it is preferable to assign both units to the same patient instead of splitting them. The reason for this is that partially satisfied requests cannot be considered partially valid because we assumed that all requested units are needed for transfusion. In practice, this may not always be the case as regularly not all requested units actually end up transfused. However, we cannot predict this beforehand and therefore we assume that all requested units are needed.

If instead a shortage is incurred for every missing RBC unit then in both assignments above the number of shortages would be two as both lack two O^- units. If we instead use our new definition of shortages, then the assignment on the right would have only one shortage whereas the assignment on the left has two as both patients are not fully satisfied with the number of units demanded.

Because a partially satisfied request is just as much a shortage as a request which is not assigned any units, we will add a constraint to our model which says that partially satisfied requests are not allowed. Thus, in every valid assignment every request is either assigned with the requested number of units or with none at all. This is mainly done for practical convenience as it does not limit the possible solutions but removes solutions with a valid equivalent.

The reason that we have chosen to work with a new definition of shortages is that we think that this definition is better suited to assign units optimally. The shortcomings of the old definition have been shown in the simple example above. However, the new definition also has some disadvantages. Firstly, the number of shortages does no longer directly correspond to the number of RBC units short, which is on its own a useful metric. Knowing how many extra units may have been sufficient to satisfy all requests could give an indication of the inventory size necessary to satisfy the given demand. However, we think that the ability to assign units more intelligently to patients who

require multiple units and thus prevent more shortages is an advantage that outweighs the loss of this insight.

Another disadvantage is that the new shortage metric is non-linear in the number of units assigned. This non-linearity makes the model harder to solve because it adds an all-or-nothing behaviour. If we consider a request for four units then the shortage penalty is one when the request is assigned 0, 1, 2 or 3 units. Intuitively one would say that satisfying 3 out of 4 units is 'closer' to the valid solution, but this is not reflected in the shortage cost. In Section 5.4 we will show what the consequences of this are for the computational complexity of the problem.

In summary, we can define the following hard-constraints of a valid assignment. Hard-constraints are constraints to which every assignment must adhere to be considered a valid solution.

- **Shortages should be minimized at all times.** Any assignment with x shortages is preferred over any assignment with $> x$ shortages.
- **Any assignment should consist of AB-RhD compatible matches only**. No mismatches on antigens A, B or RhD are allowed.
- **Request cannot be partially satisfied.** A patient should be assigned either the requested number of RBC units or none.

5.2 Integer Linear Programming

5.2.1 Linear Programming

We will use Linear Programming (LP) to compute an assignment which adheres to all the constraints that govern the validity of the assignment problem while also optimizing the quality of the matches made. Linear programming is a mathematical technique that allows the formulation of a mathematical model and subsequent optimization of an objective function. To construct a LP-model first *decision variables* are defined which form a multi-dimensional solution space. Then linear constraints are added to bound the valid solution space. Lastly, an objective function is added which specifies the value of a point in the solution space. Then a *LP-solver* can be used to find the point in the valid solution space with the lowest value of the objective function.

A typical linear programming model looks as follows:

$$\begin{aligned}
\min c_1x_1 + c_2x_2 + ... + c_nx_n & \qquad \text{(Objective Function)} \\
\text{Subjetct To:} & \\
a_{11}x_1 + ... + a_{1n}x_n \leqslant b_1 & \qquad \text{(Constraint 1)} \\
a_{21}x_1 + ... + a_{2n}x_n \leqslant b_2 & \qquad \text{(Constraint 2)} \\
\vdots & \\
a_{m1}x_1 + ... + a_{mn}x_n \leqslant b_m & \qquad \text{(Constraint m)} \\
x_1 \geqslant 0, x_2 \geqslant 0, ..., x_n \geqslant 0 & \qquad \text{(Non-negativity constraints)}
\end{aligned}$$

Here x is the vector of decision variables, c a vector of weights associated with each decision variable in the objective, A the constraint matrix specifying the coefficients of the decision variables in the constraints and b the vector with values for the right-hand sides of the constraints:

$$A = \begin{pmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \dots & a_{mn} \end{pmatrix} \quad c = \begin{pmatrix} c_1 \\ \vdots \\ c_n \end{pmatrix} \quad b = \begin{pmatrix} b_1 \\ \vdots \\ b_m \end{pmatrix} \quad x = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

In more concise notation:

$$\begin{aligned} &\min cx \\ &\text{Subject to:} \\ &Ax \leqslant b \\ &x \geqslant 0 \end{aligned}$$

Linear programs can only have linear behaviour, both in the objective function as in the constraints. Therefore, it is impossible to use the product of two or more decision variables in the objective function or constraint. Often this is no issue as many problems, including assignment problems, are easily formulated using linear constraints.

5.2.2 Integer Linear Programming

In regular linear programming models, the decision variables can take any non-negative rational fractional value (within the specified bounds). It is however not uncommon to have decision variables which represent a discrete quantity. Examples of this include the number of trucks to buy, how many people to employ or whether to select an edge in a graph. Similarly, it is impossible to supply a patient with a fractional RBC unit. Therefore, another constraint should be added to the linear program which specifies that the decision variables should be integral. An LP with integrality constraints is called an *Integer Linear Program (ILP)*. The integrality constraints make finding the optimal solution much harder. A possibility to still compute a fast solution is to ignore the integrality constraints and solve the remaining linear program, called the *LP-relaxation*. The resulting solution can then be rounded to make it integral. The problem is that there is no guarantee that this solution is optimal or even within the feasible solution space. An example of this is shown in Figure 4.

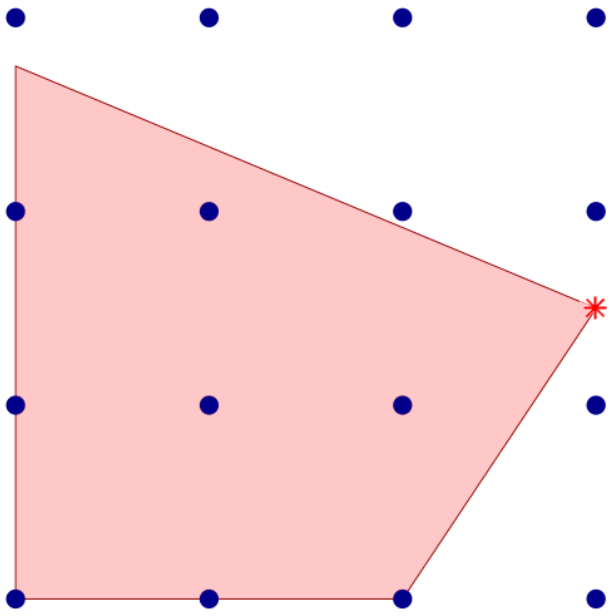

Figure 4: The vertex ✻ is the optimal solution of the linear relaxation. The blue dots represent the integral solutions. The two nearest integral solutions to the LP-relaxation are both infeasible.

Since general ILP solving is NP-hard [49] there is no polynomial algorithm which can compute the optimal solution in the general case. Even 0-1-integer programming, where all decision variables are binary, is NP-hard [50]. Although ILP solving is NP-hard it does not mean that it cannot be solved. Several techniques exist to compute optimal integer solutions. Many commercial solvers are available which can be used to compute optimal solutions to ILP problems.

5.2.3 ILP Formulation

In this section we will formulate the blood assignment problem as an ILP. First, we will introduce some notation:

Notation	Meaning
b_i	Blood phenotype of patient $i \in [1, ..., n]$
d_j	Blood phenotype of RBC unit $j \in [1, ..., m]$
$\mathcal{A}$	Set of minor antigens included
$\mathcal{A}(k)$	Minor antigen with index k
$a(k)$	Relative immunogenicity of antigen $\mathcal{A}(k)$
$b_i(k)$ or $d_j(k)$	1 if b_i or d_j is positive for antigen $\mathcal{A}(k)$; 0 otherwise
u_i	number of RBC units requested for patient i

Table 6: Mathematical notation for blood, RBC units and patients.

We can now formulate an ILP model which solves the assignment problem optimally. The model has the following decision variables:

$$x_{ij} = \begin{cases} 1, & \text{if RBC unit } j \text{ is assigned to patient } i. \\ 0, & \text{otherwise.} \end{cases}$$

$$s_i = \begin{cases} 1, & \text{if a shortage is incurred for patient } i. \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

To model AB-RhD compatibility we use the following parameters:

$$c_{ij} = \begin{cases} 1, & \text{if there is an edge between patient } i \text{ and RBC unit } j \text{ in the bipartite graph.} \\ 0, & \text{otherwise.} \end{cases}$$

Now we can construct a basic ILP that can compute a valid assignment of RBC units to requests.

Model 1

$$\min \sum_i s_i \cdot (n+1) \qquad \text{(Objective)}$$

s.t.

$$\sum_j x_{ij} + s_i \cdot u_i = u_i \ \ \forall i \tag{2}$$

$$\sum_i x_{ij} \leqslant 1 \ \ \forall j \tag{3}$$

$$x_{ij} \leqslant c_{ij} \ \ \forall i, j \tag{4}$$

$$x_{ij}, s_i \in \{0, 1\} \ \ \forall i, j \tag{5}$$

The objective function adds a penalty of $(n+1)$ for every patient which is not assigned any RBC units, as s_i is the binary indicator for a shortage for patient i. The value of $(n+1)$ seems arbitrary for now, but it is necessary to later include minor antigen compatibility. Constraint (2) ensures that every patient is satisfied with the number of units requested or a shortage is incurred and no units are assigned. Constraint (3) ensures that every RBC unit can be assigned at most once. Constraint (4) only allows pairing AB-RhD compatible matches. Constraint (5) ensures that all variables are binary.

Previously we have mentioned that linear programs are generally much easier to solve than integer linear programs. The formulation above has integrality constraints (all variables are binary) and therefore belongs to the latter category. Fortunately we can exploit a property of integer linear programs called *Total Unimodularity (TUM)*. When the constraint matrix (A) of an LP is TUM and the RHS vector (b) is integral then the optimal solution, if it exists, must be integral [51]. Because of this property it is sufficient to solve the linear relaxation of the ILP model as it will give the correct optimal integer solution.

The constraint matrix of the ILP formulation above does not have the Total Unimodularity property. This is because one requirement for Total Unimodularity is that all coefficients in the constraint matrix are either -1, 0 or 1. When we examine constraint (2) we see that the s_i decision variables have coefficient u_i which can be larger than 1.

We can now also see the consequence of our new definition for shortages. The fact that a request has to be satisfied either fully or not at all is exactly the cause of the loss of the Total Unimodularity property. Therefore it also means that the model is not guaranteed to be solvable by relaxation. Fortunately, this problem is mitigated when a sufficiently sized inventory is used. When this is the case, we expect the number of shortages to be (near-)zero. When there are no shortages in the optimal solution all s_i are equal to 0. When all s_i are 0 then constraint (2) can be rewritten as follows:

$$\sum_j x_{ij} = u_i \ \ \forall i$$

If we assume that the optimal solution does not have any shortages, the optimal solution is still found when using the rewritten constraint. Using this constraint, the resulting ILP formulation is known as a *transportation problem*, which is well known to be TUM if all supply and demand values are integral [52]. Therefore, it is sufficient to solve the linear relaxation of the model to compute the optimal solution.

Because of this observation we can simply start the solving process by solving the relaxation of the ILP formulation. If the solution of the relaxation satisfies all demand, then we know we found the optimal integer solution. Otherwise, the solver can continue to solve the ILP formulation as usual. As we expect to solve mainly instances where the inventory size is large enough to handle demand, we expect only few shortages. Therefore, the disadvantage of not having the TU property is not a big issue in practice.

5.3 Minor Antigen Compatibility

The ILP formulation above can be used to compute an assignment of RBC units to patients such that the number of shortages is minimized and only AB-RhD compatible assignments are made. We have not yet included any information about minor antigens in this model, meaning that these compatibility relations are not taken into account when computing the optimal assignment.

5.3.1 Compatibility using edges

One way to include minor antigen compatibility in the graph is to remove edges whenever the corresponding patient and unit are incompatible on a certain antigen. This is the approach used by Van Sambeeck *et al.* [6]. While this may seem like a good idea, it is far from ideal as there are two problems with this approach. First, there is the problem of which antigens to include. If all 14 minor antigens are included on top of the A, B and RhD antigens then there will only be few edges remaining as every valid match must be fully compatible on 17 antigens. Reducing the set of antigens considered will partially solve this problem but then no matching can be done for the excluded antigens as they are not included in the problem. The second problem is that the minor antigens included are treated equally to the major antigens, as both can lead to the removal of edges. For the majority of patients the minor antigen matching is not a must, whereas matching correctly for A, B and RhD is mandatory. By giving the same status to major and minor antigens shortages become intertwined with the quality of minor antigen matches. As discussed previously, minimizing shortages has a clear priority over minor antigen matching. Therefore, we propose another way to model the problem which respects this separation of priorities.

5.3.2 Compatibility using costs

We use the same graph as before with edges for every AB-RhD compatible combination. This will allow us to use the ILP formulation above to minimize the number of shortages relatively simply. To account for the compatibility of the minor antigens a cost is assigned to every edge between a unit and patient which represents how well the patient and unit are compatible on the minor antigens included. The advantage of this approach is that the quality of a solution is decoupled from the number of shortages. This is because the number of shortages is determined by the presence of edges while the quality is determined by the total cost of the edges chosen. This makes it possible to explicitly prioritize the prevention of shortages over the quality of the minor antigen matching which is in line with the previously established constraints.

To capture the compatibility of the minor antigens between a unit and a patient in a single cost we use the relative immunogenicity as an indication of the risk per mismatched antigen as described in Section 2.5. If a mismatch on a specific antigen occurs, the relative immunogenicity of this antigen is added as a penalty to the total cost of the patient-unit match. In more mathematical terms, if $b_i(k)$ and $d_j(k)$ indicate the presence of the minor antigens indexed by k in patient i and RBC unit j and $a(k)$ is the relative immunogenicity of antigen k, then the total cost of matching unit j to patient i is:

$$a_{ij} = \sum_k d_j(k) \cdot (1 - b_i(k)) \cdot a(k) \tag{6}$$

We multiply $d_j(k)$ by $(1 - b_i(k))$ because a mismatch is incurred only when the patient is negative ($b_i(k) = 0$) for antigen k while the RBC unit is positive ($d_j(k) = 1$).

Because this mismatch cost is fully determined by the antigens present in the patient's blood and those in the RBC unit, the a_{ij} values can be precomputed for every edge. We can then easily extend the objective of the ILP model above with these costs:

$$\min \sum_i s_i \cdot (n+1) + \sum_i \sum_j a_{ij} x_{ij} \qquad \text{(Objective)}$$

The cost a_{ij} of a patient-unit match is multiplied by the decision variable x_{ij} to only penalize matches included in the solution. Since changes to the objective function do not

alter the constraint matrix, the model will remain TUM if the optimum can satisfy all demand. Furthermore, it is now also clear why we use the penalty of $n+1$ per shortage as we must make sure that trading a shortage for a higher quality matching on the minor antigens may not be possible. Since the maximum theoretical alloimmunization risk cost per patient is 1 and there are n patients, the cost of a shortage should be larger than n, so we choose $n+1$. The precomputation of the mismatch costs per unit-patient combination makes optimization simple. However, if the number of units demanded per patient can be larger than one we cannot simply precompute the costs per unit anymore, as when multiple units assigned to the same patient which induce the same mismatch have a nett result of one mismatch and not two. This problem is illustrated further in the next section.

5.3.3 Multiple Units per Patient

The assumption that we can precompute the costs of the alloimmunization risk is only true when every patient requires exactly one RBC unit. This is because when multiple units are supplied to a single patient the number of mismatches is determined by the union of the antigens present, which is slightly different from the sum of all individual RBC units matched. In practice most patients receive more than one RBC unit. When a patient receives two or more RBC units, we cannot compute the costs of the alloimmunization as before. This is because in general, the fact that a mismatch on a certain antigen occurs is much more important than the number of transfused units with this mismatch. To illustrate this a simplified example is shown below.

Mismatch example Consider two patients labelled *Patient 1* and *Patient 2*. Both patients have major blood group A^+. For simplicity only one minor antigen is shown, namely Rhesus E. As can be seen in Table 7, both patients lack the Rhesus E antigen in their phenotype.

	Units needed	Major	RhE
Patient 1	1	A^+	-
Patient 2	2	A^+	-

Table 7: Two example A^+ patients. For both patients the number of units needed is known, as well as their limited phenotype.

Patient 1 requires only one unit while Patient 2 needs two. Fortunately, there are three units available for matching. These are shown below.

	Major	RhE
Unit 1	O^+	+
Unit 2	O^+	+
Unit 3	O^+	-

Table 8: Three example O+ units. For every unit the limited phenotype is shown.

All three units are O^+ and compatible with both patients on AB-RhD basis. However, both patients lack the Rhesus E antigen while only Unit 3 is Rhesus E negative. This creates two options for assigning the units to both patients: Either Unit 3 is assigned to Patient 1 or it is assigned to Patient 2. Deciding which of these is the preferred solution is the same as answering the following question: Is transfusing one additional mismatched RBC unit for an already mismatched patient equally as bad as inducing a mismatch for an otherwise correctly matched patient?

We have discussed this question with an expert on immunology. They have informed us that although it is thought that more exposure to a foreign antigen is believed to increase the risk of alloimmunization, having one 300ml mismatching RBC unit is already a very large amount of exposure. Therefore, the difference between mismatching or not is many times greater than the difference between one or multiple mismatched units. Because it is difficult to quantify the difference in scale between these different forms of mismatching we have simplified it to the following: *For every patient the first mismatch on a certain antigen is met with the full alloimmunization penalty while any additional mismatches on this antigen do not further increase the penalty.*

What this means in practice is that it is no longer individual RBC units from which the mismatches are calculated but instead we first consider the union of all matched RBC units to a patient and then sum over the existing mismatches. An example of this is shown in Table 9.

	C	c	E	e	K	k	M	N	S	s	Fy(a)	Fy(b)	Jk(a)	Jk(b)
Unit 1	+	-	-	+	-	+	+	-	-	+	-	+	-	+
Unit 2	+	+	-	+	-	+	+	+	-	+	-	-	+	+
Union	+	+	-	+	-	+	+	+	-	+	-	+	+	+

Table 9: Two RBC units with all 14 minor antigens and their union.

The union of multiple RBC units can be considered as the result of mixing them. It will show all antigens which enter the blood of a patient who is transfused with these units. The antigen mismatch penalty will now be counted once for every antigen in the union which is foreign to the patient's blood type. This means that alloimmunization is only counted once per antigen, even if more than one of the matched RBC units contain the mismatching antigen.

Going back to our simple example the result of this approach is very clear. Because Patient 2 will always receive at least one unit which mismatches on *RhE* and a mismatch will only be counted once, it is beneficial to assign both *RhE*-positive units to Patient 2. This is because it will prevent a mismatch for Patient 1. If instead both patients would receive at least one mismatching unit then in total two mismatches would be induced, with double the alloimmunization costs.

We believe that this rule captures the essence of antigen matching for multiple units per patient. The fact that there is no penalty difference between multiple mismatching units or one single mismatching unit can be considered a shortcoming. Adding a mismatching unit to an already mismatched patient could now be thought of as a "free" mismatch. However, this might also work advantageously as donor blood with many antigens can be assigned more easily leaving the more negative blood available to patients for whom any mismatching should be prevented.

5.3.4 MINRAR

We are now ready to formulate an ILP which solves the problem of assigning a set of RBC units to a given set of patients. We have identified the following hard-constraints:

- Always minimize the number of patient shortages
- All matches must be AB-RhD compatible
- Partially satisfied requests are not allowed

Furthermore we have the following soft constraint:

- Minimize the probability of antibody formation by minimizing minor antigen mismatches

Lastly, the following assumption is used:

- A patient can only be mismatched once per antigen, irrespective of the number of assigned units that mismatch on the antigen.

We will combine these four constraints and assumption in the ILP below. We will call our formulation *MINRAR* as besides minimizing shortages the main goal of our ILP is to **MIN**imize the **R**elative **A**lloimmunization **R**isk for all patients.

ILP Formulation Because the exact alloimmunization penalty depends on the union of the phenotypes of the assigned RBC units, it is no longer possible to precompute the alloimmunization values in the optimization model. The union is not known beforehand and precomputing all possible unions will lead to an explosion of the number of decision variables. To correctly model the new behaviour, we have to introduce an extra set of decision variables to represent the union of the phenotypes of the assigned units per patient. We will use one variable per antigen per patient to indicate whether any of the assigned units mismatch on this antigen. Formally:

$$y_{ik} = \begin{cases} 1, & \text{if patient } i \text{ is mismatched at least once on antigen } k. \\ 0, & \text{otherwise.} \end{cases}$$

Now we can write the ILP as follows:

$$\min \sum_i s_i \cdot (n+1) + \sum_i \sum_k y_{ik} \cdot a(k) \qquad \text{(Objective)}$$

s.t.

$$\sum_j x_{ij} + s_i \cdot u_i = u_i \;\; \forall i \tag{7}$$

$$\sum_i x_{ij} \leqslant 1 \;\; \forall j \tag{8}$$

$$u_i y_{ik} \geqslant \sum_j x_{ij} d_j(k) \;\; \forall i,k \;\; \textit{where} \;\; b_i(k) = 0 \tag{9}$$

$$x_{ij} \leqslant c_{ij} \;\; \forall i,j \tag{10}$$

$$x_{ij}, y_{ik}, s_i \in \{0,1\} \;\; \forall i,j,k \tag{11}$$

The objective function is extended with an extra term, namely $\sum_i \sum_k y_{ik} \cdot a(k)$. This term sums all the mismatches and counts the alloimmunization risk per mismatch. To

ensure the correctness of the y_{ik} variables we have added constraint (9). This constraint forces y_{ik} to take the value of 1 if one or more of the assigned units to patient i induce a mismatch on antigen k. Note that the constraint is only added when $b_i(k) = 0$, as when $b_i(k) = 1$ no mismatch can occur.

We can observe that the y_{ik} variables can have a coefficient in the constraint matrix which is not in $\{-1, 0, 1\}$ as the u_i parameter can be larger than one. The consequence of this is similar to that of the s_i variables. The model is no longer TUM, even when the number of shortages is zero. Despite this, we expect the practical performance to be still relatively good. This is because of the following reasons:

- Several antigens are either infrequent or highly frequent. Therefore, it will often be the case that all antigens mismatch or none, removing much of the non-linearity.
- More than 40% of all requests have $u_i = 1$ and therefore have no non-linearity. A further 50% of requests are for exactly two units thus limiting the prevalence of the non-linearity as the issue becomes greater when more units are demanded.

5.4 Computational Complexity

Classical assignment problems are easily solved by transforming the bipartite graph to a Minimum Cost Maximum Flow problem which is well known to be solvable in polynomial time. Although the problem of assigning RBC units to patients is also an assignment problem, it has additional constraints which prevent a simple carry over of the polynomial solvability. In our case, the constraint of having to satisfy a patient exactly with the required number of units leads to extra complexity. This gives rise to the question whether this problem is *NP-Hard* or not. In this section we will first define an abstract version of the blood assignment problem with extensive matching, then prove that the general variant is *NP-Complete* whilst the variant with a bounded number of antigens and number of units per patient is solvable in polynomial time.

5.4.1 Problem Definition

The Blood Assignment Problem with Extensive Matching (bapem) is defined as follows. We have a set of N patients and M RBC units and an assignment of units to patients has to be computed such that every patient $i \in [1, ..., N]$ is satisfied with u_i units or none. If no units are assigned to a patient a shortage is incurred. Furthermore, both patients and units have a phenotype which stores the presence or absence of K different antigens. A unit and a patient *mismatch* on antigen $k \in K$ when the antigen is present in the phenotype of the unit while absent in the phenotype of the patient. Mismatching on an antigen is allowed at a cost $a(k)$, which is given for each antigen $k \in K$. Lastly, a patient can only be mismatched once per antigen, meaning that when more than one of the assigned units for a patient mismatch on antigen k, the cost is still $a(k)$. Is there an assignment with $\leqslant d$ shortages and $\leqslant c$ total cost?

5.4.2 NP-Completeness

We will prove that this problem is NP-Complete when the number of antigens K is unbounded. To do this we will use a reduction to 3-Dimensional Matching (3DM) which is a well known NP-Complete problem [50]. 3DM is defined as follows. Let X, Y and Z be disjoint sets where $|X| = |Y| = |Z| = k$ and let T be a set of triples where each triple contains one element of X, Y and Z. The decision problem is whether there exists a subset M of T such that all elements in X, Y and Z are contained in exactly one triplet in M and $|M| = k$?

We will prove that any generic instance of 3DM can be transformed to an instance of BAPEM. First, we take an arbitrary instance of 3DM with X, Y, Z and T as described above. An example is shown in Tables 10 and 11. Now we create $n = |T|$ patients where each patient corresponds to a triple in T. Similarly, we introduce n antigens where antigen j corresponds to patient j (and thus to triple j). We let each patient be negative for the antigen that corresponds to its triple and positive for all other antigens. Furthermore, each patient has a demand of 3 units. This conversion is shown in Table 12. Now we create an RBC unit for each element in X, Y and Z. We let an RBC unit be negative for all antigens corresponding to the triples which contain the element represented by the RBC unit and positive for all other antigens. Let the mismatching cost $a(k)$ of all antigens be any positive value, for instance 1. Now we ask the question whether it is possible to satisfy k patients (thus, $n - k$ shortages) with zero mismatching cost. This conversion is shown in Table 13. When there exists a subset M of T which is solution to the 3DM instance, then it is possible to satisfy the patients corresponding to the triplets in M each with three units and no mismatching costs. This is because each of these patients can be assigned the units corresponding to the elements in the triplets represented by the patients. As each element is included in exactly one triplet in M, each RBC unit is issued exactly once and the solution has no mismatches as every unit is negative for the antigens corresponding to the triplets that contain the element represented by the unit. Similarly, if there exists a solution to the constructed BAPEM instance with $n - k$ shortages and zero mismatch costs, then this solution must correspond to a subset M of T which contains k triplets that exactly cover each element in X, Y and Z once. Note that there do not exists solutions with fewer than $n - k$ shortages as there are only $3k$ units available and thus at most k patients can be satisfied fully.

Set	Items
X	x_1, x_2, x_3
Y	y_1, y_2, y_3
Z	z_1, z_2, z_3

Table 10: Sets X, Y and Z with $k = |X| = |Y| = |Z| = 3$

Triple Index	Items
T_1	x_1, y_3, z_1
T_2	x_2, y_3, z_2
T_3	x_3, y_2, z_2
T_4	x_1, y_1, z_3
T_5	x_2, y_1, z_3

Table 11: Set of triples T with 5 triples each containing an item from X, Y and Z

Antigen	1 (T_1)	2 (T_2)	3 (T_3)	4 (T_4)	5 (T_5)
Patient 1 (T_1)		✓	✓	✓	✓
Patient 2 (T_1)	✓		✓	✓	✓
Patient 3 (T_3)	✓	✓		✓	✓
Patient 4 (T_4)	✓	✓	✓		✓
Patient 5 (T_5)	✓	✓	✓	✓	

Table 12: Constructed set of patients. Each patient corresponds to a triple in T. Furthermore, each antigen corresponds to a triple T and each patient expresses all antigens except the antigen corresponding to its triplet.

Antigen	1 (T_1)	2 (T_2)	3 (T_3)	4 (T_4)	5 (T_5)
Unit 1 (x_1)		✓	✓		✓
Unit 2 (x_2)	✓		✓	✓	
Unit 3 (x_3)	✓	✓		✓	✓
Unit 4 (y_1)	✓	✓	✓		
Unit 5 (y_2)	✓	✓		✓	✓
Unit 6 (y_3)			✓	✓	✓
Unit 7 (z_1)		✓	✓	✓	✓
Unit 8 (z_2)	✓			✓	✓
Unit 9 (z_3)	✓	✓	✓		

Table 13: Constructed set of units. Each unit corresponds to an item in set X, Y or Z. The units are negative for the antigens representing the triplets that contain their corresponding item.

Now we can see that the solution $M = \{T_1, T_3, T_5\}$ is equivalent to patients 1, 3 and 5 being satisfied with three units each and total mismatching costs of zero:

$$M = \{T_1, T_3, T_5\} \Leftrightarrow \begin{pmatrix} \text{Patient 1:} & \text{Unit 1, Unit 6, Unit 7} \\ \text{Patient 3:} & \text{Unit 3, Unit 5, Unit 8} \\ \text{Patient 5:} & \text{Unit 2, Unit 4, Unit 9} \end{pmatrix}$$

The transformation of the 3DM instance to the BAPEM instance is clearly polynomial. Therefore, we have proven that the BAPEM problem is NP-Complete. □

5.4.3 Hardness with a Bounded number of Antigens

When the number of antigens is fixed, the reduction above does no longer work as we need the number of antigens to be able to grow with N. Let U denote the maximum number of units per patient. We will now prove that for every fixed number of antigens $K = O(1)$ and maximal number of units per patient $U = O(1)$ there exists a polynomial (albeit slow) algorithm to solve the BAPEM, thus disproving NP-Completeness for this variant. To prove this, we will use the following theorem by Lenstra [53]:

Theorem 1 [53] *For each fixed natural number n, there exists a polynomial algorithm which solves the integer linear programming problem*

$$\max\{cx |\ \ Ax \leqslant b;\ \ x \text{ integral}\}, \tag{12}$$

where A has n columns (variables), and where all input data are rational.

Thus, to prove that the problem can be solved in polynomial time we have to prove that the problem can be written as an ILP with a bounded number of variables. The proof will make use of the following observation: when the number of antigens K and the maximal number of units per patient U are bounded then the number of unique patient requests is bounded by $U2^K$ as there are 2^K theoretical phenotypes and each request can be for $1, ..., U$ units.

Let i be an index enumerating all these $U2^K$ different 'patient types' and let P_i be the number of patients of type i. The number of unique units is bounded by 2^K for the same reason. Let $L = 2^K$ denote the number of different phenotypes and h index them all. As each RBC unit can be issued together with other units we will enumerate all possible combinations of l RBC units for each $l = 1, ..., U$. This gives L^U possible combinations

which we will enumerate with index by j. Let q_{hj} be a parameter indicating how many units with phenotype h are used in combination j and B_h be the number of available units of phenotype h.

Now we introduce an integer decision variable x_{ij} which specifies how many times the combination of units j is assigned to patient type i. When the number of units requested for patient type i does not match the number of units in combination j we have $x_{ij} = 0$. Also, when the phenotype of patient type i is incompatible with any of the phenotypes in combination j on one or more major antigens we have $x_{ij} = 0$. For all other x_{ij} we can compute the mismatch costs c_{ij} as we know the phenotypes of all units supplied to this patient type. Lastly, we add an integer decision variable S_i which represents the number of patients of patient type i whose demand is not satisfied (with costs C_i per shortage).

Now we can construct the following ILP:

$$\min \sum_i \sum_j c_{ij} x_{ij} + \sum_i S_i C_i \tag{13}$$

s.t

$$\sum_i \sum_j q_{hj} x_{ij} \leqslant B_h \quad \forall h \tag{14}$$

$$\sum_j x_{ij} + S_i = P_i \quad \forall i \tag{15}$$

$$x_{ij}, S_i \geqslant 0 \quad \forall i, j \tag{16}$$

The objective minimizes the mismatch costs as well as the total number of shortages. When C_i is chosen as the total number of patients then the minimization of shortages is always prioritized over mismatch costs. Constraint 14 limits the number of RBC units of phenotype h issued to the available number of units of that type. Constraint 15 ensures that all patients of 'patient type' i are either assigned units or shorted. Lastly Constraint 16 forces non-negativity for the decision variables.

The total number of x_{ij} variables is $O(U2^{2K}) = O(1)$ as $U = O(1)$ and $K = O(1)$. Furthermore the number of shortage variables is $O(2^K) = O(1)$. Thus there are $O(1)$ variables and thus $O(1)$ columns in the ILP [1]. Therefore Theorem 1 can be applied and thus the ILP is solvable in polynomial time. □

[1] For completeness, the number of constraints is also $O(1)$ for similar reasons.

6 Online Model

The MINRAR ILP formulation in the previous chapter can be used to optimally issue a given set of inventory units to a given set of patients. Shortages are minimized and patients receive blood which is compatible to their own blood as well as possible, given the units in inventory. Often there are many more RBC units available than needed. The units which are not assigned have no influence on the objective function of the MINRAR ILP, as they are not part of the solution. However, in practice the unassigned units remain in inventory and must be issued on a later moment. Matching RBC units to patients is not done once, but instead iteratively. Usually, every day units must be issued to patients and new units are ordered. Similarly, in distribution centres units are issued to hospitals and donors are invited for new donations. The fact that inventory management is a continuous process means that the scope of the assignment problem is enlarged. The decisions made on the current day affect not only current patients but also affect future patients.

To make the ILP formulation more aware of this iterative process we must take two new factors into account:

- **The age of RBC units in inventory.** As RBC units expire after 35 days, we must make sure that units with little remaining shelf life are prioritized for issuing.
- **Highly compatible products should be issued with care to prevent future shortages and mismatches.** If more antigen-negative blood is heavily substituted, then this likely leads to more shortages and mismatches over the long term.

As the impact of both these factors only becomes apparent with delay it is not obvious how they should influence the more static assignment formulation to perform well over a longer period. In this new setting there is much more information unknown. Future supply is not known and demand is also stochastic. A setting where the problem is iteratively solved while more and more information becomes available is called an *online* approach. The MINRAR ILP formulation described in the previous chapter solves an *offline* problem. All information is known at the time of computing the optimal assignment. In this chapter we describe how we will alter the model to be able to operate in an online setting.

6.1 Issuing Age

As mentioned before, RBC units have a maximum shelf life of 35 days. If a unit in inventory has only a few days of shelf life left, its issuing should be prioritized in the matching to prevent it from outdating. Therefore, it should receive a 'discount' in the matching to make it preferable to use over other units. To achieve this, we compute a discount factor based on the remaining shelf life t of a unit. Only units with a few days of remaining shelf life should receive high discount. The difference between one or two days of remaining shelf life is of much greater importance than the difference between three and four weeks. To capture this in costs we use the same exponential function as used by Van Sambeeck *et al.* [6]:

$$o(t) = \exp\left(\frac{7t}{35} \cdot \ln 2\right) = \left(\frac{1}{2}\right)^{t/5} \tag{17}$$

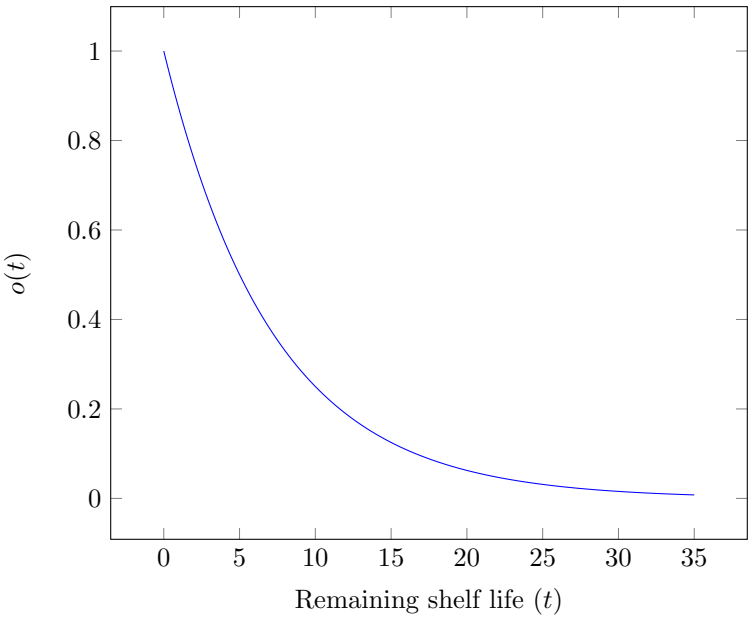

As can be seen fresh units have a low discount factor and furthermore the difference between their discounts is low. Only when units have less than 10 days of remaining shelf life the discount factor starts increasing rapidly. The discount factor per unit doubles every five days. When we want to include this term in the objective function, we want to maximize this value over all units issued. Therefore, the term must be negative, as the objective function is minimized. If we let r_j denote the age of unit j in inventory, then the term to add to the objective function is as follows:

$$\sum_i \sum_j -x_{ij} \cdot o(35 - r_j) \qquad \text{(FIFO)}$$

We will refer to this formula as the *FIFO* term, as it results in FIFO like issuing by making units which are too old more favourable for issuing.

6.2 Saving Highly Compatible Blood Types

As mentioned earlier, an issuing strategy for which the performance is measured over the long term must make sure to limit the number of substitutions of highly compatible blood types. Highly compatible blood types are phenotypes with a relatively high number of negative antigens. This makes them highly usable as they are compatible with many other phenotypes. A substitution occurs when such an RBC unit is transfused to a patient with a more antigen positive phenotype. This means that the patient could have been transfused with less usable blood without any additional alloimmunization risk. When only a limited time span is considered and no penalty is used for substitutions, it may be advantageous to perform many substitutions to minimize the number of direct mismatches. But this approach will likely result in an increase of mismatches in the long run, as more antigen positive units will accumulate in the inventory and eventually outdate or be forcibly matched. Furthermore, heavy AB-RhD substitution may lead directly to shortages in subsequent days. Thus, a way to penalize substitutions is necessary.

6.2.1 Previous work

In the predecessor of this work, Van Sambeeck *et al.* [6] used a measure called *Relative Opportunity Loss (ROL)* to account for this problem. This measure specifies how much matching potential is lost by assigning a particular RBC unit to a particular patient. To give the definition of ROL we must first define the *usability* of a blood product:

Definition. The *Usability* of a blood product, given antigen set $\mathcal{A}$, is the fraction of the population that can receive this blood without any mismatches on the antigens in $\mathcal{A}$. It is therefore equal to the probability that a random patient can be transfused with this blood product when the transfusion must be compatible on all antigens in $\mathcal{A}$. Mathematically, the usability of blood d is denoted as:

$$U(d, \mathcal{A}) = \sum_{b : b \leqslant_{\mathcal{A}} d} p(b) \tag{18}$$

In this formula $b : b \leqslant_{\mathcal{A}} d$ denotes all phenotypes b that can receive blood d without inducing any mismatch on the antigens in $\mathcal{A}$ and $p(b)$ is the prevalence of phenotype b.

In our case we include all antigens with nonzero relative immunogenicity. The set containing all these antigens will be denoted as $\mathcal{A}_{14} = \{A, B, D, C, c, E, e, K, M, S, Fy(a), Fy(b), Jk(a), Jk(b)\}$.

Note that the definition of usability is not limited to RBC units. We can just as well compute the usability of the blood type of a patient. This is useful because it allows us to compare the usability of the supplied blood for a patient with his own blood type. This is exactly what is computed in the relative opportunity loss:

Definition. The *Relative Opportunity Loss* [6] of transfusing patient i with RBC unit j is defined as:

$$R(b_i, d_j, \mathcal{A}) = \frac{U(d_j, \mathcal{A}) - U(b_i, \mathcal{A})}{U(d_j, \mathcal{A})} \tag{19}$$

Here again $\mathcal{A}$ denotes the set of antigens considered relevant. Because in their model Van Sambeeck *et al.* [6] allow only compatible matches on all antigens in $\mathcal{A}$, the relative opportunity loss can never be negative. When both blood types are identical the relative opportunity loss is 0. The more the assigned RBC unit is negative where it is not necessary, the higher the relative opportunity loss will be.

Van Sambeeck *et al.* [6] do not directly include the sum of the relative opportunity loss in the objective. First each ROL value is rescaled using the following formula:

$$\bar{R}(b_i, d_j, \mathcal{A}) = 1 - \exp(R(b_i, d_j, \mathcal{A}) \cdot -7 \ln 2) \tag{20}$$

This exponential is the same as the one used for the issuing age discount as described in Section 6.1. The result of using it is that small ROL values are amplified whereas large values are more packed together. This means that there is a large difference between identical issuing and near-identical issuing and a relatively small difference between moderate and heavy substitution. The reason for this scaling is not mentioned, only that it performs well in practice.

6.2.2 Our method

In the previous work the relative opportunity loss measure is used to prevent over-substitution. This measure is solely based on the prevalence of the antigens and ignores their immunogenicity. Another characteristic of the relative opportunity loss measure is that it is computed over both major and minor antigens conjointly. This means that highly usable blood like O^- can still have a low relative opportunity loss when assigned to a patient with a AB^- blood, when the minor antigen phenotypes are very similar.

As mentioned in the previous chapter, we think it is better to explicitly differentiate between major and minor antigens, as both influence different objectives. Major antigens determine whether matches are possible which directly relates to the number of shortages in an assignment. On the other hand, minor antigens do not affect shortages but affect the alloimmunization risk. Therefore, in an online approach we want to limit the issuing of highly compatible blood products but by issuing with two different penalty functions. One the one hand we want to discourage AB-RhD substitution to prevent shortages and on the other hand we want to discourage minor antigen substitution to prevent future mismatches.

6.2.2.1 Limiting AB-RhD substitution AB-RhD substitution makes the inventory more flexible as opposed to when only AB-RhD identical matches were allowed. However, we must take care to prevent excess substitutions which may improve the minor antigen matching but lead to shortages in the long run. Because there is a discrepancy in the distributions of major blood groups in the donor population and the general patient population, some substitution will have to take place in order to satisfy all demand. Because the shortage penalty will dominate the objective function these necessary substitutions will always be allowed. However, to make sure that when satisfying demand no more substitution than necessary is used we will add a term to the objective function that penalizes AB-RhD substitution. We call this measure the *Absolute Usability Difference (AUD)*, as the penalty of each unit-patient match is computed as the difference between the AB-RhD usability of both blood types. Mathematically:

Definition. Let $\mathcal{A}_3 = \{A, B, D\}$ denote set of major antigens. The *Absolute Usability Difference (AUD)* of transfusing patient i with RBC unit j is defined as:

$$\Delta U(b_i, d_j) = U(d_j, \mathcal{A}_3) - U(b_i, \mathcal{A}_3) \tag{21}$$

We have chosen an absolute measure instead of a relative measure like ROL, because there is no need to compute the loss in usability relative to the usability of the supplied unit. Usability is an absolute measure that expresses the probability that the unit can be used for a random request from the population. A small opportunity loss for a unit that itself already has low usability does not have to be amplified, because the actual loss is small. This is illustrated with an example below where we compare two substitutions using both the ROL and AUD penalties.

	Substitution 1			Substitution 2		
	A^+	$\xrightarrow{\text{transfusion}}$	AB^+	O^-	$\xrightarrow{\text{transfusion}}$	B^+
Usability	0.4032		0.0336	1.000		0.1092
ROL	$\frac{0.4032-0.0336}{0.4062} = 0.9167$			$\frac{1-0.1092}{1} = 0.8908$		
AUD	$0.4032 - 0.0336 = 0.3696$			$1 - 0.1092 = 0.8908$		

Table 14: Comparison of substitution penalties. Two substitutions ($A^+ \rightarrow AB^+$ and $O^- \rightarrow B^+$) are compared using *Relative Opportunity Loss (ROL)* and *Absolute Usability Difference (AUD)*.

As can be seen, both measures give equal penalties to the $O^- \rightarrow B^+$ substitution. This penalty is relatively high because O^- has much more usability than B^+. For the $A^+ \rightarrow AB^+$ substitution however we see that the ROL penalty is even larger than the $O^- \rightarrow B^+$ substitution. This is because the penalty is computed relative to the usability of the RBC unit. When using the Absolute Usability Difference measure, this problem does not exist. The absolute loss for the inventory capability is 37% and not 92%.

The total AUD penalty as it will be added to the ILP objective looks as follows:

$$\sum_j \sum_i x_{ij} \Delta U(b_i, d_j) \tag{AUD}$$

Because the total sum of the AB-RhD usability of the patient's blood types ($\sum_i U(b_i, \mathcal{A}_3)$) is constant for any assignment, this term is equivalent to minimizing the AB-RhD usability of the assigned units only. We mentioned earlier that some AB-RhD substitution is likely necessary to satisfy all demand due to discrepancies between the donor and patient populations. These substitutions will be allowed because the weight for a shortage is much larger in the objective function than the penalty for a substitution. Therefore the AUD term essentially limits the excess substitutions. However, extra substitution may improve the minor antigen matching quality of an assignment. In Section 6.5.6 we will elaborate further on the influence of the UAD penalty term on the minor antigen matching quality.

6.2.2.2 Limiting Minor Antigen Substitution To minimize minor antigen alloimmunization risk in the long run we will use an approach that focuses on immunogenicity instead of prevalence of the antigens. The reason for this is obvious: we are not explicitly minimizing mismatches, but instead we want to minimize alloimmunization risk. Therefore it is more sensible to directly use the immunogenicity, an indicator of alloimmunization risk, to minimize future alloimmunization risk.

To see how this can be done best, we will first look at the possible outcomes for a match between an RBC unit and a patient for a single antigen:

Blood	Patient		
RBC unit	Antigen	-	+
	-	Identical issuing	Substitution
	+	Mismatch	Identical issuing

Table 15: The four possible unit-patient matching options for a single antigen.

The table shows the result for the four possibilities matches for one antigen. In the previous chapter we only looked at minimizing the lower left term, namely the direct

mismatches. When only this term is minimized over all matches the model cannot differentiate substitution (negative to positive) from identical issuing (negative to negative). These options correspond to the top row of the table. The result of this is that antigen substitution is not penalized. When too many minor antigen substitutions are allowed it might cause a decrease in the availability of antigen-negative blood. When this is the case, we expect more mismatches as there is not enough antigen-negative blood to satisfy all antigen-negative requests. Therefore we also refer to a substitution as an *indirect mismatch*, because it is likely to cause a future mismatch.

To counteract this effect, we add a term to the objective that penalizes minor antigen substitution. This corresponds to the top right combination in Table 15. The term we add is the sum over all matches and all antigens where the RBC unit does not contain the antigen and the patient does. Mathematically it looks as follows:

$$\sum_j \sum_i \sum_k x_{ij} b_i(k)(1 - d_j(k))a(k) \tag{MAS}$$

We multiply $b_i(k)$ with $(1-d_j(k))$, because the result of this multiplication is only equal to one if the patient is positive for antigen k while the RBC unit is negative. We will refer to this term as the *Minor Antigen Substitution (MAS)* penalty.

Now that we penalize direct mismatches and substitution, we achieve a similar goal as relative opportunity loss: When a patient is matched with one or more RBC units with an identical phenotype to the patient, the MAS penalty is zero. When the RBC units are negative for antigens that the patient is positive for, the MAS penalty increases weighted by the relative immunogenicity of the substituted antigens. This is different from the Relative Opportunity Loss penalty, where the substitutions penalties are weighted according to the prevalence's of the substituted antigens. Our approach results in the more immunogenic antigens being preferably identically matched while the less immunogenic antigens will be more easily substituted or mismatched, which should minimize the most severe mismatches in the long run.

To illustrate this approach we give a small example.

Example We have one patient and two units available, shown schematically in Table 16

	Major	C	c	E	e	K	k	Fy(a)	Fy(b)	Jk(a)	Jk(b)	M	N	S	s
Patient	A+	+	+	+	+	-	+	+	+	+	+	+	+	+	+
Unit 1	O-	-	+	-	+	-	+	+	+	-	-	+	+	+	+
Unit 2	O+	+	+	+	+	+	+	+	-	+	+	+	-	-	+

Table 16: One example patient and two example units with their extended phenotypes shown.

At first sight it is not clear which unit should preferably be used. Both units are AB-RhD compatible, but Unit 1 is Rhesus negative and thus more compatible than necessary. On the other hand, Unit 2 induces a direct mismatch on antigen K, which is the most immunogenic antigen. Lastly, we can examine the minor antigen substitutions of both units. We then find that Unit 1 is negative for antigens C, E, $Jk(a)$ and $Jk(b)$ where the patient is not, while Unit 2 is only overcompatible for antigens $Fy(b)$, N and S where antigen N has an immunogenicity of 0 and can therefore be ignored. When all these mismatches are multiplied with their respective relative immunogenicity, we can estimate which unit is a better match. These results are summarized in the table below:

	AUD	Mismatches		Substitutions		Total Penalty
		Antigens	Penalty	Antigens	Penalty	
Unit 1	0.6010	-	0	C, E, Jk(a), Jk(b)	0.3612	0.6010 + 0 + 0.3612 = 0.9622
Unit 2	0.4330	K	0.384	Fy(b), N, S	0.0262	0.4330 + 0.384 + 0.0262 = 0.8434

Table 17: Summarized comparison of penalties for Unit 1 and Unit 2 when transfused to the example Patient. For each unit we show the AUD penalty, the penalty of the direct antigen mismatches, the MAS penalty (antigen substitutions) and the combined total penalty.

As the total penalty for Unit 2 is lower than that for Unit 1 we see that all in all, Unit 2 is preferred for issuing to the patient. This is mostly because it is less 'overcompatible' compared to Unit 1. Unit 1 is a useful unit worth saving as it is AB-RhD compatible for all patients and furthermore negative for $K, E, Jk(a)$ which are the three most immunogenic antigens. In this example, this outweighs one mismatch on antigen K which is induced by Unit 2.

6.3 MINRAR-Online

We can now add all the penalty terms we have previously discussed to the MINRAR ILP Formulation we have constructed in the previous chapter. We will call this new variant of the MINRAR formulation *MINRAR-Online* as we have adapted the original MINRAR formulation to perform in an online setting by including the FIFO, UAD and MAS terms in the objective.

Below we first show the notation used for the different parameters. We have introduced one new parameter r_j that represents the age of RBC unit j.

Notation	**Meaning**
b_i	Blood phenotype of patient $i \in [1, ..., n]$
d_j	Blood phenotype of RBC unit $j \in [1, ..., m]$
$a(k)$	Relative immunogenicity of antigen $\mathcal{A}(k)$
$b_i(k)$ or $d_j(k)$	1 if b_i or d_j is positive for antigen $\mathcal{A}(k)$; 0 otherwise
u_i	number of RBC units requested for patient i
c_{ij}	AB-RhD compatibility of unit j and patient i
r_j	age of RBC units j

Table 18: Notation for blood

$$\min \sum_i s_i \cdot (n+1) + \sum_i \sum_k y_{ik} \cdot a(k) + \sum_i \sum_j -x_{ij} \cdot o(35 - r_j)$$
$$+ \sum_j \sum_i x_{ij} \Delta U(b_i, d_j) + \sum_j \sum_i \sum_k x_{ij} b_i(k)(1 - d_j(k))a(k) \tag{22}$$

s.t.

$$\sum_j x_{ij} + s_i \cdot u_i = u_i \ \ \forall i \tag{23}$$

$$\sum_i x_{ij} \leqslant 1 \ \ \forall j \tag{24}$$

$$u_i y_{ik} \geqslant \sum_j x_{ij} d_j(k) \ \ \forall i, k \ \ \textit{if} \ b_i(k) = 0 \tag{25}$$

$$x_{ij} \leqslant c_{ij} \ \ \forall i, j \tag{26}$$

$$x_{ij}, y_{ik}, s_i \in \{0, 1\} \ \ \forall i, j, k \tag{27}$$

The structure of this ILP is the same as the MINRAR ILP in the previous chapter. Because the constraints are left unchanged the model has the same solution space and therefore the hard constraints discussed earlier are still honoured. The difference lies in the objective function. Originally, the objective only consisted of two terms: minimizing the number of shortages and minimizing the direct alloimmunization. These correspond to the first two terms in the objective above. This is now extended with three extra terms which maximize the age of the units issued (FIFO), minimize the AB-RhD usability of the issued units (UAD) and minimize the number of minor antigen substitutions, weighted by their relative immunogenicity (MAS).

6.3.1 Solving the Model

Because the model is a general Integer Linear Program it can be solved by all generic ILP solvers. Ideally, the optimal solution should be used as it minimizes the sum of all objective terms mentioned above. However, when solving for large instances under a time constraint, intermediate sub-optimal solutions can be used if they are feasible.

6.4 Computational Experiments

We will use simulation experiments to investigate the long-term performance of the proposed MINRAR-Online model. The performance will be measured using three indicators: average number of shortages, average number of outdates and average relative alloimmunization risk per patient. We will compare the performance of the MINRAR-Online model against the FIFO/MROL model of Van Sambeeck *et al.* [6]. Their FIFO/MROL model has an objective function that minimizes a combination of FIFO cost and Relative opportunity loss. It further requires that before execution a set of antigens is chosen on which all matches must be compatible. As it is unclear which subset of antigens will perform optimally and enumerating all subsets would be exponential in the number of antigens (2^{14} for 14 antigens), we will use similar subsets as the authors. These are the following:

$$\mathcal{A}_3 = \{\text{A,B,D}\}$$
$$\mathcal{A}_8 = \{\text{A,B,D,C,c,E,e,K}\}$$
$$\mathcal{A}_{14} = \{\text{A,B,D,C,c,E,e,K,Fy(a),Fy(b),Jk(a),Jk(b),M,S}\}$$

Antigen set $\mathcal{A}_3$ only includes the three most important antigens for which currently all patients receive compatible matching. The performance of this set will give a good indication of the performance of an issuing strategy with no extended antigen matching. The second antigen set $\mathcal{A}_8$ extends the first one with the four remaining Rhesus antigens (C, c, E, e) as well as K. The total relative alloimmunization of these antigens combined is 81.27%, meaning that any allowed match using this variant can have a mismatch penalty of at most 18.73%. In practice this value will be lower as for both the Duffy and Kidd blood group systems a mismatch on both antigens is extremely unlikely as in both cases the double-negative phenotype is very rare among Caucasians. The last variant considered ($\mathcal{A}_{14}$) uses all antigens with nonzero relative immunogenicity (as well as A, B and D which have a relative immunogenicity of zero as they may never be mismatched). Using this antigen set, all mismatches that may lead to any relative alloimmunization risk are prohibited. Therefore, all patients receive compatible products on all these antigens. We do expect however the number of shortages to be high.

The three variants described above will be referred to as FIFO/MROL $\mathcal{A}_3$, FIFO/MROL $\mathcal{A}_8$ and FIFO/MROL $\mathcal{A}_{14}$. These three models are solved by reducing the extended blood matching problem to a min-cost-max-flow problem. In the experiments we will use the MINROL-Online ILP model as described in Section 6.3 and we will include all antigens with nonzero alloimmunization. To be able to make a fair comparison between the different issuing strategies we will test them on fixed supply and demand scenarios. How these are generated is explained in the sections below.

6.4.1 Supply Scenarios

The stochasticity of blood supply is determined by the echelon of the inventory in the blood supply chain. Hospitals, for example, have predetermined order-up-to levels for each AB-RhD blood group and the units that are ordered from the distribution centre are almost always available for supply. However, the supply of units to a distribution centre is already much less certain. Distribution centres cannot place AB-RhD specific orders as they are not supplied by another bigger distribution centre. Instead, blood supply chain managers have to ensure that enough donors are invited for donation to keep a steady supply for all AB-RhD blood types.

Because the MINRAR-Online model can be used on any scale, it is not immediately clear whether we should use a relatively deterministic model of supply, resembling a hospital inventory, or a more stochastic type of supply as seen on regional or national level. For simplicity we have chosen for random order-up-to supply, meaning that at the end of every day the inventory is replenished with random units up until a fixed maximum inventory size M. This type of supply has two advantages. Firstly, it is the simplest type of supply and therefore the performance in the simulation is not dependent on specific order-up-to levels or other more complex methods of supply. Secondly, because random order-up-to supply does not allow explicit orders for each individual AB-RhD blood type, it allows us to better test the highly compatible blood saving capabilities of the different issuing strategies. If we were to use AB-RhD order up-to-levels, then an issuing strategy that is not well able to preserve rare blood groups like O^- is rewarded as a low stock for these blood groups is immediately compensated for in the supply of the next day. A possibility would be to allow this and present aggregate results of the supplied units, but this makes an equal comparison on the minor antigen matching difficult as when more Rhesus D negative units are supplied it also means that more C and E negative units are supplied due to Rhesus antigen linkage. This would make it more difficult to draw clear conclusions from the results and therefore we have chosen not to use this type of supply.

In a supply method of random order-up-to supply we assume that at the end of each day the inventory is supplied with exactly enough units such that the total inventory size is again equal to M. Although it is unrealistic that the supply of each day is exactly enough to replace the issued units in a distribution centre, it does make sure that the outdating percentage is a direct result of the issuing strategy and not caused by oversupply. Lastly, our supply model does not contain any supply shocks that can stress-test the inventory. However, it is not the scope of this research to investigate the availability of blood under abnormal circumstances.

To model the random order-up-to supply, we have generated several supply scenarios. A supply scenario is a large queue of random blood types, with each item corresponding to a single RBC unit. When the inventory needs to be replenished with m units, the first m units in the queue are supplied and removed from the queue. The major blood types of the units are sampled using the known distribution of the eight AB-RhD blood types in the donor population as shown in Figure 5. The remainder of the phenotype of each RBC unit is randomly generated based on the major blood group and minor antigen blood group prevalences among Caucasians, which can be found in Appendix A. As donors are not selected in any way on their minor antigen phenotype, we can assume that the prevalence of minor antigens in the donor population is the same as in the general Caucasian population. Furthermore, we have not explicitly modelled repeated donations of the same donor. Because of the large number of donations, we do not expect that this simplification will introduce a large bias in the sampled supply scenarios. Lastly, the supply scenarios are generated before the simulations such that we can compare different issuing strategies on the same supply scenario.

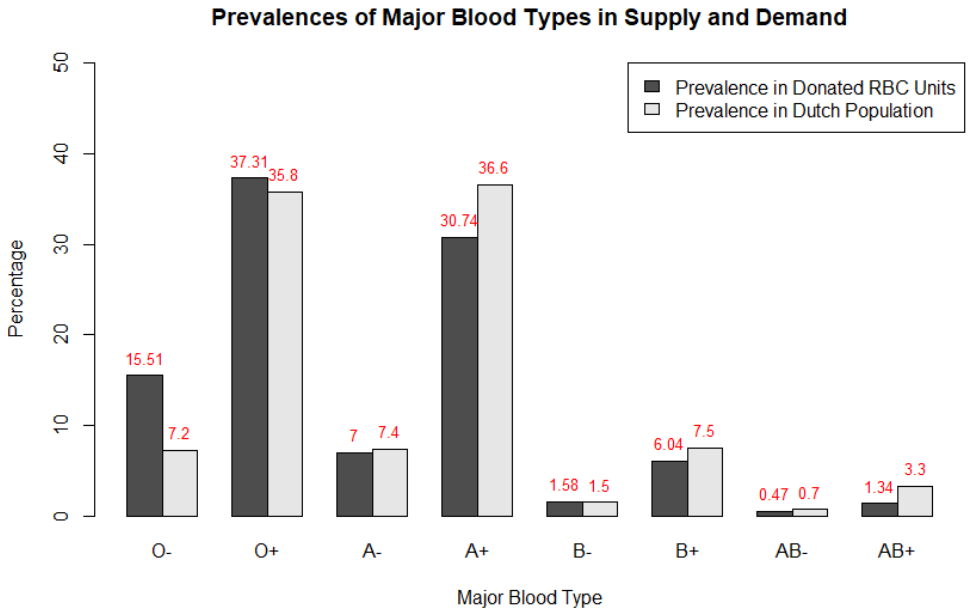

Figure 5: AB-RhD blood type prevalence of donated RBC units and the Dutch population.

6.4.2 Demand Scenarios

Demand for RBC units is also to a large extent stochastic. Similar to the supply scenarios we have also generated demand scenarios. A demand scenario specifies for each day the patients who demand blood. For each patient we store their extended phenotype and the number of units demanded. The phenotypes of the patients are generated at random based on the prevalence of different blood groups among Caucasians (Appendix A). The number of units per patient is sampled according the distribution shown in Section 3.1.3 where we only include requests up to four units. Furthermore, we have not explicitly modelled recurring request for individual patients. To estimate the number of patients per day, a dataset containing the total daily number of units demanded in the Netherlands from 2009 up to 2019 was used. Because the daily demand has a cyclic pattern of 7 days, we have estimated a distribution for each day of the week. To estimate these distributions correctly we first removed a slight decrease trend in the data, showing that the average daily demand for weekdays decreased with 60.41 ($p < 2.2 \cdot 10^{-16}$) units per year. The demand for Saturdays decreased with 14.57 ($p < 2.2 \cdot 10^{-16}$) units per year and on Sundays there was no decrease observed. Then we removed all holidays and other drops in demand that showed near-zero demand for a day. The resulting empirical distributions per day were approximated using the method of Adan *et al.* [54] which fits a discrete distribution preserving the mean and variance of the empirical distribution. The result is a mix of two negative binomial distributions per day of the week, shown in Figure 23 which can be found in Appendix B. The average national demand is roughly 1500 units per day. To be able to simulate smaller inventories we use scaled-down versions of the fitted distributions. Because the national distributions are the sum of individual demand distributions for hospitals it is not possible to sample the demand for a single hospital by first sampling the national distribution and then dividing the number of units sampled by the difference in size between the distributions. The reason for this is that the national distributions have relatively less variance, because they are the sum of smaller individual distributions. Under normal circumstances it is for example very unlikely that all individual hospitals have a demand spike on the same day. To still be able to sample demand for smaller inventories we will use the assumption that the national distributions above are the sum of a number of *identical and independent* smaller distributions. Clearly, this does not exactly reflect the practical situation, but in absence of more concise data it is the most general assumption that allows sampling smaller demand.

When we want to sample a distribution for a given day for a given average demand $\mu_{individual}$ based on the corresponding national distribution for that day with $\mu_{national}$ and $CV_{national}$, we first calculate the factor in size difference: $\alpha = \frac{\mu_{national}}{\mu_{individual}}$. Now we compute the coefficient of variation for the individual distribution as $CV_{individual} = \sqrt{\alpha} CV_{national}$. This is because when α independent and identical distributions with variance σ^2 are summed, the resulting distribution has α times the original mean but $\frac{\sqrt{\alpha}}{\alpha}$ times the variance.

When we have obtained the mean and coefficient of variation for the individual scaled down distribution, we can use the same procedure as used earlier [54] to estimate the distribution preserving the first and second moment.

6.4.3 Simulation Setup

To test an individual issuing strategy, we perform a one-year simulation to estimate its long-term average performance. For a single simulation we have the following parameters: inventory size M, demand scenario $\mathcal{D}$, supply scenario $\mathcal{S}$, duration T and initialization period T_0.

The simulations were started by filling the inventory with units from the supply scenario. All units have 35 days of remaining shelf life when they enter inventory. We simulate an inventory where requests must be assigned with units on a daily basis. This means that per day, all requests become known simultaneously and they do not have a specific time of day assigned to them. The daily procedure will look as follows:

1. **Obtain requests for the current day**. From the demand scenario we retrieve all requests that have their due date equal to the current day. Per request we have the following information: Number of units requested and phenotype for the blood groups ABO, Rhesus, Kell, Duffy, Kidd and MNS.

2. **Construct the ILP model.** We build the MINRAR-Online ILP model as described in Section 6.3.

3. **Solve the model.** The solver is run to compute the optimal solution to the ILP model. When an optimal solution is found we read the values of the x_{ij} variables to see which assignments are made as well as which requests are left unsatisfied.

4. **Issue units and log statistics.** The assignment computed in the previous step is used to issue units to the patients and remove them from inventory. Furthermore, if the initialization period has passed various statistics are logged such that these can later be used for analysis.

5. **Increase age of remaining units.** The age of all unissued units is increased with one day. Units that have an age of 35 days are removed from inventory and marked as 'expired'.

6. **Replenish inventory.** From the supply scenario new units are added to replace all issued or outdated units such that the total new inventory size is equal to the original size. These new units have 35 days of remaining shelf life.

7. **Increase day count by one.** The day count is increased by one and we start at step 1 one again.

This process is repeated until the end of the demand scenario. The duration of each simulation was 396 days. This was done to allow for an initialization period of 31 days at the start of the simulation. The reason for this initialization period is to allow the issuing strategy to reach a steady state, where the performance is no longer influenced by initial conditions.

In all simulations we have assumed an inventory size equal to five times the average daily demand. The reason for this is that this is the ratio of inventory size to average daily demand that is most common in Dutch distribution centres. In hospitals the inventory size is usually smaller compared to the average daily demand. The reason that this is possible is that most hospitals use AB-RhD specific order-up-to levels, which allow the inventory to be stabilized every day. We have previously discussed that we will not use AB-RhD specific order-up-to levels in our experiments to be able to better assess to compatible blood saving capabilities of the issuing strategy. When hospitals do not have the luxury of using AB-RhD specific order-up-to levels then a larger inventory size is to be expected to be able to cope with irregularities in supply and demand.

All simulations were run on a laptop running Windows 10, with an Intel Core i5-5200U CPU (Dual Core, 2.20GHz) CPU and 8 GB RAM. The simulations were implemented in C# and Gurobi 9.0.1 was used to solve both the MINRAR-Online ILP as well as the FIFO/MROL formulations. Unless stated otherwise the ILP models were

solved to optimality with an optimality gap of $1.0 \cdot 10^{-4}$, meaning that a solution was declared optimal if the value of the objective function was within 0.01% of the highest found lower bound.

6.5 Results

In this section we present the results of the simulations of the four considered issuing strategies. We have tested the performance of the different issuing strategies for five values of average daily demand. These are 25, 50, 100, 200 and 500 units per day. A demand of 25 units per day, corresponds to a small hospital. A large academic hospital has a demand of about 100 RBC units per day. Average daily demands for 200 and 500 units per day correspond to a small and large distribution centres. In each case, the inventory size is equal to five times the average daily demand. Each issuing strategy was tested on the same 25 combinations of five supply and five demand scenarios per value for average daily demand. We have constructed various figures to highlight where the strategies differ.

6.5.1 Overall Performance

The overall average results in terms of shortages, outdates and relative alloimmunization risk are shown numerically in Table 19 as well as graphically in Figure 6. The performance of the four issuing strategies considered was compared using four performance indicators: number of unsatisfied/partially satisfied patients, number of units short, outdates and relative alloimmunization risk per patient. The reason that we use two indicators for the number of shortages is that the MINRAR model uses a different method of defining shortages than the FIFO/MROL models. The MINRAR model minimizes the *Patient Shortages*, for which a shortage is defined as a patient not receiving the requested number of units, irrespective of the number of units requested. The FIFO/MROL models do not explicitly model patients that can have a demand for multiple units. Instead, these requests are split into singular requests and then a shortage is defined as the failure to satisfy one of these singular requests with blood. This is called a *Unit Shortage* and it was not computed for the MINRAR model as it is not optimized for and more importantly, not relevant. Furthermore, we know that when the MINRAR model has zero percent patient shortages, the number of unit shortages is automatically also zero. Similarly, when the FIFO/MROL models have zero unit shortages, the number of patient shortages must also be zero.

For the simulations with an average daily demand of 500 RBC units we have solved the MINRAR-Online ILP only to 1% optimality. For most instances the solver was able to quickly find an optimal solution within the default optimality gap ($1.0 \cdot 10^{-4}$). However occasionally some instances turned out to be hard to solve optimally and took considerable time. Because we wanted to simulate 25 one year (=365 + 31 days) runs, we have chosen to accept a solution if it was within 1% of the best known lower bound of the optimum. Therefore, it may be that in practice the average performance of this issuing strategy is slightly better than the value shown in the table.

	FIFO/MROL $\mathcal{A}_3$				
Average Daily Demand	25	50	100	200	500
Patient Shortages (%)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)
Unit Shortages (%)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)
Outdates (%)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)
Relative Alloimmunization Risk per Patient (%)	17.55 (17.43-17.67)	17.59 (17.52-17.67)	17.55 (17.52-17.58)	17.50 (17.47-17.53)	17.57 (17.55-17-59)
	FIFO/MROL $\mathcal{A}_8$				
Average Daily Demand	25	50	100	200	500
Patient Shortages (%)	8.42 (7.51-9.32)	4.60 (4.43-4.77)	3.94 (3.85-4.03)	3.39 (3.23-3.55)	2.68 (2.62-2.74)
Unit Shortages (%)	7.06 (6.32-7.79)	3.75 (3.61-3.89)	3.24 (3.17-3.32)	2.85 (2.71-2.98)	2.67 (2.61-2.72)
Outdates (%)	0.74 (0.68-0.81)	0.53 (0.47-0.60)	0.17 (0.13-0.20)	0.05 (0.04-0.07)	0.00 (0.00-0.00)
Relative Alloimmunization Risk per Patient (%)	4.19 (4.17-4.20)	4.20 (4.18-4.21)	4.20 (4.19-4.22)	4.19 (4.18-4.20)	4.19 (4.19-4.19)
	FIFO/MROL $\mathcal{A}_{14}$				
Average Daily Demand	25	50	100	200	500
Patient Shortages (%)	49.77 (49.33-50.21)	39.38 (38.94-39.83)	29.24 (28.92-29.57)	20.91 (20.71-21.11)	13.32 (13.23-13.40)
Unit Shortages (%)	46.32 (45.92-46.72)	36.19 (35.74-36.64)	26.50 (26.20-26.80)	18.72 (18.54-18.89)	11.71 (11.64-11.79)
Outdates (%)	5.66 (5.53-5.79)	3.41 (3.30-3.51)	2.00 (1.96-2.04)	1.03 (1.00-1.06)	0.48 (0.46-0.51)
Relative Alloimmunization Risk per Patient (%)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)
	MINRAR-Online				
Average Daily Demand	25	50	100	200	500
Patient Shortages (%)	0.01 (0.01-0.02)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)
Unit Shortages** (%)	-	-	-	-	-
Outdates (%)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)
Relative Alloimmunization Risk per Patient (%)	4.04 (3.98-4.10)	3.19 (3.16-3.22)	2.37 (2.36-2.39)	1.66 (1.65-1.67)	1.30* (1.29-1.30)

Table 19: Performance of the four issuing strategies considered. Values shown are % (95% CI) averages over 25 one-year simulations with fixed combinations of supply and demand scenarios. Five different values of average daily demand are tested and for each combination we report the patient shortages, unit shortages, outdates and relative alloimmunization risk per patient.

*Due to time constraints the model was not solved to optimality but instead, all feasible solutions with an optimality gap of 1% were accepted. A 1% optimality gap means that optimization was terminated when a feasible solution was found that was within 1% of the highest found lower bound.

**The unit shortages were not computed for the MINRAR-Online model as it is not possible to compute these, since they were not optimized for. Furthermore, whenever the number of patient shortages is zero, the unit shortages will also be zero.

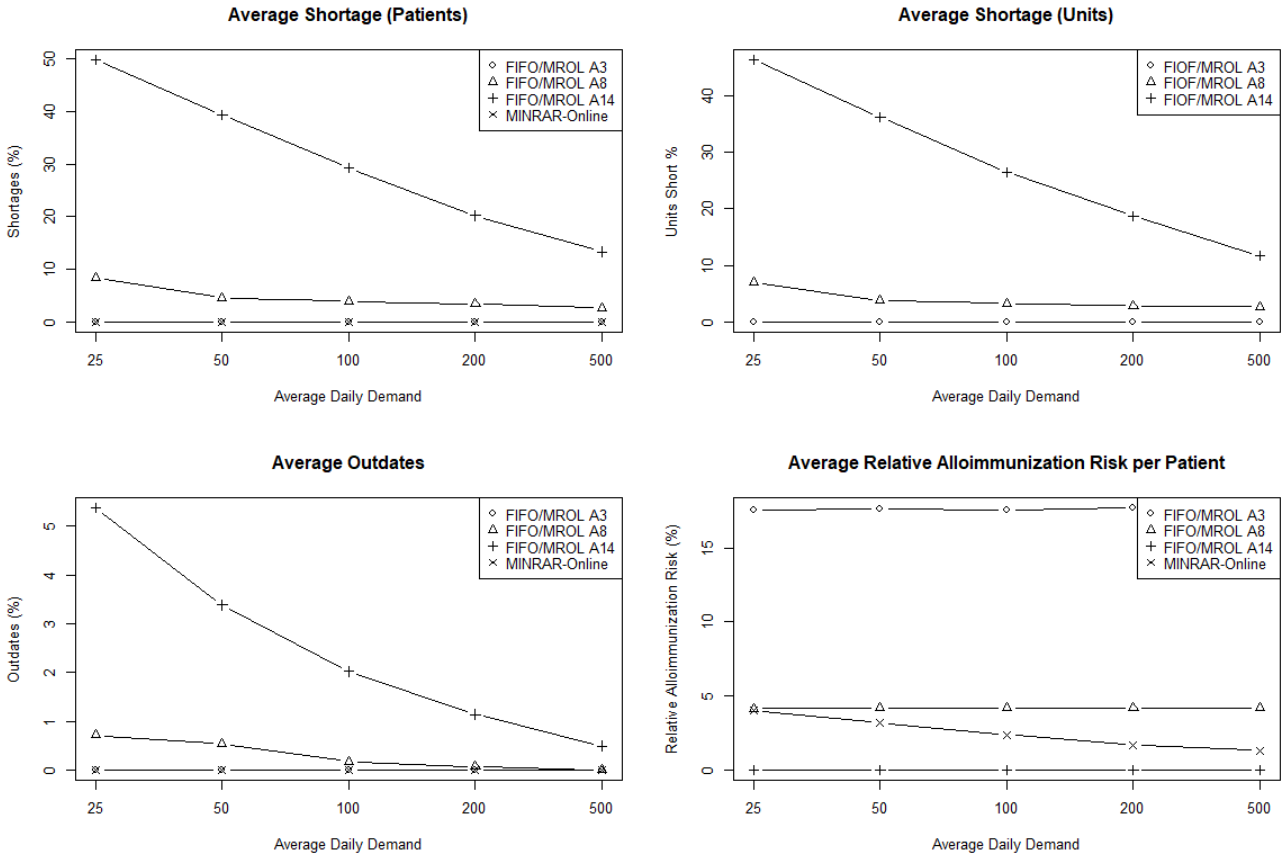

Figure 6: Results of 25 one-year simulations for five different values of average demand. In every case the inventory size is five times the average daily demand. The four issuing strategies considered are compared on three performance indicators: Average percentage of shortages (patients), average percentage of outdates, average relative alloimmunization risk per patient. Additionally, we show the shortage percentage in terms of units which we can only compute for the FIFO/MROL issuing strategies.

6.5.2 Inventory Composition

The composition of the inventory is directly related to the issuing strategy. Any issuing strategy has to work with the units present in inventory and which units are present is influenced by the issuing on the previous day and the new units that have been supplied. As the distribution of antigens in the supplied units is the same for every issuing strategy, we can look at the average inventory composition to see the effect that the particular issuing strategy has on the prevalence of the different antigens in inventory.

6.5.2.1 Antigen Prevalence In Figure 7 we show for each antigen the average fraction of units in inventory that express this antigen. This was calculated at the start of each day by counting for each antigen the number of positive units. After the simulation these values were divided by the duration of the simulation and by the total inventory size for normalization. In Figure 7 we also show the prevalence of the antigens over all units supplied. This value, denoted as "supply", gives an indication of the prevalence of the different antigens in the units as supplied and can therefore be used as a reference point to see whether an issuing strategy increases or decreases the prevalence

of this antigen in inventory.

We can see that using the MINRAR-Online ILP the prevalence of the antigens A, B, D, C, E and K is reduced significantly compared to the prevalence among the supplied units. This indicates that the MINRAR-Online ILP favours issuing these antigens, thereby keeping units that lack these antigens in stock. This is to be expected as the MINRAR-Online ILP explicitly tries to reduce AB-RhD substitution as well as the substitution of the more highly immunogenic antigens. It may be surprising that some other antigens, like c, have an increased prevalence. This can be explained by the linkage of the Rhesus antigens: Because the MINRAR-Online ILP heavily favours saving Rhesus-D negative blood, we can look at the conditional probability of a phenotype with antigen c given the absence of antigen D. For all practical purposes this probability (among Caucasians) is 100%, as all variants that lack antigens c and D have (near) zero probability of occurrence.

6.5.2.2 Issuing Age Another way to get an idea of the inventory composition resulting from a particular issuing strategy is to look at the average issuing age of the RBC units. More specifically, we look at the average issuing age of the units for each of the eight major blood groups individually. When a particular blood group has a high average issuing age it implies that a relatively large part of the inventory consists of units with this blood type. Similarly, a low average issuing age means that units with this blood group are issued fast and therefore occupy a relatively small fraction of the entire inventory. An overview of these values is shown in Figure 8. In this figure we show for three values of average daily demand the average number of days that units stayed in inventory for each of the four different issuing strategies.

As mentioned earlier, the MINRAR-Online ILP heavily penalizes AB-RhD substitution and this is also reflected in the figure when we see that the average issuing age of O^- blood is the highest of all major blood groups in the MINRAR-Online model. Furthermore, we see similar behaviour for the FIFO/MROL $\mathcal{A}_3$ strategy, which only focuses on antigens A, B and D. Next, we can see that when the average demand and corresponding inventory size are small (25 and 125 units respectively) the FIFO/MROL $\mathcal{A}_8$ and FIFO/MROL $\mathcal{A}_{14}$ issuing strategies have an accumulation of low usability blood groups in inventory such as AB^- and AB^+. This is to be expected when the absolute inventory size is small, as we expect more substitution is needed to satisfy demand. Furthermore, because these strategies use the relative opportunity loss measure with more antigens than A, B and D, the saving of more usable blood groups as O^- and $O+$ is not as explicitly prioritized anymore.

When daily demand and inventory size increase, we see that all strategies reduce the issuing age of the less usable blood groups. For the MINRAR-Online ILP we see that the average issuing age of O^- units increases to roughly three weeks, while all other units have an average issuing age of less than five days. This effect is also observed for the FIFO/MROL strategies with the exception of FIFO/MROL $\mathcal{A}_{14}$, which struggles to accumulate these units in inventory. The reason for this is that we still see many shortages with this strategy, even when the average daily demand is 500 units. Having more shortages means issuing less units and thus the average issuing age will increase.

Finally, it may appear as if some issuing strategies have an average issuing age over all units which is higher than five days, which would not be possible when the inventory size is five times the average daily demand. The reason that we still see this is because with many shortages there will be less units issued and therefore also less new units supplied per day. When the inventory size is constant but the average number of units supplied per day decreases, the average issuing age increases as the same units must be kept in inventory longer.

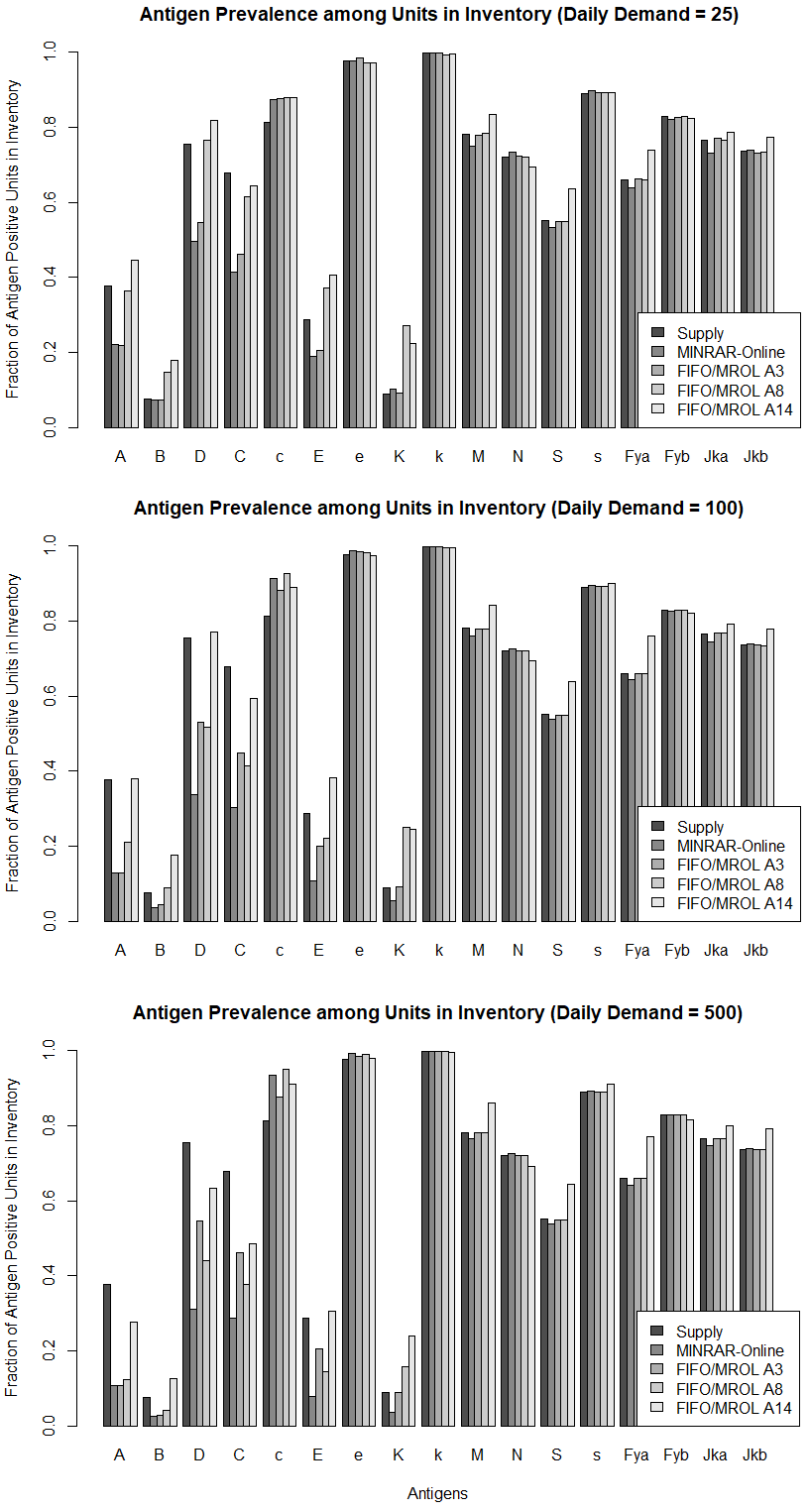

Figure 7: The bar plots show the average proportion of RBC units in inventory that are positive for the different antigens. These proportions are the result of the issuing strategy and how well it can prevent the accumulation of antigen-positive units. The left most bar of each column shows the prevalence of each antigen in the donor population. This value can be used as a reference to see if the issuing strategy has a higher, lower or the same proportion of antigen positive units as the donor population.

Figure 8: Average issuing age for all eight major blood groups for the issuing strategies considered. The average is computed over all units issued, while outdates units are ignored.

6.5.3 Relative Alloimmunization Risk per Patient

We have used the relative alloimmunization risk as a measure for determining the quality of a match between a unit and a patient. In the overall performance table with results (Table 19) we have seen that the MINRAR-Online issuing strategy reduces the average relative alloimmunization risk per patient. However, averages do not provide any insight into which mismatches do occur frequently and which do not. We therefore constructed Figure 9 which shows for three values of daily demand the cumulative fraction of patients that is mismatched with a total relative immunogenicity less than or equal to a given value. On the x-axis we show the relative immunogenicity where 0 corresponds to a compatible match on all antigens with nonzero relative immunogenicity. To make these relative immunogenicity values easier to interpret, we added vertical lines to indicate the relative immunogenicity values of all antigens with nonzero relative immunogenicity. Furthermore, we included a subplot where we show an amplified version of the most left section of the figure as most antigens have a relative immunogenicity between 0 and 0.1.

First of all, we can observe that the cumulative line for all FIFO/MROL issuing strategies is the same for all three values of average daily demand. This is because these strategies use a predetermined set of antigens for which every patient must be compatibly matched. The remaining antigens are ignored and therefore the quality of matching is determined by chance, as no matching is done for these remaining antigens. The difference in issuing strategies can most clearly be seen in the first figure where the average daily demand is lowest (25 units). When we compare the MINRAR-Online ILP to the FIFO/MROL $\mathcal{A}_8$ issuing strategy, we see that initially the MINRAR-Online ILP can match more patients with zero mismatch costs. This trend continues until a mismatch cost of about 0.09, at which point the MINRAR-Online ILP can match fewer patients with the given mismatch cost. However, we must note that these values are only computed over the patients who are assigned units. Patients who are not matched any units are not included in this graph. Now we can also clearly see why the MINRAR-Online ILP has fewer shortages than the FIFO/MROL models. This is because we allow severe mismatches (on the Rhesus and K antigens) if they cannot be avoided. In the figure we can see that the fraction of all patients who receive such heavy mismatches is small, especially when the average daily demand is increased. The fact that for different values of average demand the cumulative line for the MINRAR-Online ILP differs exactly reflects the flexibility of the MINRAR-Online ILP compared to the FIFO/MROL strategies.

6.5.4 Major Blood Group Substitution

To investigate the amount of substitution of the eight different major blood groups we have constructed substitution matrices for the four issuing strategies considered in Figures 10 and 11. Figure 10 shows how the demand for each major blood group was satisfied. Each row represents the total demand for a certain blood group, and every cell in that row is the percentage of the total demand that was satisfied by the blood group corresponding to the column. For instance, all O^- demand must be satisfied with O^- RBC units and therefore in every matrix we see that the combination $O^- \rightarrow O^-$ is always 100%. Figure 11 shows the same process but from the supply side. Every column represents the total supply of a certain major blood group. Each cell in this column represents the percentage of supplied RBC units used to satisfy demand with a blood group corresponding to the row the cell is in. For example, we can see that 100% of the supplied AB^+ units were assigned to AB^+ patients, as these are the only patients who can be transfused these units. The values in all matrices are averages over 25 one-year simulations of a 1000 unit inventory with an average daily demand of 200 units. In the figures we can see that overall, in every issuing strategy identical AB-

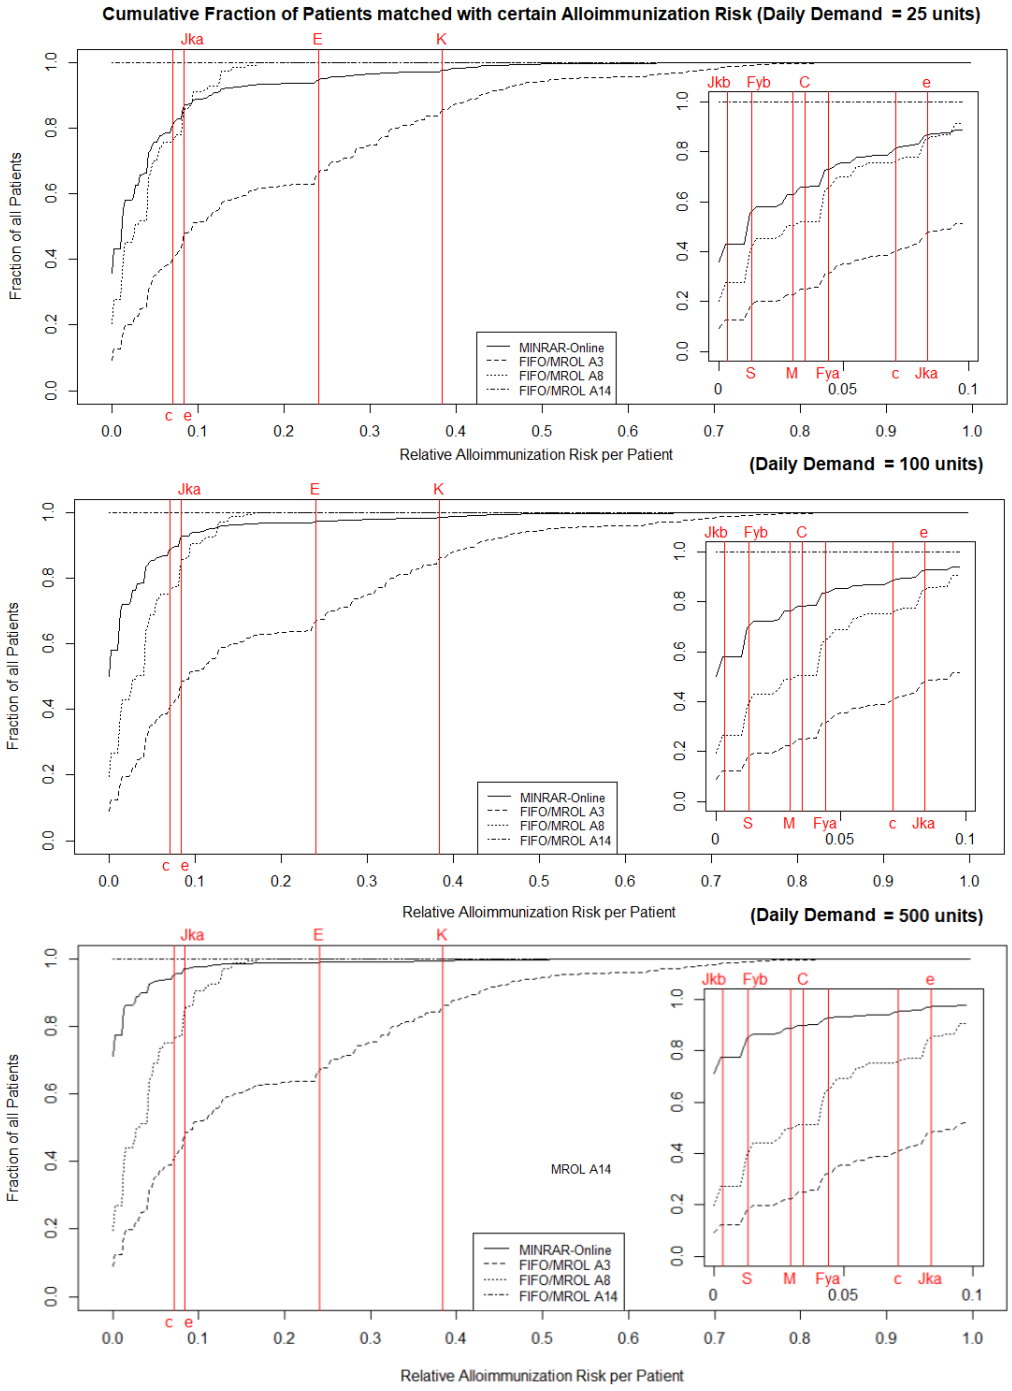

Figure 9: Cumulative proportion of patients matched with alloimmunization equal or less than the shown values. The relative immunogenicity of the antigens considered is shown to give an idea which mismatches correspond to different alloimmunization levels. Furthermore, the plots contain a sub plot which highlights the leftmost parts of the main plots with relative alloimmunization values between 0 and 0.1.

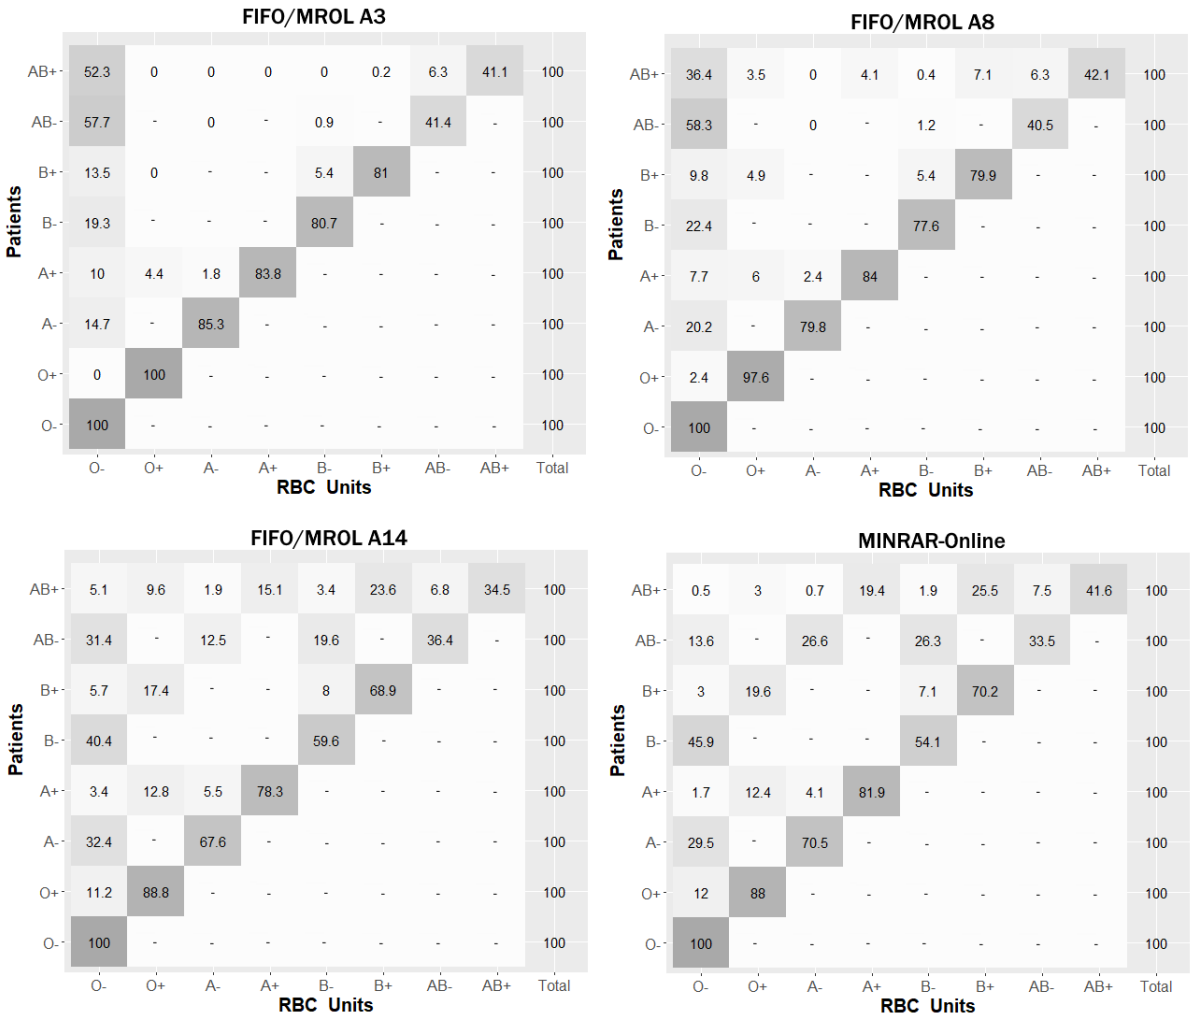

Figure 10: Substitution matrices as seen from the demand side. The values per row show the percentage of demand for the corresponding blood group that was satisfied with an RBC unit with a given major blood group. These numbers were computed by taking the average over 25 one-year simulations all with a daily demand of 200 units and an inventory size of 1000 units. Combinations that are incompatible are marked with -.

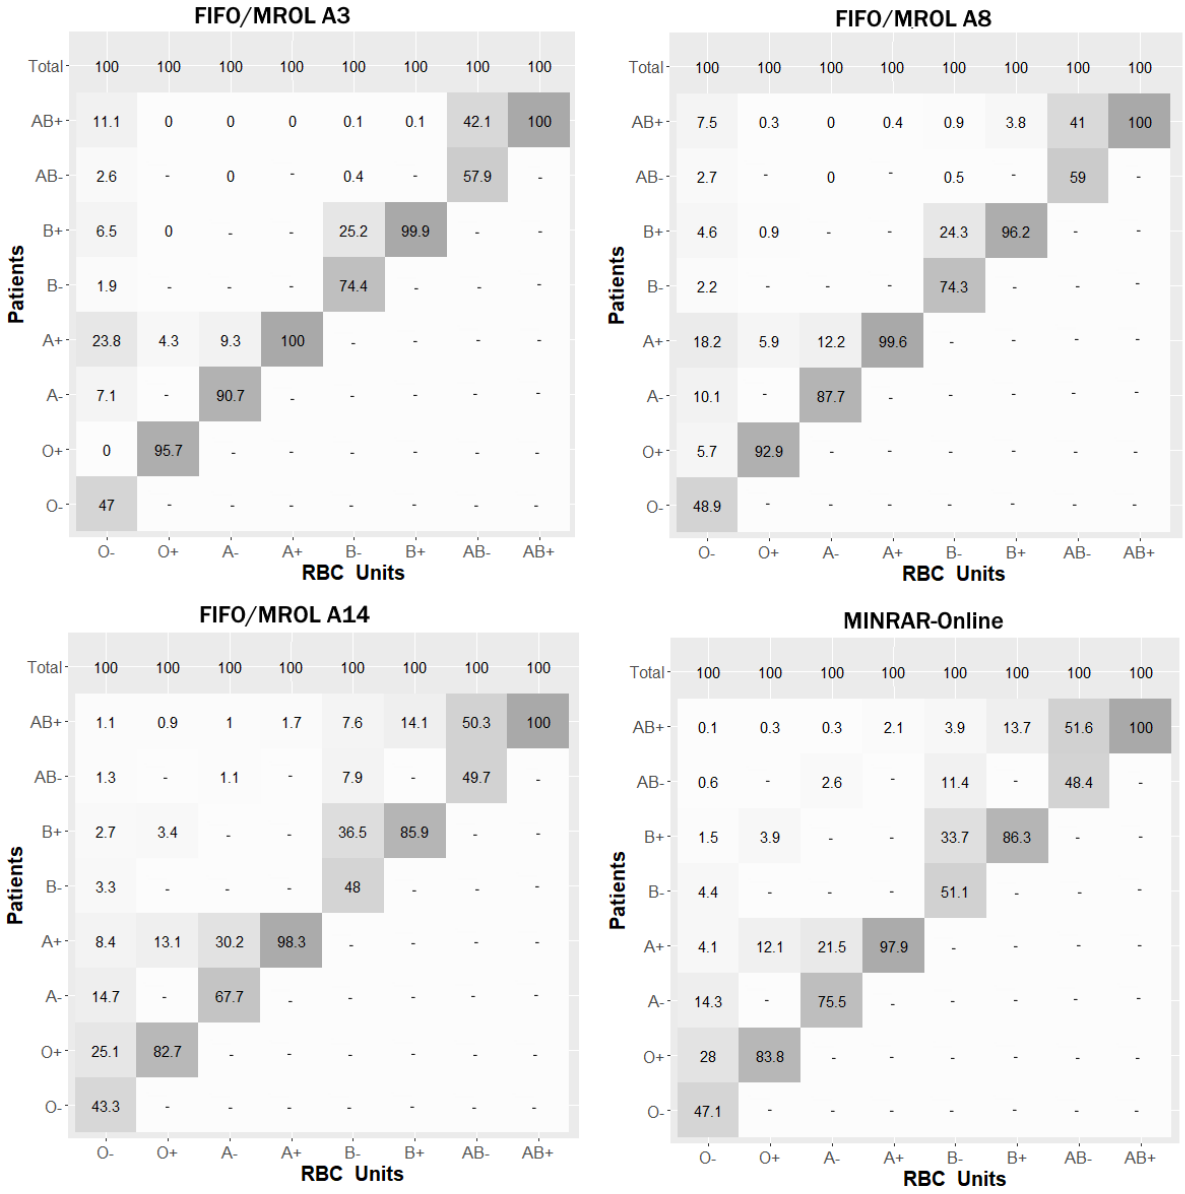

Figure 11: Substitution matrices as seen from the supply side. The values per column show the percentage of supplied units with the corresponding major blood group that were issued to patients of a particular major blood group. These numbers were computed by taking the average over 25 one-year simulations all with a daily demand of 200 units and an inventory size of 1000 units. Combinations that are incompatible are marked with -.

RhD issuing is preferred. Furthermore, Figure 10 shows that when the set of antigens is increased for the FIFO/MROL strategies, the number of substitution also increases. When we compare the substitution in the MINRAR-Online strategy to the FIFO/MROL strategies we see that pattern of substitutions is most similar to the FIFO/MROL $\mathcal{A}_{14}$ strategy. To compare the total number of substitution between the different issuing strategies, we have constructed Table 20. This table show for each strategy the sum of all non-identical substitutions made. We can see that when more antigens are considered in the FIFO/MROL strategies, the percentage of substitutions increases. Ignoring the FIFO/MROL $\mathcal{A}_{14}$ issuing strategy, we see that the MINRAR-Online ILP has the highest percentage of substitutions, which is explained by the fact that the strategy is well able to prevent shortages by substituting when necessary. Furthermore, we see that when the daily demand increases the percentage of substitutions decreases for the FIFO/MROL strategies, likely because the strategy is better equipped to satisfy all demands. On the other hand, we can see a slight increase in the number of substitutions for the MINRAR-Online ILP, which is likely a result of the fact that when more units are available for matching, more alloimmunization risk can be prevented by substitutions.

	Average daily demand 25	Average daily demand 200
FIFO/MROL $\mathcal{A}_3$	13.4% (86.6% identical)	11.1% (88.9% identical)
FIFO/MROL $\mathcal{A}_8$	17.4% (82.6% identical)	12.3% (87.7% identical)
FIFO/MROL $\mathcal{A}_{14}$	30.1% (69.9% identical)	20.0% (80.0% identical)
MINRAR-Online	18.1% (81.9% identical)	18.5% (81.5% identical)

Table 20: Percentage of major blood group substitution for the four issuing strategies considered.

6.5.5 Effect of Penalizing Minor Antigen Substitution

To verify that minimizing substitution alloimmunization risk on top of direct alloimmunization risk is better in the long run, we have run some additional simulations to verify this. The same setup was used as described above, but now the issuing strategies compared are two variants of the MINRAR-Online ILP, one normal and one without the MAS term, responsible for penalizing minor antigen substitution. We used an inventory of size 200 units and five demand scenarios with an average daily demand of 40 units. We used five different supply scenarios. We have simulated one year of inventory management for all 25 combinations of five randomly generated demand scenarios and five randomly generated supply scenarios. The results over these 25 combinations have been averaged and are shown in Figure 12 below.

We know that ignoring the MAS term and only minimizing the alloimmunization caused by the mismatches on the current day is beneficial for the performance on that day. There is no limit on substitution and therefore a solution can be computed that is most suitable for the patients on that day. However, Figure 12 clearly shows that such a strategy is inferior to a strategy with the MAS term, which focuses more on long term performance. This means that for single day in the simulation the minor antigen matching quality can likely be improved by ignoring the MAS term, it is beneficial for the average over all days to include it. Furthermore, we did not observe any significant difference in shortages or outdating when removing the MAS term. Therefore these results are not included in Figure 12. We have not experimented with variations of the weight of the MAS term in the objective. The reason for this is that the penalty of a substitution of an antigen is currently equal to the penalty of the mismatch associated with it. If we were to make the penalty of a substitution larger than the penalty of the corresponding mismatch, then this would lead to strange situations. For example, a

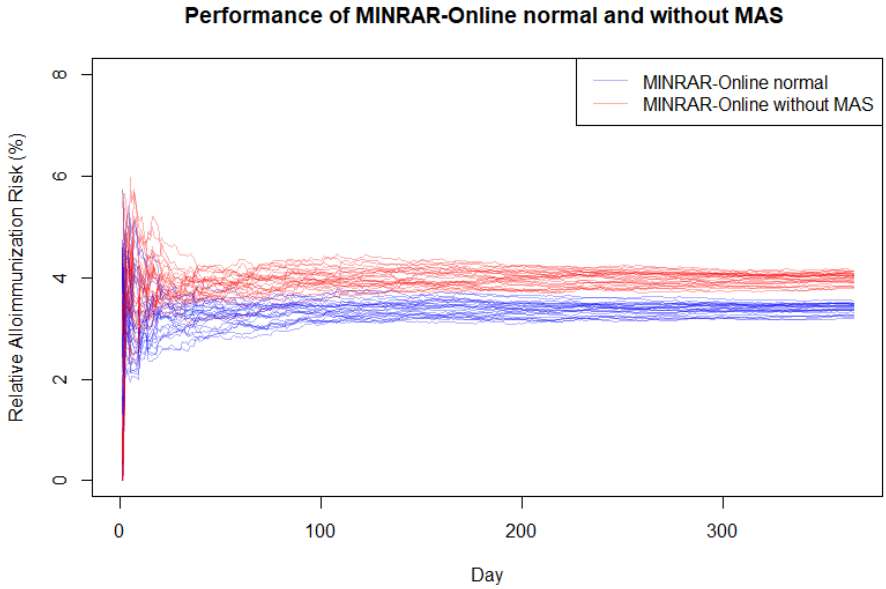

Figure 12: Moving average of the relative alloimmunization risk for 25 simulations using MINRAR-Online and 25 simulations using MINRAR-Online without the MAS term. Both sets of simulations were run on the same combinations of supply and demand scenarios.

patient that is positive for antigen E and negative for antigen K might be mismatched on antigen K to prevent a substitution on antigen E while the relative immunogenicity of antigen K is larger than E. The possibility of reducing the weight of the MAS term in the objective was not investigated but might give slightly better performance if tuned perfectly. However, when set too low (near zero), Figure 12 shows that the average performance is again worsens. We estimated that the gain of tuning the parameter was therefore only small and no experiments were performed.

6.5.6 Influence of the UAD Term on Compatible Blood Saving

To investigate the influence of the UAD term in the objective we ran additional simulations. We ran one particular one-year simulation 100 times, each time with a different weight for the UAD term in the objective function. The results are shown in Figure 13.

The main goal of the UAD term is to limit the issuing of O^- blood to patients with a different major blood type to prevent future shortages of O^- units. In the left part of the figure we can see that when the weight of the UAD term is very low, some shortages occur in the simulation. When this weight is increased the shortages disappear, but the average relative alloimmunization risk also increases. In the right part of the figure we can see that the average issuing age of all major blood groups decreases when the weight of the UAD term is increased, with the exception of O^- units. Their average issuing age can be seen to keep increasing with the UAD term.

The figure indicates that a high weight for the UAD term is preferrable when the more compatible blood units are to be saved. This is especially the case in smaller hospitals, which may not have a large stock and therefore are less well equipped to handle disruptions is supply or demand. On the other hand, larger inventories could benefit from a lower UAD weight as this will improve the minor antigen matching

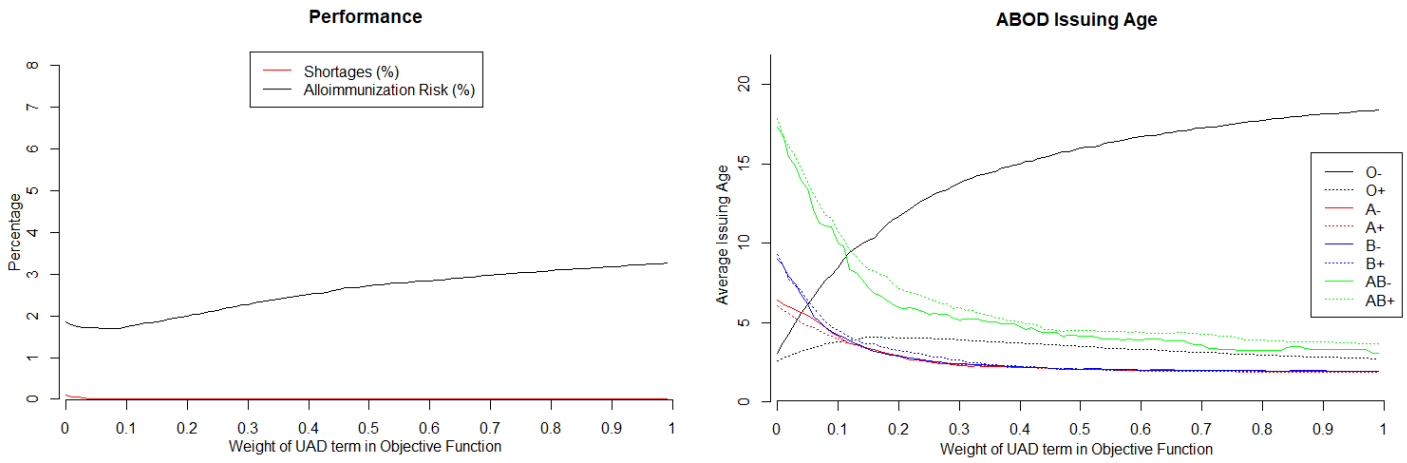

Figure 13: Trade-off between stockpiling O^- to prevent shortages and average alloimmunization risk when changing the weight of the UAD term in the objective. Left: Average performance over the entire simulation in terms of shortages and alloimmunization risk. Right: Average RBC unit issuing age per major blood group. Results are obtained by repeated simulation of a single one-year scenario with an average daily demand of 50 RBC units.

quality. However, care should be taken to make sure that it does not lead to an increase in shortages or possibly an increase in outdating.

6.6 Discussion

In general, the results show that the proposed MINRAR-Online ILP is better able to prevent shortages and minimize alloimmunization risk than the FIFO/MROL strategies. Furthermore, we have eliminated the antigen set parameter that must be specified for the FIFO/MROL issuing strategies. Removing this parameter allows us to consider all AB-RhD compatible matches and thus enabling us the flexibility to issue whatever is needed to prevent shortages when the daily demand and inventory size is low. When demand is high, we can accumulate O^- blood to make future shortages as unlikely as possible. This is done by penalizing the absolute usability difference instead of the relative opportunity loss which the FIFO/MROL models do. As illustrated, the relative opportunity loss can assign a high penalty to substitutions which in practice are not important because the absolute usability of the issued units is low. Furthermore, we do not penalize minor antigen substitution based on prevalence but based on immunogenicity instead. The results show that the MINRAR-Online approach essentially outperforms all variants of the FIFO/MROL issuing strategies. Although this gain is good, it is also as expected. When we compute an assignment for a single day, we see that the solution space of the FIFO/MROL strategies is by definition a subset of the solution space of the MINRAR-Online ILP. This larger solution space allows us to consider all solutions allowed by the FIFO/MROL $\mathcal{A}_3$ issuing strategy with the precision of the FIFO/MROL $\mathcal{A}_{14}$ strategy. We do not discard matches where some minor antigens are incompatible. Instead, we allow mismatches at a cost. To determine the cost per mismatch we have chosen to use the relative immunogenicity. The reason for this is that immunogenicity is the likelihood that a mismatch leads to alloimmunization. This makes it a logical value to use because in the end, the goal of minor antigen matching is to prevent alloimmunization for as many patients as possible. A second reason for using the relative immunogenicity is that the availability of empirical values [3], which are likely correspond to the true

relative immunogenicity of the minor antigens. If instead historical data on antibody formation was used, then the immunogenicity of some antigens may be overestimated while others are underestimated. This is because antigens with high prevalence are only rarely mismatched and therefore the occurrence of antibody formation may be low. The same argument can be made for low frequent antigens.

By using the relative immunogenicity, we can assign a numerical value to a solution that expresses the total mismatching penalty. After consulting with an expert we have decided to count a mismatch per patient at most once, even if more than one of the assigned units have this mismatch. The reason for this was that there is no data on how much more likely alloimmunization becomes when the number of mismatched units increases. Although it is known that the probability of a mismatch does increase with exposure to more than one unit, how much this increase is compared to the first unit of exposure is not known. An important note here is that there is a difference between receiving two mismatching units in one transfusion episode versus receiving two mismatching units in separate transfusion episodes. It is possible that after a first mismatch no antibody formation occurs but the immune system becomes primed. This means that during a later transfusion, the same mismatch is much more likely to lead to antibody formation. In their study Evers *et al.* [3] do not distinguish whether mismatching units were transfused in the same or in separate transfusion episodes. Although they compute alloimmunization incidences after one and after two units transfused, this data cannot directly be used to estimate the difference between exposure to one or two units within the same transfusion episode. Thus, we argue that using the MINRAR method of counting alloimmunization once per patient will capture the most important dynamic of penalizing mismatches. Furthermore, if mismatches cannot be avoided then the MINRAR approach will try to minimize the number of patients who are mismatched as illustrated in Section 5.3.3.

The relative immunogenicity values are parameters of the MINRAR model. This means that any set of values can be used that are deemed representative of the risk introduced by mismatching on the given antigens. When we compare the MINRAR model to the FIFO/MROL models we use these relative immunogenicity values to assess the quality of the matching on the minor antigens. As the MINRAR approach explicitly uses these costs in the optimization, one could think that the comparison is biased, as the FIFO/MROL models do not use this information. However, we think that that the novelty of the MINRAR approach is exactly that the immunogenicity values are included in the optimization process, thereby removing the problem of choosing a set of antigens to include like in the FIFO/MROL strategies. In their work Van Sambeeck *et al.* [7] also discuss a solution for this problem. They use the relative immunogenicity values together with dynamic programming to compute an optimal sequence in which antigens should not be matched for when no compatible unit can be found. However, this sequence can only be used when units are sequentially issued, meaning that the assignment problem is solved for one patient at the time. This makes it impossible to use such a sequence in the FIFO/MROL strategy because both solve a different problem. The FIFO/MROL strategy minimizes the number of unit shortages by computing a maximum flow over all patients simultaneously. On the other hand, when sequentially issuing units using the optimal antigen exclusion sequence, the number of shortages will likely increase as there is no longer a global allocation of all patients to units. Therefore, when computing a global allocation of units to patients an optimal antigen exclusion sequence cannot be used for each patient individually. In theory it is possible to combine the optimal antigen exclusion sequence with the FIFO/MROL allocation strategy by iteratively excluding antigens when the min cost max flow problem still has shortages. However, this is still impractical, as there is no clear way to determine the optimal antigen exclusion

sequence for such a scenario. On top of that, in every iteration, the FIFO/MROL strategy is still limited to a subset of antigens and thus still has the same issue with the limited solution space. All in all, combining an optimal antigen exclusion sequence with a FIFO/MROL approach can make it easier to find a suitable subset of antigens for the FIFO/MROL approach. However, it does not remove the theoretical limitations that prohibit the FIFO/MROL strategy from using the complete solution space. Therefore, any FIFO/MROL strategy inferior to the MINRAR formulation.

The improved quality of matching using the MINRAR model does come at a cost. We have introduced non-linearity in our new definition of shortages, as well as in the definition of the mismatch penalty. Because of this, the MINRAR-Online ILP no longer has the Total Unimodularity property which the FIFO/MROL strategies do have. Therefore, solving the LP-relaxation no longer suffices to compute the optimal integer solution. In practice the MINRAR model is still solved quickly when the total demand is low ($\leqslant$ 200 units). When we consider more patients (and thus also an inventory with more units) then the time to solve the model optimally increases. However, our model is not intended to be solved sequentially and thus some time for the solving is justified. Furthermore, the average daily demand for a regular distribution centre in the Netherlands is about 200 units per day which is still easily solved.

Furthermore, the MINRAR-Online model still has some hyper parameters that require tuning to perform as well as possible in the long run. For instance, the sum of all terms in the objective function is minimized and thereby we assume that the terms are balanced nicely. This does not have to be the case at all. We have investigated how changing the weight of the UAD term influences the average performance and issuing age of the RBC units. The tuning of this weight is not trivial in practice. Whether a low or high weight for the UAD term should be used depends on multiple things: the inventory size, the daily demand, they method of supply (AB-RhD order up to or not), the supply interval, the desired amount of O^- blood in inventory for emergency use, the discrepancy between the AB-RhD prevalences of the supplied RBC units and the patient population etc. Most importantly, when a hospital works with AB-RhD order-up-to levels, then the probability of shortages is low as every day the inventory levels are stabilized. Therefore, it could be argued that these hospitals could not even include the UAD term in the objective at all. In practice however, although hospitals do have the safety of order-up-to levels, they still must make sure to limit the number of unnecessary O^- substitutions. The national blood bank knows the average yearly demand for most hospitals and when the yearly demand for O^- units heavily increases this is investigated as there is only a limited national supply of O^- units and therefore they are still considered a scarce resource.

Like the weight of the UAD term in the objective, the MAS and FIFO terms can also be given alternative weights in order to save more rare antigen negative units or issue units at younger age respectively. Tuning the weights of all terms to fit the exact practical circumstances best is non-trivial. This is because during operation the priorities may change, depending on seasonality, holidays or other external factors. In our simulations we have found the terms to be adequately balanced, as we see no shortages, outdating and relatively little alloimmunization risk. However, there is always room for improvement. When it becomes clearer how large-scale extensive matching will fit into the RBC supply infrastructure, larger and more accurate simulations can give more insight into how the MINRAR-Online ILP can be best tuned for use in such a scenario.

6.7 Usage in Practice

The MINRAR ILP as described in Chapter 5 was designed to be solved *offline*, meaning that the entire problem was known at the time of solving. In this chapter we constructed

the MINRAR-Online model, which is an extension of the MINRAR ILP suitable for use in an *online* setting. It can be used in any RBC unit inventory that must satisfy requests on a regular basis. This includes both hospitals and distribution centres. Currently, hospital and distribution centre inventories already have inventory management systems that keep track of all units in inventory.

6.7.1 Hospitals

Hospital inventory management systems already have the option to search for a unit that satisfies certain criteria including minor antigen compatibility. One or more units can be selected and assigned to a patient and then new criteria can be entered to search units for a next patient. Such a sequential approach is not suitable for the MINRAR-Online model. It is possible to solve the MINRAR-Online ILP model for each patient sequentially and select the best unit(s) every time and remove them from the available inventory, although it is not recommended. When matching is computed for a batch of patients its matching quality will increase for the simple reason that more information is presented to the model before the products are assigned. This is easy to see because the resulting assignment of a sequential solving approach is by definition also a valid solution when the model is run once over all patients considered, assuming that no new supply has arrived during the sequential solve. When patients are batched, many more solutions are available and thus we expect the overall matching quality to improve. Increasing the batch size will continue to improve the overall matching quality, up until the point where the available inventory is no longer able to satisfy the demand of all patients in the batch.

In practice we recommend solving the model after a new set of units is supplied to the inventory. All patients whose demand is still not satisfied can be included in the optimization. When the optimal solution has been computed these patients can be assigned their units and the units can be removed from the available inventory.

6.7.2 Distribution Centre

The optimization can be performed in hospitals as well as in distribution centres. Current practice for hospitals is to order units from their assigned distribution centre on AB-RhD basis only, with the exception that extensively typed units can be ordered for patients who require additional matching. As we have discussed earlier, the matching quality will increase when more patients are considered simultaneously. In a future scenario where extensive matching is possible for the majority of patients, hospitals may forward all their requests to the regional distribution centre. The distribution centre can then compute an optimal assignment of these units to patients before shipping the assigned products to various hospitals. As most hospitals receive daily supply and most patient requests are known more than a day in advance such an implementation is certainly foreseeable.

Several logistical problems will have to be assessed however to make such a system work in practice. Unused units in hospitals will either have to be sent back to the distribution centre or alternatively kept in hospital inventory for issuing. Furthermore, hospitals will still want a small inventory of available units to satisfy unforeseen demand or for emergency use. An important question is then when a patient's request should be forwarded to the distribution centre or satisfied from the hospital inventory.

Despite these logistical challenges we believe the presented model can form the basis of an extensive matching scenario as its flexibility allows it to be used in smaller hospital inventories as well as larger distribution level inventories.

6.8 Extensions

Because of the flexibility of ILP modelling, the MINRAR-Online model can easily be extended with various additional features. Some of these are listed below:

6.8.1 Allowing Partially Satisfied Requests

Currently, the model has a constraint that ensures that partially satisfied requests are not valid. This is because it is not known beforehand how many units end up actually transfused and therefore we assume that all are needed. As we expect that the model will be used in situations where shortages are rare, we have not investigated how partially satisfied requests could be included as valid solutions. However, when shortages become more frequent, either due to demand spikes or more specific antigen negative demand for certain types of patients, it may be worth extending the model such that partially satisfied requests are rewarded. One way of doing this could be to introduce a new decision variable z_i per request that represents the number of units short for this request. We can still leave the binary shortage variable s_i as is. However, we now formulate the demand constraints as follows:

$$\sum_j x_{ij} + z_i = u_i \ \ \forall i \tag{28}$$

$$u_i s_i \geqslant z_i \ \ \forall i \tag{29}$$

These sets of constraints ensure that patients who are not satisfied with the demanded number of units still lead to full shortages. This is because whenever the value of z_i is larger than zero, s_i is forced to be equal to one. Using these new constraints, we can reformulate the shortage term in the objective function as follows:

$$\min \sum_i s_i S - z_i \tag{30}$$

When we choose $S = 3n+1$ this objective will always minimize the number of shortages, as well as maximize the number of units issued. As z_i can be at most three ($u_i \leqslant 4$), we must choose $S = 3n + 1$ to make sure that preventing a single shortage is always preferred over partially satisfying multiple requests. The mismatch variables y_{ik} can keep their current definition and constraints. This will ensure that partially satisfied requests can still lead to mismatches, which is also the desired behaviour.

6.8.2 Considering Demand for Multiple Days

In practice most demand for RBC units is already known several days in advance. Therefore, it would be useful to include information about these future requests in the MINRAR-Online model. However, care should be taken as it is possible that when all future requests that are known are included in the optimization it may well be that the total number of units demanded is larger than the current inventory size. This could influence performance as well as create apparent shortages that do not yet actually exist. One way to mitigate this is to use a weighted objective function where requests that are further in the future have lower shortage and alloimmunization costs. Therefore, when there are too many requests considered, the optimal solution will still try to satisfy as much of the demand on the current day, as that is the most important objective.

Another possibility is to aggregate all future requests by major blood group and compare these aggregates to the expected supply (based on donor major blood types). It can then be estimated per major blood group whether shortages are likely to occur or not. When shortages are likely the size of the usability penalty can be increased such

that the issuing strategy will focus less on preventing minor antigen mismatch and more on preserving the most compatible blood types. Similarly, this strategy can be applied to assess the influence of minor antigen substitution. If, for example, the future requests have relatively fewer K negative phenotypes than expected, the substitution penalty of K can be reduced as fewer K negative units are needed. The same argument can be applied in reverse when future demand shows a relatively high percentage of K negative requests. This method of tuning the substitution and usability parameters based on aggregated future demand does not influence the size of the problem at hand, as only parameters are changed. Therefore it leaves performance unaffected.

6.8.3 Categorizing Different Patient Types

In the proposed the MINRAR-Online model, all patients are considered similar. Mismatching one patient on K causes the same penalty as mismatching any other patient on K. However, in practice there can be large differences in the severeness of mismatches for different types of patients. When it is known beforehand that a patient may need other blood transfusions later in life, there is a larger incentive to prevent mismatching as antibody formation should be prevented in order not to complicate these later transfusions. Furthermore, patients with certain (chronic) illnesses like SCD, Autoimmune Hemolytic Anemia (AIHA) or Thalassemia should also be prioritized in the matching process. In the MINRAR-Online model we can use a different mismatch penalty for every patient and therefore we can easily extend the model to account for different patient types. We will further explore this approach in Chapter 8.

6.8.4 Accounting for Unused Units

Many requests for blood transfusion are for more units than actually transfused. This is because it is often not known exactly how many units will be used and doctors rather have too many units available than too few. The MINRAR model currently does not account for this. One way to extend the model to allow for this is to make the units assigned to a request positional, meaning that there is an order in the units assigned. For example, a request for three units can have a first unit that is most important, a second unit that is less important and a third unit that can have more mismatches as it is likely not used. Note that we cannot divide the mismatch penalty by the number of units assigned as this does not reflect what happens in practice. When only 2/3 units are used it does not mean that we use 2/3 of every unit, but instead we choose two out of the three assigned units and use those.

To adapt the MINRAR ILP formulation of Section 5.3.4 to assign units in a particular order for a patient, the simplest solution is to add a dimension $l \in \{1, ..., u_i\}$ to the x_{ij} decision variables that indicates the position of the unit:

$$x_{ijl} = \begin{cases} 1 & \text{if unit } j \text{ is assigned to request } i \text{ in position } l \\ 0 & \text{otherwise} \end{cases} \tag{31}$$

$$y_{ikl} = \begin{cases} 1 & \text{if the unit assigned to request } i \text{ at position } l \text{ mismatches on antigen } k \text{ and} \\ & \text{no assigned units at a lower positions mismatch on this antigen} \\ 0 & \text{otherwise} \end{cases} \tag{32}$$

To be able to let the order of how units are assigned influence the quality of the solution, we need a parameter that specifies how important a unit is for a given position when assigned to a patient. For example, when a patient requires three units, the

mismatches induced by the first assigned units should be counted fully. Then additional mismatches (antigens that were not mismatched by the first unit) can be counted with a factor of 0.8. The remaining mismatches induced by the third unit can then be counted with a factor of 0.4. To specify these factors per patient we need another parameter that specifies the importance of a unit at position l for a request i:

$$v_{il} = \text{factor by which the mismatches induced by unit at position } l \text{ should be multiplied for request } i \tag{33}$$

The resulting ILP will then be:

$$\min \sum_i s_i \cdot (n+1) + \sum_i \sum_k \sum_l y_{ikl} v_{il} \cdot a(k) \qquad \text{(Objective)}$$

s.t.

$$\sum_j \sum_l x_{ijl} + s_i \cdot u_i = u_i \quad \forall i \tag{34}$$

$$\sum_i \sum_l x_{ijl} \leqslant 1 \quad \forall j \tag{35}$$

$$y_{ikl} = \sum_j x_{ijl} \cdot d_j(k) - \sum_{l'=1}^{l-1} y_{ikl'} \quad \forall i,k,l \;\; if \;\; b_i(k) = 0 \tag{36}$$

$$\sum_l x_{ijl} \leqslant c_{ij} \quad \forall i,j \tag{37}$$

$$x_{ijl}, y_{ikl}, s_i \in \{0,1\} \quad \forall i,j,k,l \tag{38}$$

The most important constraint here is Constraint 36. It ensures the correctness of the y_{ikl} variables. This is done by observing that $y_{ikl} = 1$ if for request i, the unit assigned at position l mismatches on antigen k and none of the units assigned at lower positions induce this mismatch. This is exactly what the constraint expresses. $\sum_j x_{ijl} \cdot d_j(k)$ is equal to one when an assigned unit at position l mismatches on antigen k. Then we must subtract 1 if any of the earlier units already mismatches on this antigen, which is equal to $\sum_{l'=1}^{l-1} y_{ikl'}$.

Because the v_{il} parameters can be specified per request it is possible that for some requests where we know that the full demand is used, we can have $v_{il} = 1 \; \forall l$. In that case the model is equivalent to the original MINRAR formulation in Section 5.3.4. However, for other requests (especially those with a large number of units requested) we can have decreasing values for v_{il}.

6.8.5 Mismatch Penalty based on Number of Units Mismatched

Thus far we have worked with the assumption that for a patient is does not matter how many of the assigned units mismatch, but only whether they mismatch or not. This approach was chosen because of simplicity and absence of actual data specifying how much worse more exposure during a single transfusion episode is. It may be that in the future evidence regarding the impact of mismatches will become available or can be estimated (especially once extended matching is applied in practice). The model could then be extended to use a mismatch penalty that is a function of the antigen mismatched and the number of units that mismatch, to better minimize the actual alloimmunization risk for the patients involved.

7 Offline Assignment

The MINRAR-Online ILP proposed in the previous chapter is suitable for *online* use, meaning that it should be iteratively solved to compute assignments when these must be made. To test the performance of this model we have compared it to three variants of an earlier model. We ran simulations to assess the performance of these issuing strategies. This gave an indication of how much better the proposed model performed compared to other model, but it does not give an indication of the absolute performance of the MINRAR-Online model compared to the optimal issuing strategy. To gain insight in the absolute performance of the model, we will compare it to an *offline* model. An offline model is not run iteratively while the problem presents itself, but instead it is run when the whole problem is known. In our case this means computing an optimal assignment of RBC units to patients over a given period in the past. In this chapter we will create such a model and use it to compare the performance of the (MINRAR-)online model to the maximum theoretical performance achievable.

7.1 Increasing the Solvability

The main issue with an offline model is the size of the solution space. If we want to compute an optimal issuing in hindsight over one year in a scenario with on average 100 patients per day, then the number of possible patients that a single RBC unit can be assigned to is roughly $100 \cdot 365 = 36,500$. Furthermore, in the total simulation there will be about $100 \cdot 365 \cdot 2 = 73,000$ RBC units, as there are 100 patients per day who on average (roughly) require two units. Scale this up to all possibilities for units and patients and the total number of possibilities in a naive implementation is roughly $36,500 \cdot 73,000 = 2,664,500,000$ possibilities as every patient can be matched any unit over the entire simulation. To make sure that the model is still solvable, we must make some adjustments to limit the size of the solution space.

7.1.1 Maximum Shelf Life

When we use a naive approach and create a decision variable for every combination of unit and patient, the number of decision variables will become very large and therefore consume a lot of memory. To reduce this effect, we can observe that every request can only be assigned units which have arrived at most 35 days before the request must be satisfied. If a unit is supplied earlier, it cannot be assigned to the request as it has already expired. This means that for each request we only have to create $L = MSL \cdot \bar{d}$ variables where MSL is the maximum shelf life and $\bar{d}$ is the average daily demand. We use the average daily demand because at the end of each day, the inventory is resupplied with a number of units equal to the demand size of the current day. Therefore, the average supply size is equal to the average demand size except for the first day, when the entire inventory is filled.

Limiting the number of possible units per request to only those which are within inventory and not expired will considerably reduce the number of variables per request. This effect increases when the MSL parameter is further reduced. We have chosen to compute the performance for the offline model for MSL values between 1 and 14 days. The reason for this is twofold: firstly, computing the optimal assignment for different values of MSL will give us insight into how the minor antigen matching quality improves when units are allowed longer in inventory. Secondly, it will allow us to reduce the solution space further, which is necessary to make the optimization practically feasible.

7.1.2 Remove Non-linearity in Alloimmunization Penalty

In the proposed MINRAR model in Section 5.3.4 we use the binary y_{ik} variables to indicate whether one or more of the assigned units for patient i mismatch on antigen k. As mentioned before, the introduction of these variables was necessary as the mismatching behaviour could not be modelled as a linear cost per RBC unit. Because the offline model will only be used to get an indication of the optimal performance of an issuing strategy, we will not use the y_{ik} variables, but instead use an upper and lower bound on the mismatching costs. These bounds are linear costs per RBC units and can therefore be precomputed, making optimization much more efficient. Furthermore, the y_{ik} variables and constraints are removed from the ILP, reducing the overall size and complexity.

7.1.2.1 Lower Bound If we want to compute a lower bound on the alloimmunization risk, we must make sure to never count too much alloimmunization per patient-unit combination. We do this by multiplying the alloimmunization risk of a match between patient i and unit j by $\frac{1}{u_i}$ if the request is for u_i units in total:

$$a_{ij}^{lower} = \frac{1}{u_i} \sum_k a(k) \cdot d_j(k) \cdot (1 - b_i(k)) \tag{39}$$

Now if for patient i all the u_i units assigned have the same mismatch then the alloimmunization will sum to 100% of the original value. If, on the other hand, only some of the assigned units mismatch where others do not, then only a fraction of the alloimmunization cost will be included. Another way to think of these costs is that it is the result of relaxing the integrality constraint for the y_{ik} variables. This immediately proves that the alloimmunization computed using these fractional values is a lower bound as the relaxation of the variables will by definition yield the lowest possible value for the objective function and the integral solution is still an element in the solution space.

7.1.2.2 Upper Bound To compute an upper bound we can use a similar approach, but now we must never underestimate the alloimmunization value. We can do this by counting the full alloimmunization penalty per mismatch per patient. Thus, we ignore the fact that if two products mismatch on the same antigen for a single request the alloimmunization should be counted only once.

$$a_{ij}^{upper} = \sum_k a(k) \cdot d_j(k) \cdot (1 - b_i(k)) \tag{40}$$

The computation of this upper bound is equivalent to splitting all requests into requests for single products only and then computing the optimum assignment for these singleton requests (given that there are no shortages). As is also already discussed in Section 5.3.2, this variant is easily solvable as there is no non-linear behaviour per patient. Furthermore, it is easy to see that the total alloimmunization penalty of this model is an upper bound to the total alloimmunization penalty of the optimal model, as the optimal solution will always have a total alloimmunization penalty less than or equal to the optimal upper bound.

7.1.3 Limited Time Span

To limit the solution space further we have chosen to compute the optimal performance for the issuing strategy not for a full year but only for 225 days. We will show why this value is chosen later.

7.2 Offline Considerations

7.2.1 Boundaries

To compare the offline model to the online model we will use the same supply and demand scenarios as described in the previous chapter. However, to make a fair comparison we cannot simply run the offline model over the entire scenario. The reason for this is that the offline model has fixed start and endpoints. These will influence the performance of the model, as at the start the inventory is filled with random units and we can expect the matching quality to be different since the inventory has not yet reached a steady state. Similarly, at the endpoint of the timespan the offline model can use up all the antigen negative units which the online model might save for later use. Ideally, we only want to measure the performance of the online model between these initialization and exit periods, as they are most representative of the steady state performance.

7.2.2 Variation in Demand

In our model we assume that the number of units supplied per day is equal to the number of units demanded on the previous day. Earlier we mentioned that the number of variables per request was equal to L, which was the product of the maximum shelf life and the average daily demand. However, because the daily demand may vary from day to day, it could be the case that by chance, for certain days the total number of units which have been supplied in the past MSL days is more than L. For simplicity we will limit the number of units considered for the requests on these days at L. By doing this we possibly exclude the option of matching some units which had not yet expired. However, we expect that this will not have a significant influence on the average performance as only those units will be affected which were unlikely to be assigned to such requests in the first place.

7.2.3 Outdating

Because we have set the size of the supply for each day equal to the size of the demand on the previous day, the outdates and shortages become intertwined. This is because when a unit outdates, the supply for the next day does not increase. Thus, when units outdate it will cause the inventory size to decrease. There is no easy way to assess this problem without increasing the solutions space, as it would no longer be known beforehand on which day a unit is supplied. We have chosen not to alter the model to account for this. Instead, we will penalize outdating just as we penalize shortages. When we solve the model with small values of MSL we expect that there may be substantial outdating and therefore this may lead to a decreased effective inventory size which in turn might cause shortages. However, some partially inaccurate results in testing for these small values of MSL is not a real issue as these values were already unsuitable for comparison to the online model. Still, we have included them in the figures such that the consequences of maintaining such a small inventory size become visible. For larger values of MSL we see that outdating tends to zero and thus the problem is no longer relevant.

Note that we cannot ignore outdating in our model as otherwise an optimal solution could be found where the decrease of inventory size caused by units outdating is compensated by letting those units outdate which are hardest to match. Because the online model has an outdating percentage of 0%, it means that even those antigen positive units must be issued. Therefore, we have to include outdating in the objective to prevent the non-issuing of hard to match RBC units.

7.3 Offline ILP Formulation

We will now formulate the offline ILP used to compute an upper and lower bound on the optimal performance of any issuing strategy for a given supply and demand scenario and timespan. First, we will introduce some notation.

Notation	Meaning
T	Number of time periods (days) considered
T_{init}	Number of time periods (days) of initialization period
T_{exit}	Number of time periods (days) of exit period
$T_{effective}$	Number of effective time periods (days) $T_{effective} = T - T_{init} - T_{exit}$
MSL	Maximum shelf life considered (days)
$\mathcal{S}$	Supply Scenario
$\mathcal{D}$	Demand Scenario
$\bar{d}$	Average daily number of units demanded in $\mathcal{D}$
b_i	Blood phenotype of patient $i \in \mathcal{D}$
d_j	Blood phenotype of RBC unit $j \in \mathcal{S}$
r_i	Day that patient's request i must be satisfied
u_i	Number of RBC units requested for patient i
q_j	Day at which unit j is supplied
U	Shortage penalty per patient
O	Outdating cost per RBC unit
L	Maximum number of units considered per request ($L = MSL \cdot \bar{d}$)

Table 21: Mathematical notation for blood, RBC units and patients.

To make sure that for every request we only consider L units we will only introduce a decision variable x_{ij} if unit j is one of the L units can be assigned to patient i. This means that the earliest day that unit j must have entered inventory is day $r_i - MSL + 1$ and the latest day is day r_i. Furthermore, no more than $L - 1$ units should satisfy the same condition and have a higher index than j. Lastly, unit j must be AB-RhD compatible with request i in order to create an x_{ij} variable.

Parameters

$$a_{ij} = \text{Mismatch cost. Either } a_{ij}^{lower} \text{ or } a_{ij}^{upper}$$

Decision Variables

$$x_{ij} = \begin{cases} 1 & \text{if unit } j \text{ is assigned to request } i. \\ 0 & \text{otherwise.} \end{cases}$$

$$s_i = \begin{cases} 1 & \text{if a shortage is incurred for request } \mathcal{D}_i. \\ 0 & \text{otherwise.} \end{cases}$$

$$o_j = \begin{cases} 1 & \text{if unit } j \text{ outdates.} \\ 0 & \text{otherwise.} \end{cases}$$

ILP

$$\min \sum_i s_i U + \sum_{j | q_j \leqslant T - MSL} o_j O + \sum_i \sum_j x_{ij} a_{ij} \qquad \text{(Objective)}$$

s.t.

$$\sum_{l} x_{ij} + s_i \cdot u_i = u_i \quad \forall i \tag{41}$$

$$\sum_{i} x_{ij} + o_j = 1 \quad \forall j \tag{42}$$

$$x_{ij}, s_i \in \{0, 1\} \quad \forall i, j \tag{43}$$

The objective function minimizes the number of shortages, outdates and the total alloimmunization cost. Note that we only sum the outdate variables when the supply day q_j is smaller than or equal to $T - MSL$. This is because when a unit is supplied later, it cannot outdate within the duration of the scenario. Constraint 41 ensures that demand is satisfied fully or a shortage is incurred. Constraint 42 allows a unit to be issued at most once and if it is not issued it must outdate. Constraint 43 forces all decision variables to be binary.

7.4 Computational Experiments

The goal of the offline model is to establish a reference point to estimate the maximum achievable performance of the online approach. In order to do so, we must ensure that the performance computed by the offline model is representative of the steady state performance of the online model. Therefore, we have to ignore the performance in the beginning and at the end of the offline model. The reason for this is clear: in the beginning the inventory is filled randomly and therefore the prevalence of each antigen in inventory is likely not as it would be in a steady state. Thus, the performance might be worsened. Similarly, at the end of the scenario the online model will not know that the simulation is ending and therefore it will still try to save useful antigen negative units. The offline model on the other hand, has complete information and can therefore adjust its 'issuing strategy' at the end of the scenario to use up all useful units, leaving behind more antigen positive units, which are harder to issue without inducing mismatches.

To assess these problems, we ignore the first and last four weeks of the scenario ($T_{init} = T_{exit} = 28$). As we only test with a maximum shelf life of at most 14 days, these initialization and exit periods should be of sufficient length. Although we ignore the performance in these periods, we should still include them in the objective function as usual. Thus, we optimize over the entire duration (T), but for the results we only consider requests that occurred during the effective duration ($T_{effective}$).

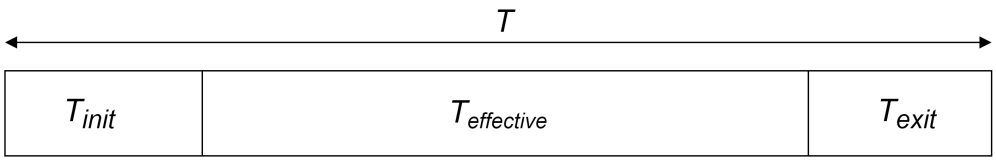

Figure 14: Schematic view of duration of the offline problem.

7.4.1 Setup

In order to make an equal comparison between the performance of the offline model and online model we first investigated the performance of the offline model as a function of the duration of the scenario. For ten different demand scenarios with $\bar{d} = 25$ we computed the average performance over different durations ranging from 70 days up

until 280 days. In all cases the same supply scenario was used. The lower and upper bounds of these ten demand scenarios are plotted in Figure 15.

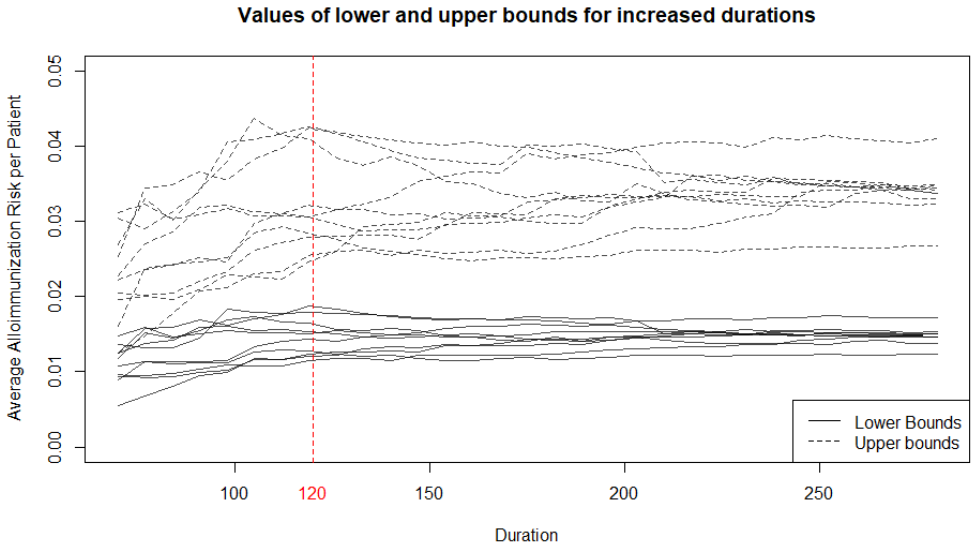

Figure 15: Performance of the offline model for ten different demand scenarios in terms of their upper and lower bound on relative alloimmunization risk for different values of total duration and a fixed MSL of 14 days.

We can see that when the duration is shorter than approximately 120 days, the average alloimmunization risk per patient is still increasing. The reason for this is that the total duration of the offline problem is likely too short to give an accurate estimate of the long-term behaviour. There are four weeks of initialization period and four weeks of exit period. This means that 56 days of the total duration are not included when calculating the results. We can see that after 120 days the lower bound values seem to have reached their final value. The values of the upper bound are less well behaved and seem to take longer to settle on a final value. However, as we are most interested in the lower bound this is acceptable. In our computations we will use a duration of 225 days for the scenarios with average daily demand 25 and 50 and a duration of 120 days when the average daily demand is 100 units. The reason we used a smaller value for the latter is to limit the memory use even further. Furthermore, all these durations include a four week initialization and a four week exit period.

7.5 Results

To compare the performance of the online model to the offline model we have run both models on the same combination of supply and demand scenario. Therefore, both models are presented with the same patients and supplied with the same RBC units. Figure 16 shows the average performance measured for an average daily demand of 25, 50 and 100 units per day. Each row of the figure consists of two parts. On the left we show the average performance of the offline model for different values of MSL. The grey area is bounded by the lower and upper bound alloimmunization risk per patient, thus implying that the true optimum must lie somewhere within these bounds. This figure also includes the percentage of shortages and outdates for the offline model. In each row on the right we show the performance of the online model on the same scenario.

We have computed the performance of the MINRAR-Online model for different weights of the UAD in the objective function. The reason for this is that this hyper parameter influences the minor antigen matching quality. When we use a weight of 1, then the issuing strategy will heavily penalize major blood group substitution and therefore the minor antigen matching quality is lowered. To compare the best version of the online model we have included these figures to show where the optimal choice of this weight lies. This is at the point where the shortages are zero and the alloimmunization per patient is as low as possible, which is indicated with a red line. The y-axes of both plots have the same scale for the Relative Alloimmunization risk, which allows evaluating the performance of the online model against the offline model by comparing the height of the red line to the bounds of the grey area.

It is important to know how to interpret the x-axis in the plots on the left. The value of MSL does not only limit the maximum shelf life, but it also limits the inventory size. We start with an inventory size five times the average daily demand, but if $MSL \leqslant 4$ then units can stay in inventory at most four days so the effective inventory size can never be five times the daily demand. This is because having an inventory size of five times the average daily demand implies that the average issuing age is five days, which is impossible with MSL values lower than five. Therefore, only the results for values of $MSL \geqslant 5$ can be meaningfully compared to the online approach.

The results should be interpreted as follows. First, we look at the figures on the right to find the minimum value for average alloimmunization risk per patient where the number of shortages is zero. This point is shown with a red line. Next, we can compare this value to the figure on the left to see how close to the optimal average alloimmunization risk the online model is able to match patients. Most interesting is the comparison between the optimal performance of the online model and the computed performance of the offline model with $MSL = 14$. This is because this value of MSL is most representative for the MSL value in practice, which is 35. Although 35 is much higher than 14, we can still see that the curve of the bounds computed by the offline model flattens. We do not expect the alloimmunization risk to lower much further because the far majority of RBC units are issued within 14 days.

When looking at the first row in Figure 16, where a comparison is made on a scenario with an average daily demand of 25 units, we can see that the online model performs best when the weight of the UAD term in the objective function is equal to 0.26. The corresponding average alloimmunization risk per patient is just over 3%. Now we can compare this to the bounds computed by the offline ILP as shown in right graph. Here we can see that the mismatch penalty of the optimal issuing strategy when the MSL is equal to 14 days is between 1.8% and 4.2%. It is not clear whether this curve is going to flatten off further, although the lower bound of the grey area seems to stop decreasing. The online model most likely performs close to the optimal performance and is at worst about 1.5 percentage point off.

When we look at the second set of graphs of Figure 16 (average daily demand of 50 units), we see a similar result. The lowest relative alloimmunization risk per patient is reached when the weight of the UAD term in the online ILP is equal to 0.08. Shortages are then eliminated and we see an average alloimmunization risk of 1.7% per patient. When compared to the optimal offline performance computed we see that again the performance of the offline model is within the bounds of the online model. We estimate that the worst-case difference to the optimal lower bound will be about one percentage point. Thus, we conclude that again the performance of the online model is most likely close to optimal performance.

Lastly, we compare the performance for the scenario with an average daily demand of 100 RBC units. This is shown in the bottom row of Figure 16. The optimal performance

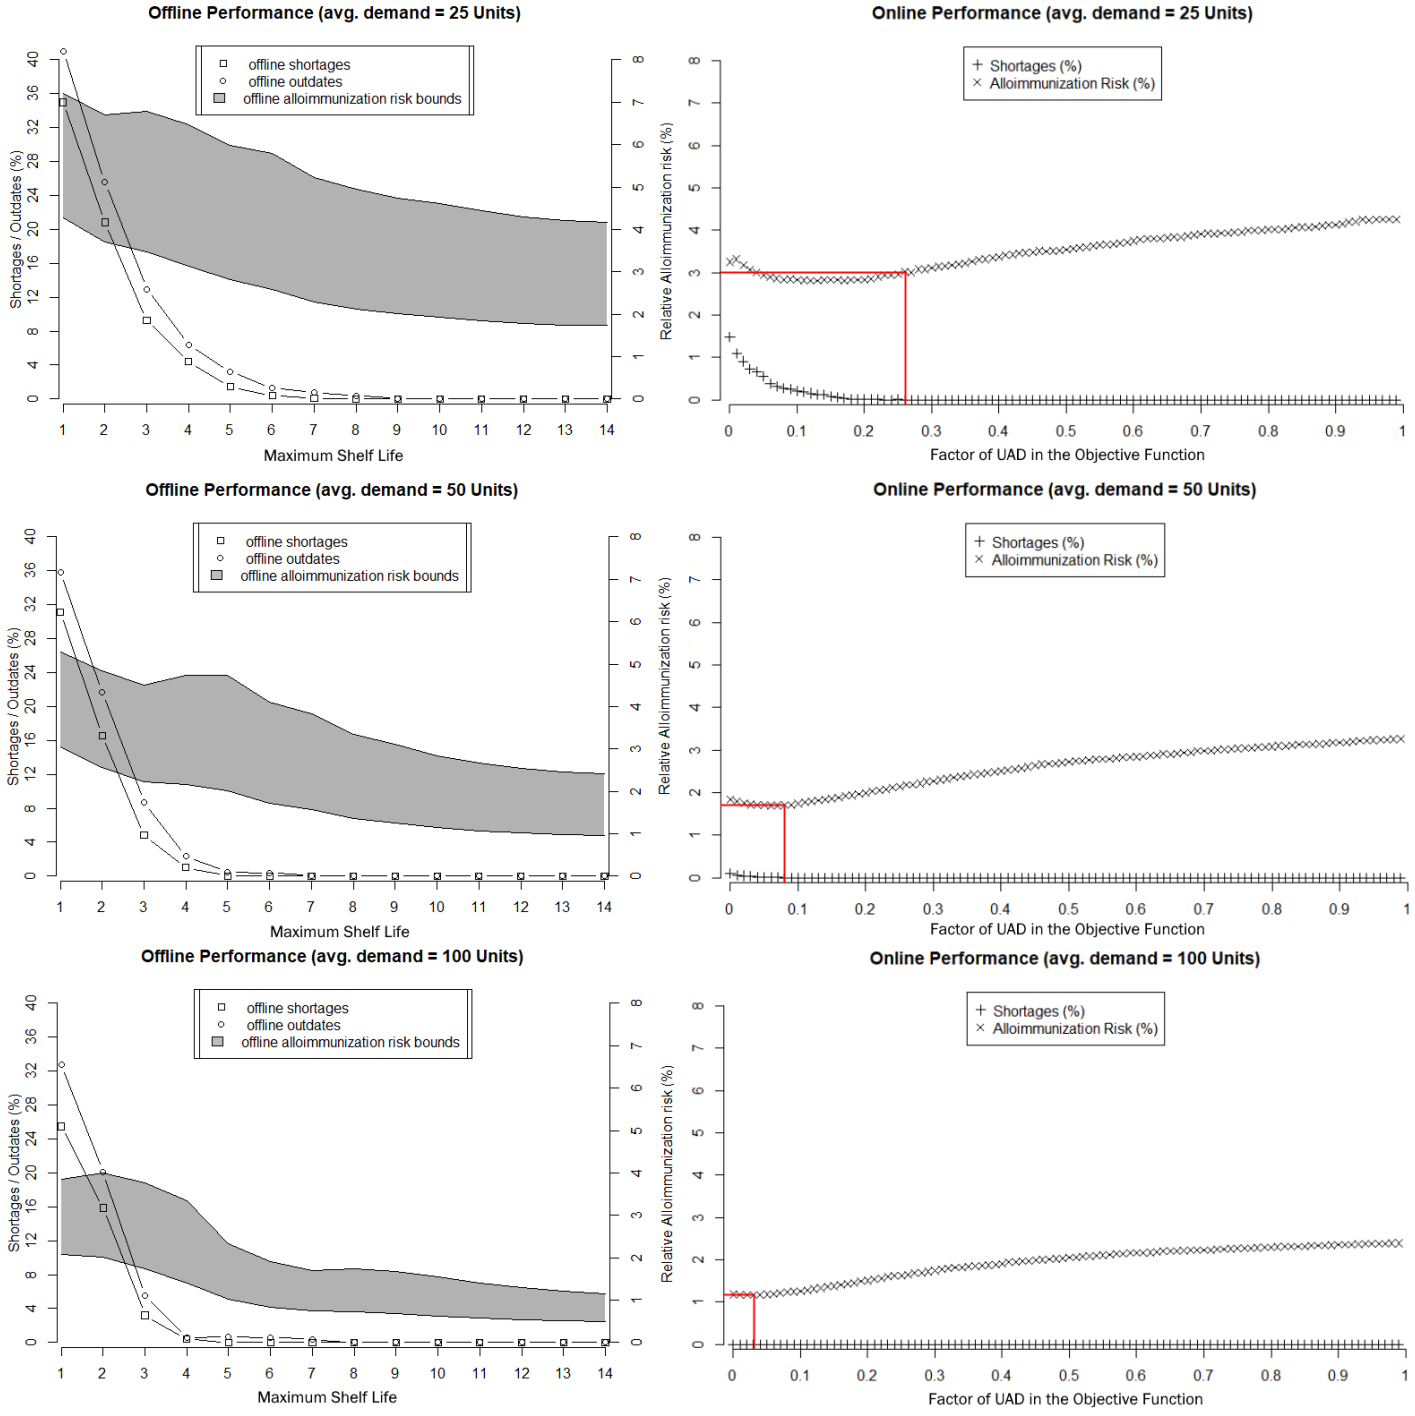

Figure 16: Comparison of the performance of the online issuing strategy to the optimal issuing strategy as computed by the offline ILP on the same supply and demand scenario for three values of average daily demand. Left: The performance of the offline ILP, computed for different values of MSL from 1 to 14 days. Right: The average performance of the online model for different weights of the UAD term in the objective function (0 to 1). The red line shows the lowest value for alloimmunization risk where the number of shortages is zero.

for the offline model lies around 1.2%. Comparing this to the offline model we can see that this is only slightly worse than the upper bound of the optimum for $MSL = 14$ which is 1.1%. It is unclear whether the optimal lower bound will stop decreasing after $MSL = 14$. However, again the worst-case difference seems to be about one percentage point.

7.6 Discussion

The results shown above must be put into perspective. First of all, the comparison that was done used the same supply and demand scenario for both the online and offline models. This allowed us to make a valid comparison of the performance of for both models accurately. However, only one combination of scenarios was used in the comparison. Therefore, it is possible that by chance the specific scenario that was used has some abnormalities or other irregularities. This means that absolute performances shown in the figure are not necessarily representative of the average performance over all possible scenarios. However, as the same scenarios were used in the comparison for both models, the relative performance gap between the models is still a representative indication of the actual difference between the performance of the online model and the optimal issuing strategy.

The main goal of this chapter is to give an indication of the optimal performance of an issuing strategy, such that we can compare the performance of the MINRAR-Online model to this optimum. In order to do so we have constructed an ILP which computes the optimal issuing strategy by solving an assignment problem over an entire scenario. To limit the size of this ILP we have chosen to use a lowered value of MSL to reduce the number of possible assignments between units and patients. The curves in the figures show that although the performance of the offline models seems to improve for higher values of MSL, the rate of this improvement becomes smaller and even seems to come to a stop. Furthermore, as the tuning of the weight of the UAD term in the objective function of the MINRAR-Online model has a large impact on the actual performance, we have included Pareto figures showing the trade-off between shortages and relative alloimmunization risk. We have set the weight as to optimize the online model to perform optimally in terms of shortages and alloimmunization risk. The results show that if this weight is chosen optimally, the online minor antigen matching quality is likely within one percentage point of that of the optimal offline lower bound. One percentage point of relative alloimmunization risk roughly corresponds to a mismatch on an antigen S or $Fy(b)$ as these antigens both have a relative immunogenicity of 1.31%. Furthermore, these two antigens are among the three antigens with the lowest immunogenicity. The conclusion that can be drawn from this is that although slightly better antigen matching may be possible, the relatively simple MINRAR-Online ILP performs close to optimal antigen matching for the majority of all eleven minor antigens with nonzero relative immunogenicity.

Finally, we would like to remark that in practical settings we recommend that the UAD weight in the objective function is rather set too high than too low. In the analysis above we used the smallest weight for the UAD term which leads to zero shortages. This means that some shortages are only just avoided. In practice, it is much safer to use a more conservative variant, such that rare blood groups like O^- are substituted with caution. The higher the weight of the UAD term in the objective of the MINRAR-Online ILP, the more O^- blood will be stocked which increases the flexibility of the inventory. This is most relevant when the MINRAR-Online strategy is used on a distribution centre level where supply is uncertain. Hospitals have the luxury of AB-RhD specific order-up-to levels and therefore to focus less on stockpiling O^- units.

8 Matching Patient Groups

In the Netherlands, extended antigen matching is currently only available for certain patient groups such as patients with Sickle Cell Disease (SCD), Thalassemia, Autoimmune Hemolytic Anemia (AIHA), Myelodysplastic Syndrome (MDS) or patients with irregular antibodies in their blood. Also, women under 45 years of age are matched on antigens c, E and K to prevent possible complications in future pregnancies. When in the future extended matching is a possibility for a larger group of patients, we do not want the availability of extensively matched products for these special patient groups to decrease because of this. While we assume that genotyping all donor blood will lead to the discovery of more highly usable blood groups among donors, this will also lead to the discovery of more patients with these blood groups. Since it is far more relevant for the afore mentioned patient groups to receive extended matching, we want to investigate how to prioritize these various patient groups for matching using our MINRAR-Online issuing strategy as a basis.

8.1 Matching Priorities for Special Patient Groups

To assess which patient groups should have priority in an extended matching approach we have consulted with experts on clinical immunology. They have kindly constructed the Table 22. This table categorizes seven different patient groups. For each group the severeness of mismatching on every antigen is classified into four levels. Level four matching is a must and these patients can therefore only receive RBC units that are compatible on these antigens. Level three matches are classified as important and should only be mismatched if no other unit is available. Level two matches are preferable and level one matches are only if possible.

	C	c	E	e	K	Fy(a)	Fy(b)	Jk(a)	Jk(b)	S	s
With antibodies	4	4	4	4	4	3	3	3	3	2	1
Sickle Cell Disease	4	4	4	4	4	4	3	4	4	3	2
Thalassemia	4	4	4	4	4	2	2	2	2	1	1
MDS	4	4	4	4	4	2	2	2	2	1	1
AIHA	4	4	4	4	4	3	3	3	3	2	2
Women <45	2	4	4	2	4	1	1	1	1	1	1
Remaining Patients	2	2	2	2	2	1	1	1	1	1	1

Table 22: Desirability of antigen matching for various patient groups, distinguishing four levels: (1) if possible, (2) preferred, (3) important (4) must. The level of each antigen per group was estimated by immunology experts and based on immunogenicity, pathogenicity and clinical relevance for the patient group. Antigens that are not included are deemed irrelevant.

The matching levels in the table cannot directly be used as weights in an optimization algorithm. This is because the four levels cannot be interpreted as linear penalties. A level two mismatch may be 10 times more severe than a level one mismatch, while a level four mismatch is possibly 100 times more severe than a level three mismatch.

Therefore, we combined the relative immunogenicity weights with estimated weights describing the relative importance between the patient groups to derive the weights for matching preferences. These estimated weights are shown in Table 23.

	Category Weight
With antibodies	60
Sickle Cell Disease	100
Thalassemia	40
MDS	40
AIHA	60
Women <45	1
Remaining Patients	1

Table 23: Estimated matching weights for individual patient groups. These weights express how much more severe a mismatch on the same antigen is for the different categories, provided that the antigen is not in the "Must" category.

When we multiply these weights with the relative immunogenicity values, we already obtain a good approximation of the relative importance of matching for different antigens for each of the patient groups. However, we still need to correct for the pathogenicity of antigen mismatches. The pathogenicity expresses the severeness of alloimmunization against a particular antigen. We will mainly use this to correct for antigens that have low clinical relevance in an extended matching strategy. For the pathogenicity we did not have empirical values that we could use. Instead, we have estimated pathogenicity values based on the supplied matrix. These values are shown below in Table 24.

	C	c	E	e	K	Fy(a)	Fy(b)	Jk(a)	Jk(b)	M	S	s
Relative Pathogenicity	1.5	1.5	1.5	1.5	1.5	1	1	1	1	0	0.5	0.5

Table 24: Estimated (Relative) Pathogenicity per antigen derived from Table 22.

We estimated the Rhesus antigens and K to receive a higher pathogenicity. This is because these antigens are given a higher desirability level in Table 22. Antigen M is given zero pathogenicity. This is because the clinical relevance of antibody formation against M is negligible and therefore an extensive matching approach should not need to match patients for M. On the contrary, antigen s, which was estimated to have zero relative immunogenicity [3] is said to have little but not zero importance in the matching, especially for patients with SCD. Therefore, we have artificially given antigen s an immunogenicity of 0.005. This value was estimated based on the known immunogenicity of the other antigens and Table 22. We have given antigens S and s a pathogenicity of 0.5 as their overall priority in the matching is low.

As we can see in the original table constructed by the experts consulted (Table 22), there are several antigen-patient combinations categorized as "must". We can either give these a very high weight in the matching, or not allow any mismatches in a valid solution. We choose for the latter approach as these matchings really are essential and therefore incompatible matches should not be allowed.

We can combine all the individual components to construct a final weight matrix that approximates the original matrix supplied (Table 22). First, we have again normalized the relative immunogenicity of the antigens for which we used the same value as described in Section 2.5. However, we have used re-normalized values since we artificially added an immunogenicity of s of 0.005. Furthermore, in this normalization step we have given M an immunogenicity of zero for convenience. Next, we constructed the final matrix of mismatch penalty weights using the following procedure:

- Take a combination of patient category and antigen.
- If the given antigen-patient combination is a must in Table 22, denote an 'M' and stop.
- Otherwise, compute the weight of this combination as a multiplication of category weight (Table 23), pathogenicity (Table 24) and immunogenicity (adapted Table 5).

The result is the following matrix:

	C	c	E	e	K	Fy(a)	Fy(b)	Jk(a)	Jk(b)	S	s
With antibodies	M	M	M	M	M	1.3635	0.4040	2.5756	0.1010	0.2020	0.0769
Sickle Cell Disease	M	M	M	M	M	M	0.6734	M	M	0.3367	0.1282
Thalassemia	M	M	M	M	M	0.6818	0.2020	1.2878	0.0505	0.1010	0.0384
MDS	M	M	M	M	M	0.6818	0.2020	1.2878	0.0505	0.1010	0.0384
AIHA	M	M	M	M	M	1.3635	0.4040	2.5756	0.1010	0.2020	0.0769
Women < 45	0.0265	M	M	0.0644	M	0.0227	0.0067	0.0429	0.0017	0.0034	0.0013
Remaining Patients	0.0265	0.0543	0.1843	0.0644	0.2954	0.0227	0.0067	0.0429	0.0017	0.0034	0.0013

Table 25: Mismatch penalty matrix for the various antigens and patient group combinations. When a combination is marked with 'M' mismatching on this antigen is not allowed for the corresponding patient group. For the remaining combinations a weight is calculated based on the immunogenicity and pathogenicity and patient category weights.

We can compare this matrix to the original matrix supplied by the experts (shown in Table 22). The new matrix is constructed such that all "must" combinations are still a must. Note that when the number of must combinations is increased this will lead to more shortages as fewer assignments are considered valid. Furthermore, in general the constructed weights do correspond to the desirability levels as described in the original expert matrix. However, the weights of different combinations that are assigned the same desirability level in Table 22 can still have widely different values. The main cause of this is that we have used the relative immunogenicity values in the construction of Table 25. For example, when we look at the mismatch cost of $Jk(a)$ and $Jk(b)$ for *"Thalassemia"* patients we see that their corresponding weights are 1.2878 and 0.0505, whereas they are both assigned level 2 ("preferable") in the original matrix. The reason for this large difference is that the corresponding relative immunogenicity values of these antigens are 0.0837 and 0.0033 respectively. This means that the immunogenicity of $Jk(a)$ is estimated to be roughly 25 times greater than $Jk(b)$. This essentially says that a mismatch on $Jk(b)$ is far more unlikely to lead to alloimmunization. We think that including these empirically estimated relative immunogenicity values into this matrix is a worthful addition as it will likely more accurately correlate with the probability of alloimmunization, which is what we want to prevent after all. Another thing to remark is that patients in the group *"With Antibodies"* must be correctly negatively matched for the antigen(s) against which they have previously formed antibodies. This means that when the matrix above is used in practice, the antigen against which the *"With Antibodies"* patient has formed antibodies should also be treated as "must". Another possibility in practice is that some patients can belong to multiple categories. For example, it could be that a Thalassemia patient also has some antibodies. We recommend that for such patients the highest weight for each antigen is used.

8.2 Simulation Experiments for Single Hospital

To investigate how effective the constructed matrix of mismatch weights works in prioritizing special patient groups in an extended matching issuing strategy, we will perform simulation experiments. For two types of hospitals (regular and academic) we will run multiple one-year simulations to assess the average performance. To investigate the effect of the constructed matrix as shown in Table 25, we will run two variants of the same issuing strategy per hospital. Both variants use the same "must" combinations. Furthermore, one variant will use the normal relative immunogenicity weights as described in Table 5 and the other variant will use the patient category based weights as proposed in Table 25. We will use an inventory size equal to three times the average daily demand in these simulations. The reason for this is that this inventory size roughly corresponds to the inventory size that is found in these hospitals. Using this inventory size will give a more accurate indication of the practical matching possibilities within such hospitals.

8.2.1 Demand Scenarios

To simulate the demand for different types of patients we used data from two Dutch hospitals. One was a smaller regular hospital (OLVG Oost) and the other was an academic hospital (AMC). In Table 26 below we show the distribution of the different patient groups in these hospitals. These distributions are based on records of issued RBC units. When a unit was issued to a patient who was in multiple categories it was recorded to be in the category with the highest level of extensive matching. In the MINRAR model we explicitly model patients with requests for multiple units. The data in Table 26 is based on units issued and not on the number of patients. Therefore, this distribution may not be an accurate representation of the distribution of patients. However, we have no data available on the actual number of patients. This means that we will have to use the data available to us. Although the actual distribution of patients may differ, we argue that the distribution shown in the table is likely to correspond with the distribution of patients as it is not the case that the average number of units demanded per patient varies a lot between the different patient groups. Even if these percentages may not be exactly accurate, they should still give a meaningful estimate of the availability and matching quality of RBC transfusions for the patient groups considered.

Category	OLVG Oost	AMC
Patients without extended matching	88.69%	64.61%
Women <45	4.94%	10.25%
MDS	0%	4.54%
AIHA	1.49%	5.67%
With antibodies	3.16%	6.54%
SCD / Thalassemia	1.72%	11.33%

Table 26: Distribution of different patient categories for OLVG Oost and AMC.

For *"MDS"*, *"AIHA"*, *"SCD"* and *"Thalassemia"* patients we generated requests with one week between the request becoming known and the date that the demand must be satisfied. This is because in practice these transfusions are planned and therefore it would be incorrect to assume that these requests would become known on the day that they are needed. The lead time of one week is chosen as it is a period long enough ahead to make a request plannable, but also short enough to prevent an accumulation of future requests for these patient types. For the groups *"Patients without extended*

matching" and *Women<45* we assume that requests become known either on the day they must be satisfied or one day ahead, both with 50% probability. Lastly, patients from the *"With Antibodies"* group could have short or long lead times. Therefore, we have chosen to sample the lead times for this patient group at random between 1 and 7 days.

Because the data only contains a combined aggregate for the *"SCD"* and *"Thalassemia"* patient groups we have assumed that half of these patients are *"SCD"* patients and the other half are *"Thalassemia"* patients. Furthermore, all patient phenotypes were sampled from the Caucasian population with exception of the *"SCD"* patients for whom we used antigen prevalences for individuals of African descent, because SCD only occurs in that population. Both variants of antigen prevalences can be found in Appendix A.

To determine the number of units required per request for each patient group we used the distribution as described in Section 3.1.3 for the categories: *"Patients without extended matching"*, *Women<45* and *"With Antibodies"*. For the remaining categories we have generated requests with a demand of two units every time. The reason for this is that this is the most often requested quantity of units for these patient groups.

Lastly, we have assumed a daily demand of 50 for the regular hospital and a daily demand of 100 for the academic hospital. As the average number of units per request is about two, this comes down to roughly 25 patients per day for a regular hospital and about 50 patients per day for an academic hospital. The distributions used to sample demand for these values are the same as mentioned in Section 6.4.2. However, one change was made to the demand values sampled. This is because in our simulations we use an inventory with size equal to three times the average daily demand and therefore there is a small chance that the demand sampled for some days is larger than the total inventory size. This will lead to shortages that cannot be prevented. Therefore, when a particular day has a total demand larger than 2.5 times the average daily demand, we will omit random requests until the total demand is lower than or equal to 2.5 times the average daily demand. This will make sure that the probability of encountering shortages that cannot be prevented is low, and unlikely to affect the average performance.

8.2.2 Supply Scenarios

We used the same supply scenarios as described in Section 6.4.1. This means that the phenotype of the supplied RBC units is sampled according to the prevalence of antigens A, B and D in a historical dataset of transfused RBC units. The remaining antigens in the phenotype were sampled based on the prevalence of these antigens in the Caucasian population. We have chosen not to use AB-RhD specific order-up-to levels for these simulation for similar reasons as mentioned in Section 6.4.1.

8.2.3 Setup

As mentioned earlier, the inventory size in all simulations is assumed to be three times the average daily demand. At the start of the simulation the inventory was filled with units of age zero only. After an initialization period of 31 days the logging of results commenced and the simulation was run for one year from then onwards. Issued or outdated units are replaced at the end of each day with new units from the supply scenario.

In the simulation we will have to account for requests that are available but do not yet have to be satisfied with units. The way these are included in the optimization is as follows: We construct the ILP as described in Section 6.3 and we include all available requests. Then we alter the shortage penalty slightly, by making it twice its original size for all requests that have their due date on the current day. This will prevent that requests for the current day are left unsatisfied to accommodate for future requests. We

also add a constraint such that no units can be assigned to future requests if they will be outdated by then. Furthermore the MAS term, which is the term that penalizes minor antigen substitution, is only counted for *"Patients without extended matching"* patients and *"Women <45"*. This is because we do not want to prevent substitution for special patient groups, as the priority is to prevent mismatches for these groups and some substitution should be allowed if this can decrease the number of mismatches.

Now the ILP can be solved as before. When an optimal assignment is computed we discard all matches for requests that have their due date in the future. The remaining matches form the total assignment for the current day and are processed as such. The corresponding units are removed from inventory, we check for outdates and then the inventory is refilled with units from the supply scenario such that the total number of units again equals the original inventory size.

8.2.4 Results

Tables 27 and 28 show the percentage of shortages and antigen mismatches for each patient group when using the patient group specific mismatch weights and relative immunogenicity weights for all groups respectively. These results were obtained by averaging 25 simulations in which used all combinations of five demand and supply scenarios. The demand for the specific patient groups was generated according to the distribution of the OLVG Oost hospital, as described in Table 26.

We can compare Table 27 to Table 28 to see the improvement of minor antigen matching for special patient groups using the proposed weight matrix. Firstly, we can note that all must combinations are correctly matched in both variants. Therefore we focus on the antigens $Fy(a)$, $Fy(b)$, $Jk(a)$, $Jk(b)$, S and s. When comparing the two tables we can see that in Table 27 for every minor antigen and special type of patient the percentage of mismatches is reduced. To be able to better compare the difference we have constructed Figure 17. This figure shows the percentage decrease in mismatches that is the result of using the patient group specific weights compared to using the relative immunogenicity weights. We can see for antigens $Jk(a)$ and $Fy(a)$ the decrease is 85% or more for the patient groups *"With Antibodies"*, *"AIHA"* and *"Thalassemia"*. Patients from the *"SCD"* group must be compatibly matched on these antigens and therefore a decrease is not possible. Antigens $Fy(b)$ and S have a decrease varying between 60 and 80 percent and antigens $Jk(b)$ and s have the lowest decreases, varying between 20 and 60 percent. The magnitude of these observed decreases per antigen corresponds to the relative immunogenicity of the antigens as we see that the antigens with the highest immunogenicity have the largest percentage decrease in mismatches. We can also see that although the *"SCD"* patient group is given the highest weight of all groups, the decrease in mismatches for this group for antigens S and s is smaller than the corresponding decrease for the *"AIHA"* and *"With Antibodies"* groups. This is likely caused by that fact that the *"SCD"* group has the most "must" antigens that heavily reduces the number of units that are eligible for matching to this category, limiting the ability to perform high quality matching on antigens S and s. This shows that requiring more antigen combinations to be matched more compulsory does not necessarily improve the overall matching quality. Instead, it will likely lead to more shortages and lower quality matches for the remaining antigens.

Figure 17 also shows that when using the patient group specific mismatch weights, the minor antigen matching quality for patient groups *"Patients without extended matching"* and *"Women<45"* decreases slightly. This is to be expected when the mismatch costs for these antigens are relatively low compared to the other patient groups. Also, interesting to note is that the number of mismatches on antigens C and e for patients from the group *"Women<45"* heavily increases. This is caused by the dependencies between the

Regular Hospital Using Patient Group Specific Mismatch Weights

	Shortages	C	c	E	e	K
With antibodies	0.42 (0.32-0.52)	0	0	0	0	0
Sickle Cell Disease	1.56 (1.15-1.97)	0	0	0	0	0
Thalassemia	0.50 (0.22-0.79)	0	0	0	0	0
MDS	No patients	-	-	-	-	-
AIHA	0.28 (0.14-0.41)	0	0	0	0	0
Women <45	0.03 (0.00-0.07)	6.97 (6.57-7.37)	0	0	1.68 (1.39-1.97)	0
Remaining Patients	0.00 (0.00-0.00)	4.56 (4.40-4.72)	4.80 (4.61-5.00)	2.40 (2.31-2.50)	1.41 (1.35-1.48)	2.01 (1.94-2.08)

	Fy(a)	Fy(b)	Jk(a)	Jk(b)	S	s
With antibodies	1.02 (0.79-1.25)	2.83 (2.44-3.21)	0.48 (0.29-0.67)	7.95 (7.49-8.42)	4.64 (4.01-5.26)	6.78 (6.51-7.04)
Sickle Cell Disease	0	68.77 (66.44-71.10)	0	0	12.93 (11.03-14.82)	5.73 (4.65-6.80)
Thalassemia	1.39 (0.78-2.00)	3.83 (3.02-4.64)	0.45 (0.12-0.78)	13.09 (11.27-14.92)	6.99 (5.84-8.14)	6.80 (6.12-7.49)
MDS	-	-	-	-	-	-
AIHA	0.60 (0.34-0.85)	3.15 (2.54-3.76)	0.47 (0.23-0.72)	9.72 (8.94-10.50)	4.47 (3.75-5.19)	4.85 (4.43-5.27)
Women <45	9.06 (8.46-9.66)	8.67 (8.09-9.26)	4.18 (3.84-4.51)	13.63 (13.04-14.22)	17.32 (16.52-18.13)	8.88 (8.29-9.47)
Remaining Patients	10.29 (10.15-10.42)	8.72 (8.65-8.80)	6.31 (6.20-6.42)	14.57 (14.42-14.71)	18.34 (18.20-18.48)	8.81 (8.70-8.91)

Table 27: Average percentage shortages and antigen mismatches (95% CI) over 25 one-year simulations for a regular hospital (average daily demand = 50 units, inventory size = 150 units) where mismatches are penalized according to the weights in Table 25. Antigens which must be compatibly matched for certain patient groups are denoted with 0.

Regular Hospital Using Relative Immunogenicity Weights For All Groups

	Shortages	C	c	E	e	K
With antibodies	0.44 (0.34-0.55)	0	0	0	0	0
Sickle Cell Disease	1.29 (0.80-1.78)	0	0	0	0	0
Thalassemia	0.68 (0.38-0.98)	0	0	0	0	0
MDS	No patients	-	-	-	-	-
AIHA	0.30 (0.16-0.43)	0	0	0	0	0
Women <45	0.02 (0.00-0.04)	4.53 (4.12-4.95)	0	0	1.18 (0.94-1.42)	0
Remaining Patients	0.00 (0.00-0.00)	4.53 (4.37-4.69)	4.28 (4.10-4.47)	2.17 (2.08-2.26)	1.44 (1.38-1.51)	1.95 (1.88-2.02)

	Fy(a)	Fy(b)	Jk(a)	Jk(b)	S	s
With antibodies	9.16 (8.62-9.70)	10.39 (9.57-11.20)	3.86 (3.41-4.31)	17.74 (17.12-18.35)	15.94 (15.31-16.58)	9.79 (9.46-10.11)
Sickle Cell Disease	0	75.28 (73.65-76.92)	0	0	35.60 (32.78-38.41)	7.11 (5.46-8.76)
Thalassemia	10.72 (9.28-12.16)	9.81 (8.62-11.00)	4.23 (3.18-5.27)	18.34 (16.21-20.47)	19.96 (18.06-21.87)	8.54 (7.75-9.34)
MDS	-	-	-	-	-	-
AIHA	9.93 (8.68-11.17)	11.97 (11.02-12.92)	4.70 (3.78-5.62)	21.43 (20.60-22.27)	19.93 (18.62-21.23)	6.99 (6.45-7.52)
Women <45	9.38 (8.78-9.97)	8.62 (8.02-9.22)	5.52 (5.11-5.92)	13.06 (12.36-13.77)	16.82 (16.28-17.37)	8.57 (7.97-9.16)
Remaining Patients	9.56 (9.43-9.69)	8.28 (8.18-8.38)	5.98 (5.91-6.05)	14.21 (14.08-14.35)	17.35 (17.19-17.51)	8.74 (8.63-8.85)

Table 28: Average percentage shortages and antigen mismatches (95% CI) over 25 one-year simulations for a regular hospital (average daily demand = 50 units, inventory size = 150 units) where mismatches are penalized according to the weights in Table 5. Antigens which must be compatibly matched for certain patient groups are denoted with 0.

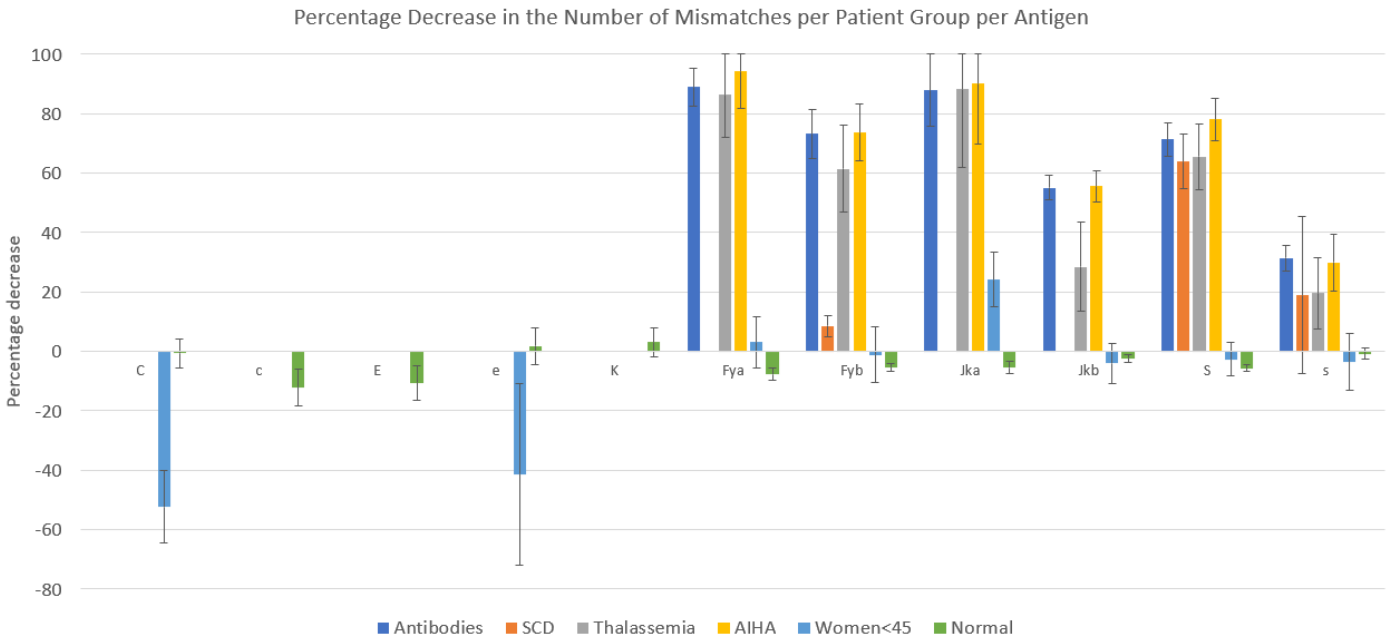

Figure 17: Percentage decrease (95% CI) in mismatches for each patient group for the different antigens when using the proposed category specific weights instead of equal weights for all patients. Results are calculated using simulations with demand scenarios sampled from the regular hospital demand (OLVG Oost).

different Rhesus antigens and is a direct result of the fact that this patient group must be compatibly matched on antigens c and E.

One value that stands out is the percentage of the *"SCD"* patient group that is mismatched on $Fy(b)$. In both Table 27 and Table 28 we see that roughly 70% of all SCD patients are mismatched on $Fy(b)$. This is caused by the fact that the phenotype of the SCD patients is sampled according to antigen prevalences for individuals of African descent. Among those, the $Fy(a-b-)$ double negative Duffy phenotype is common with a 68% probability of occurrence. In Caucasians however, this phenotype is so uncommon that the effective probability of occurrence is 0%. Because all supplied RBC units in our simulation are sampled according to antigen prevalences among Caucasians, it is therefore impossible to perform compatible matching on these two antigens for SCD patients with the double negative phenotype. Because the genotype $Fy(a-b+)$ is the most common and $Fy(a)$ has a higher relative immunogenicity we can expect that most SCD patients will be transfused with $Fy(a-b+)$ RBC units, and thus a mismatch will be induced on antigen $Fy(b)$. Fortunately the SCD patients with the $Fy(a-b-)$ phenotype can be safely matched $Fy(a-b+)$ RBC units as the same mutation causing the absence of both antigens also causes an inability to produce antibodies against $Fy(b)$[43].

Lastly, we can look at the availability of extensively matched blood for the different patient groups. Therefore we must look at the first columns in the upper parts of Tables 27 and 28. First, we can observe that the percentage of shortages is very low for both issuing strategies. Both have 0% shortages for regular patients and less than 0.5% shortages for every special patient group. When we compare the shortage percentages for these two issuing strategies, we do not see that one clearly outperforms the other. This is as expected as the number of shortages is a result of the MAS term in the

objective function, which minimizes minor antigen substitutions. The weights of the antigens in this term in both approaches was the same as the mismatch penalty of that antigen for regular patients. This means that the relative weight of the MAS term in both variants was equal and therefore no meaningful difference should be expected.

Academic Hospital Using Patient Group Specific Mismatch Weights

	Shortages	C	c	E	e	K
With antibodies	0.18 (0.15-0.21)	0	0	0	0	0
Sickle Cell Disease	5.38 (5.04-5.73)	0	0	0	0	0
Thalassemia	0.10 (0.05-0.15)	0	0	0	0	0
MDS	0.13 (0.08-0.19)	0	0	0	0	0
AIHA	0.13 (0.06-0.19)	0	0	0	0	0
Women <45	0.00 (0.00-0.00)	4.43 (4.25-4.62)	0	0	1.33 (1.27-1.39)	0
Remaining Patients	0.00 (0.00-0.00)	3.98 (3.91-4.05)	4.88 (4.79-4.98)	2.72 (2.64-2.81)	1.24 (1.20-1.28)	2.77 (2.70-2.84)
	Fy(a)	Fy(b)	Jk(a)	Jk(b)	S	s
With antibodies	0.89 (0.82-0.96)	1.08 (1.00-1.16)	0.23 (0.19-0.28)	5.72 (5.52-5.92)	3.13 (2.95-3.31)	4.03 (3.87-4.20)
Sickle Cell Disease	0	67.98 (67.56-68.39)	0	0	13.25 (12.72-13.77)	4.85 (4.70-5.00)
Thalassemia	1.66 (1.51-1.81)	1.69 (1.56-1.82)	0.29 (0.24-0.34)	8.40 (8.11-8.68)	4.74 (4.39-5.10)	6.38 (6.12-6.64)
MDS	1.80 (1.58-2.01)	1.68 (1.54-1.83)	0.37 (0.30-0.44)	8.77 (8.36-9.17)	4.58 (4.29-4.87)	5.67 (5.43-5.91)
AIHA	0.91 (0.79-1.04)	0.99 (0.82-1.16)	0.35 (0.26-0.44)	5.88 (5.44-6.32)	3.05 (2.73-3.38)	4.56 (4.21-4.90)
Women <45	11.06 (10.78-11.35)	7.14 (6.98-7.31)	5.08 (4.88-5.27)	14.07 (13.88-14.27)	16.70 (16.31-17.09)	8.06 (7.85-8.27)
Remaining Patients	11.05 (10.91-11.18)	7.86 (7.76-7.95)	5.69 (5.61-5.77)	14.01 (13.85-14.16)	17.14 (17.03-17.24)	8.26 (8.16-8.36)

Table 29: Average percentage shortages and antigen mismatches (95% CI) over 25 one-year simulations for an academic hospital (average daily demand = 100 units, inventory size = 300 units) where mismatches are penalized according to the weights in Table 25. Antigens which must be compatibly matched for certain patient groups are denoted with 0.

Academic Hospital Using Relative Immunogenicity Weights For All Groups

	Shortages	C	c	E	e	K
With antibodies	0.21 (0.16-0.25)	0	0	0	0	0
Sickle Cell Disease	5.13 (4.79-5.46)	0	0	0	0	0
Thalassemia	0.11 (0.06-0.16)	0	0	0	0	0
MDS	0.14 (0.08-0.20)	0	0	0	0	0
AIHA	0.08 (0.03-0.14)	0	0	0	0	0
Women <45	0.00 (0.00-0.00)	3.39 (3.27-3.50)	0	0	1.17 (1.10-1.25)	0
Remaining Patients	0.00 (0.00-0.00)	3.59 (3.52-3.65)	3.76 (3.67-3.85)	2.09 (2.02-2.15)	1.17 (1.13-1.21)	2.35 (2.27-2.43)
	Fy(a)	Fy(b)	Jk(a)	Jk(b)	S	s
With antibodies	7.77 (7.59-7.95)	4.10 (3.92-4.28)	2.39 (2.28-2.50)	11.68 (11.49-11.88)	9.95 (9.54-10.36)	6.85 (6.73-6.98)
Sickle Cell Disease	0	71.04 (70.63-71.45)	0	0	27.16 (26.68-27.65)	6.06 (5.92-6.19)
Thalassemia	8.27 (7.99-8.55)	4.63 (4.33-4.93)	2.47 (2.26-2.68)	12.62 (12.17-13.06)	10.89 (10.42-11.35)	8.38 (8.12-8.64)
MDS	8.38 (8.06-8.70)	4.69 (4.42-4.96)	2.74 (2.55-2.92)	12.89 (12.48-13.31)	10.74 (10.30-11.19)	7.69 (7.32-8.06)
AIHA	8.87 (8.24-9.50)	5.14 (4.72-5.56)	2.87 (2.63-3.12)	12.69 (12.04-13.34)	11.34 (10.73-11.96)	7.66 (7.25-8.06)
Women <45	8.28 (7.96-8.60)	6.28 (6.11-6.44)	4.68 (4.46-4.89)	12.73 (12.53-12.93)	13.34 (13.03-13.65)	7.92 (7.74-8.11)
Remaining Patients	8.28 (8.18-8.38)	6.56 (6.48-6.65)	4.60 (4.54-4.66)	12.44 (12.31-12.57)	13.57 (13.48-13.67)	7.85 (7.74-7.96)

Table 30: Average percentage shortages and antigen mismatches (95% CI) over 25 one-year simulations for an academic hospital (average daily demand = 100 units, inventory size = 300 units) where mismatches are penalized according to the weights in Table 5. Antigens which must be compatibly matched for certain patient groups are denoted with 0.

We have also simulated the use of the new antigen weights on scenarios that sampled patient groups from the distribution of an academic hospital. These contain more patients of special groups, but also have a larger demand volume and corresponding

inventory size. In Tables 29 and 28 below the average results are shown for these simulations. The average daily demand was 100 RBC units per day and the inventory size was 300 RBC units.

Table 29 shows the proportion shortages and antigen mismatches for the patient groups considered when using the patient group specific weights. Table 30 shows the same results of identical simulation which did not use the patient group specific weights. Figure 18 shows the relative decrease in the number of mismatches for the different patient groups. Overall we can observe a similar trend to that of the regular hospital. However, in Figure 18 we do see that the percentage of mismatches on all antigens for the *"Patients without extended matching"* group is increased more in the academic hospital compared to the regular hospital (Figure 17). The reason for this is that a larger proportion of the patient population of the academic hospital requires some form of extended matching. Therefore it will occur more often that one or more regular patients will have to be mismatched in order to prevent mismatches for these patient groups. Furthermore, we can observe in both Table 29 and 30 that the percentage of shortages is low for all patient groups. Only the *"SCD"* patient group cannot be allocated compatible RBC units in roughly 5% of all cases. This shows that the inventory size is not yet large enough to be able to safely accommodate for these requests without the risk of shortages. We also see that the *"MDS"* patient group, which was absent in the regular hospital, receives minor antigen matching of equal quality to the *"Thalassemia"* patient group. This is as expected since these patient groups were given equal importance in the matching in Table 23.

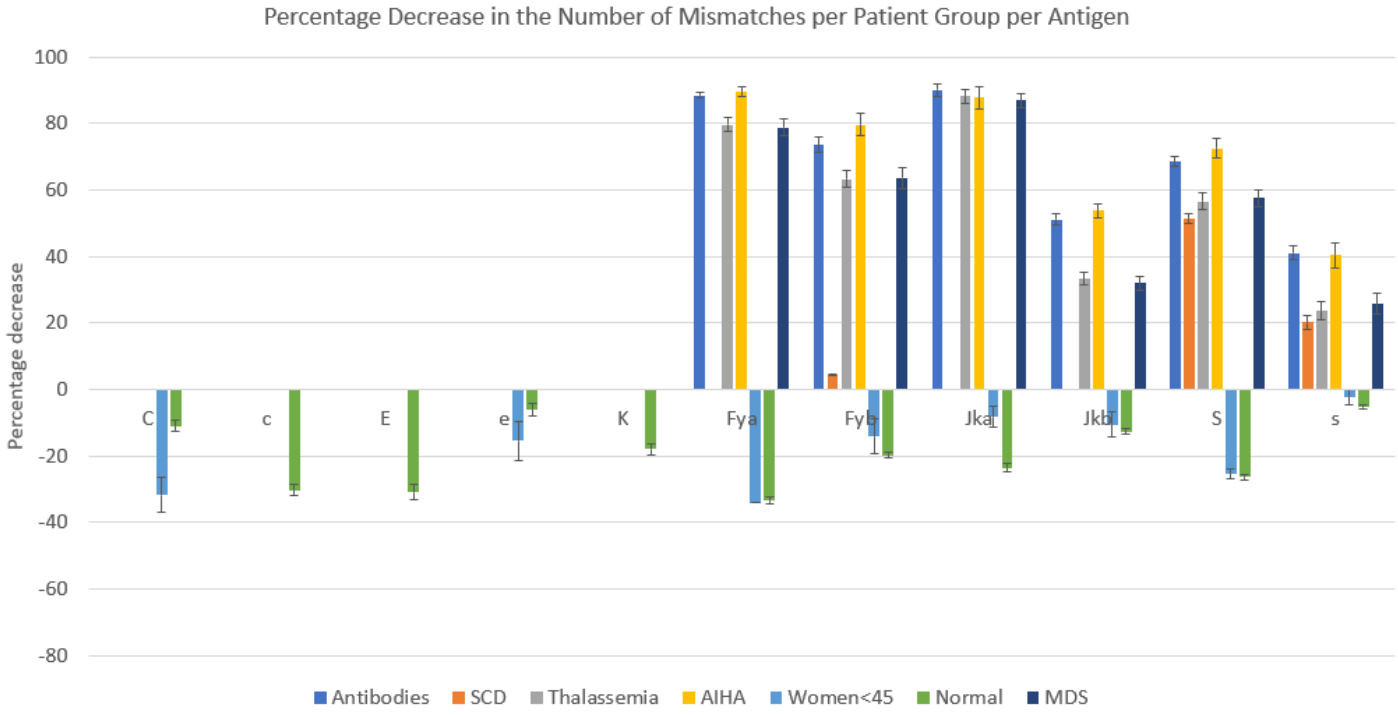

Figure 18: Percentage decrease (95% CI) in mismatches for each patient group for the different antigens when using the proposed category specific weights instead of equal weights for all patients. Results are calculated using simulations with demand scenarios sampled from the academic hospital demand (AMC).

8.3 Simulation Experiments for Multiple Hospitals

In practice, hospitals are not limited to their own inventory in a search for compatible RBC units for special patient groups. We have already seen in Chapter 6 that a larger volume of RBC units leads to a decrease in average relative alloimmunization risk per patient. In this section we will investigate how the availability and quality of compatible RBC units can be improved by including a distribution centre in the optimization process. This is achieved by allowing the cascading of patient's request to the distribution centre such that they can be directly matched with units from the inventory of the distribution centre. As such an approach concerns the management of more than one inventory, it is more complicated than the scenarios we have previously analysed.

The goal of this proposed strategy is to prioritize special patient groups. Therefore, they are given the privilege to improve their matching by being allowed to be matched with units from the distribution centre. However, we will not allow these patients to be directly assigned units from the distribution centre. Instead, these units are only allocated to the corresponding hospital where they can then be assigned to these patients on the next day. By not matching all patients on the distribution centre level we hope to be able to further improve the RBC unit availability and minor antigen matching for the special patient groups. To increase this effect even more we will penalize the allocation of highly usable units to the hospital when they must be resupplied. In this way we hope to keep these units in the distribution centre and available to special patient groups in all three hospitals.

8.3.1 Setup

In our analysis we will simulate three hospitals and one distribution centre. The distribution centre receives all supply and distributes it over the three hospitals (and no other hospitals). For two of these three hospitals we will use demand scenarios based on the patient distribution for regular hospitals (data from OLVG) and an average daily demand of 50 RBC units. The third hospital will have its demand based on the academic variant (data from AMC) and will have an average daily demand of 100 RBC units. All three hospitals will have an inventory size equal to three times the average daily demand, which roughly corresponds to the ratio seen in practice (usually between two and three). We also introduce a distribution centre, which supplies only these three hospitals. The average daily demand of the distribution centre is the sum of the average daily demands of the individual hospitals and is therefore equal to 200 RBC units per day. The inventory size of the distribution centre is five times this demand, in accordance with current practice in distribution centres (average issuing age in distribution centres is about five days). In the setup we allow requests for the patient groups *"Women<45"*, *"Thalassemia"*, *"AIHA"*, *"SCD"*, *"With Antibodies"* and *"MDS"* to be cascaded from the hospitals to the distribution centre to improve the matching. This setup is shown schematically in Figure 19.

Each hospital is assigned a demand scenario as described in Section 8.2.1 that specifies the demand for each day, as well as on which day requests will become available. Each hospital has to compute an assignment of RBC units to patients to satisfy the demand of the current day. We have assumed that there is no supply during the day and therefore requests that are due on a particular day must be satisfied with units present in the hospital inventory on that day.

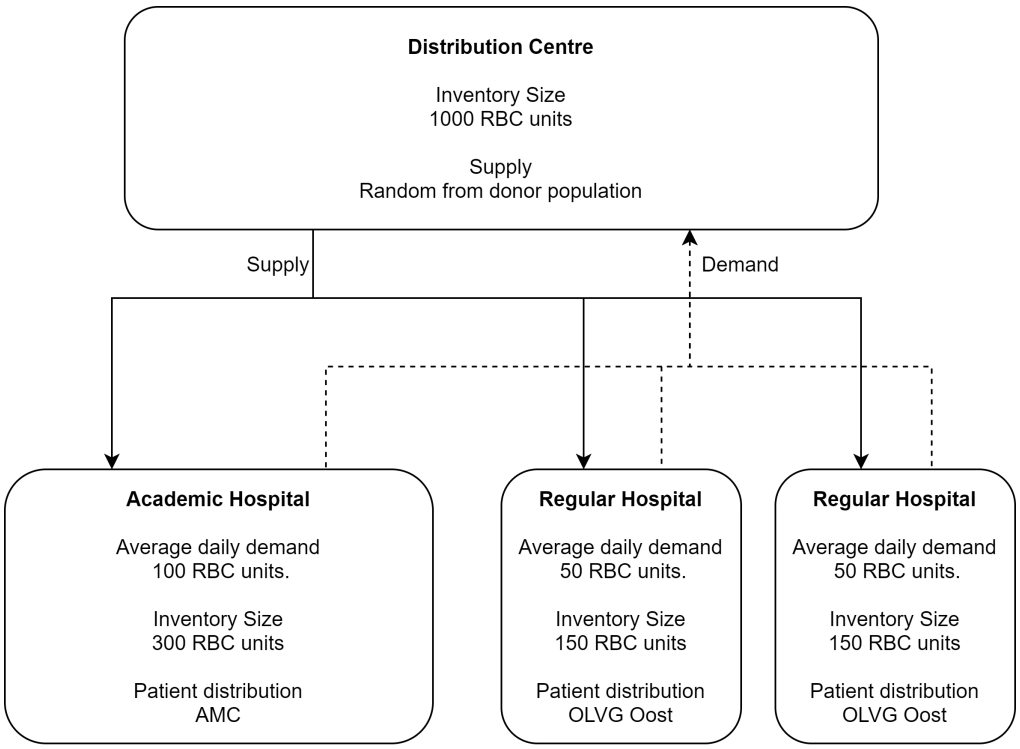

Figure 19: Schematic setup of the multi-hospital simulation study. The three hospitals receive their supply from the distribution centre. Requests for the patient groups *"Women<45"*, *"Thalassemia"*, *"AIHA"*, *"SCD"*, *"With Antibodies"* and *"MDS"* can be cascaded down from the hospitals to the distribution centre to improve the ability of matching.

On each day we will use a four-step process to assign RBC units to patients in the hospitals and resupply the hospitals with units from the distribution centre.

- **Step 1.** Compute best possible matches for all patients within each hospital.
- **Step 2.** Propagate unmatched or non-perfectly matched patient requests to the distribution centre such that better units may be ordered.
- **Step 3.** Compute a final assignment in each hospital, using the information about ordered units in the distribution centre.
- **Step 4.** Compute an assignment of units from the distribution centre to the hospitals to replenish their inventories.

We will elaborate further on how each of these steps are modelled.

Step 1 In this first step each hospital use the MINRAR-Online algorithm with weights as specified in Table 25 to compute an assignment of RBC units to all patient requests currently known. However, an altered objective function is used such that requests on the current day have twice the shortage weight compared to future requests. This will ensure that there are no shortages incurred on the current day because of units that are reserved for future requests. Furthermore the FIFO and UAD terms are given a weight of zero and the MAS term is only counted for "Patients without extended matching" and *"Women <45"*. The reason for this is that we want to get an indication of

the best possible match of the units in inventory for these patients. After optimization, for each patient in the categories *"SCD"*, *"AIHA"*, *"MDS"*, *"Thalassemia"* and *"With Antibodies"* we store whether or not they are satisfied with RBC units and if they are we also store their current total mismatch penalty. Furthermore for *"Women <45"* we store whether or not their requests are satisfied.

Step 2 In this step all non-regular patients whose requests were not satisfied or not optimally satisfied are gathered. Only those patients are considered who have a future due date, as patients whose due date is on the current day cannot be assigned units form the distribution centre. *"Women <45"* are only included if there are no units assigned in their hospital. We again use the MINRAR-Online algorithm to compute an assignment of RBC units from the distribution centre to this new set of requests. For each patient we add a constraint that the total mismatch penalty should be less than the mismatch penalty that was obtained in the hospital. If a patient was unsatisfied in its hospital then this constraint is not added, because any assignment will be an improvement. Again we use an altered objective function. Requests that are due the next day are given double shortage cost, for similar reasons as mentioned earlier. Furthermore the FIFO and UAD terms are again given zero weight and the MAS term is only included for requests for *"Women <45"*. An optimal assignment is computed and all assigned units are marked as such. All requests that are satisfied are marked 'assigned unit(s)' to prevent that a unit in the hospital is allocated as well.

Step 3 This step is very similar to Step 1. The difference is that now in each hospital it is known which future requests have units reserved in the distribution centre. These requests are ignored and for the remaining requests the MINRAR-Online algorithm is used to compute an assignment. Again, requests that must be satisfied on the current day have a double shortage penalty. Furthermore, all requests for which we previously stored their mismatch penalty are now given the constraint that the mismatch penalty to be computed may not exceed this value. This time, the FIFO and UAD terms are given full weight. The MAS term is only counted for *"Patients without extended matching"* patients and *"Women <45"*, similar to Step 1. Next, the model is optimized and an assignment is computed. When extracting the solutions, all matches made for future requests are ignored. The remaining assignments are processed and the remaining units in inventory are check for outdating.

Step 4 In this final step an assignment of units from the distribution centre to the hospitals must be computed. First, all units assigned to a request with a due date equal to the next day are automatically added to the delivery to the corresponding hospital. Next, per hospital the number of units to be supplied is counted in order to refill their inventories to the predetermined size. Then a simple ILP is constructed that computes an assignment of RBC units to hospitals such that all hospitals receive the requested number of units. The objective function of this ILP consists of two terms: FIFO and UAD. The FIFO term is equal to how it was defined in the MINRAR-Online ILP, but the UAD term is slightly different. Instead of using antigen set $\mathcal{A}_3$ (A, B, RhD) to compute the usability loss, we use $\mathcal{A}_{minor} = \{C, c, E, e, K, Fy(a), Fy(b), Jk(a), Jk(b)\}$. The reason for this is that we would like more antigen negative units to stay in the distribution centre. Antigens A, B and RhD are excluded is to prevent an undersupply of the O and Rhesus negative blood groups to the hospitals. Antigens S and s are also excluded as they have less relevance in the matching and they are no must for any patient group. Furthermore, including these would also mean that units that were otherwise regarded as highly compatible will have a reduced compatibility if they are positive for both S and s. We think this is undesirable as these units are compatible for all patient

groups and thus we do not want to reduce the incentive to keep them in the distribution centre. Including antigens S and s in the antigen set may or may not result in better overall performance, however this was not tested. Note that these antigens are not ignored in the matching of units to patients. Instead, they are not specifically included in the process of saving antigen negative units in the distribution centre. Lastly, when an optimal assignment has been computed, the distribution inventory will be replenished with random units from the supply scenario.

8.3.2 Results

We have run 25 simulations of the multi-hospital setup. We will first compare the availability of RBC units for the different patient groups. Table 31 shows the average percentage of shortages for the small regular hospital in the single-hospital setup that was investigated in the previous section, as well as the average percentage of shortages in the similarly sized two regular hospitals in the multi-hospital setup scenario. The percentage decrease in shortages is shown, as well as the percentage of requests in the regular hospitals that were allocated and reserved in the distribution centre. Table 32 shows the same results for the academic hospital.

RBC Availability For Patient Groups in Regular Hospitals

	Shortages single-mospital Setup	Shortages multi-hospital Setup	Percentage Decrease	Percentage Units Reserved in the Distribution Centre
With antibodies	0.42 (0.32-0.52)	0.18 (0.09-0.26)	57.93 (53.42-62.44)	20.80 (19.62-21.97)
Sickle Cell Disease	1.56 (1.15-1.97)	0.27 (0.03-0.52)	82.60 (78.29-86.92)	38.31 (36.49-40.13)
Thalassemia	0.50 (0.22-0.79)	0.04 (0.00-0.13)	91.61 (83.18-100.00)	26.51 (23.95-29.06)
MDS	-	-	-	-
AIHA	0.28 (0.14-0.41)	0.03 (0.00-0.09)	89.73 (82.04-97.41)	22.77 (21.41-24.14)
Women <45	0.03 (0.00-0.07)	0.17 (0.10-0.24)	-412.99 (-446.70 – -379.29)	0.26 (0.19-0.32)
Remaining Patients	0.00 (0.00-0.00)	0.03 (0.00-0.05)	-	-
Total	0.04 (0.03-0.04)	0.04 (0.02-0.06)	-5.37 (-14.45-3.70)	1.53 (1.47-1.59)

Table 31: Percentage of shortages (95% CI) for simulations of a regular hospital (OLVG Oost / average daily demand = 50 RBC units). Per patient group we show the shortages for the single- and multi-hospital setup, the percentage decrease when using a multi-hospital setup and the percentage of patients who received units that were allocated and reserved at the distribution centre.

RBC Availability For Patient Groups in Academic Hospital

	Shortages for single-hospital Setup	Shortages for multi-hospital Setup	Percentage Decrease	Percentage Units Reserved in the Distribution Centre
With antibodies	0.18 (0.15-0.21)	0.06 (0.04-0.09)	65.13 (61.98-68.28)	17.71 (17.32-18.10)
Sickle Cell Disease	5.38 (5.04-5.73)	0.48 (0.41-0.54)	91.15 (90.23-92.07)	40.27 (38.71-41.84)
Thalassemia	0.10 (0.05-0.15)	0.01 (0.00-0.03)	87.00 (79.81-94.20)	29.69 (29.01-30.38)
MDS	0.13 (0.08-0.19)	0.00 (0.00-0.00)	100.00 (94.65-100.00)	29.32 (28.64-30.00)
AIHA	0.13 (0.06-0.19)	0.01 (0.00-0.03)	89.52 (81.80-97.24)	21.00 (20.12-21.89)
Women <45	0.00 (0.00-0.00)	0.01 (0.00-0.02)	-	0.33 (0.28-0.38)
Remaining Patients	0.00 (0.00-0.00)	0.01 (0.00-0.01)	-	-
Total	0.34 (0.31-0.36)	0.04 (0.03-0.04)	88.53 (87.58-89.48)	7.16 (7.02-7.30)

Table 32: Percentage of shortages (95% CI) for simulations of a regular hospital (AMC / average daily demand = 100 RBC units). Per patient group we show the shortages for the single and multi-hospital setup, the percentage decrease when using a multi-hospital setup and the percentage of patients who received units that were allocated and reserved at the distribution centre.

Both Tables 31 and 32 show that the number of shortages in all patient groups decreases when the individual hospitals can propagate patient requests to the distribution centre. There is only one exception, as we see an increase from 0.03% to 0.17% in shortages for the *"Women <45"* category in the regular hospitals. This increase of roughly 14 in 10,000 patients is most likely caused by the fact that 50% of the requests for this patient group must be satisfied on the day of request. Thus, a compatible unit must be in stock as it cannot be ordered from the distribution centre. Because the distribution centre tries to save antigen-negative units we expect fewer cEK negative units to be in the hospital inventories, which is the likely cause of the increase in shortages. We also see that in the academic hospital there is an increase from 0.00% to 0.01% for *"Women <45"*, which probably has the same cause, but a less pronounced impact as the inventory size of the academic hospital is larger. In both tables we can also see a slight increase in the number of shortages for the *"Patients without extended matching"* group. For both types of hospitals this increase is not significant. It is not immediately clear why there should be an increase, but a possible cause is that the patient groups matched at the distribution centre have no UAD substitution penalty and therefore can be matched with relatively much O^+ and O^- blood, thereby decreasing the availability of these blood groups for regular patients.

The bottom rows of Tables 31 and 32 show the overall results for both hospital types. For the regular hospital we can see that the percentage of shortages remains constant, at about 0.04%. A five percent increase in shortages is computed, but the 95% confidence interval of (-14.45% – 3.70%) shows that a decrease is not certain. On the other hand, the total number of shortages at the academic hospital does decrease. The percentage of shortages drops from 0.34% (0.31% – 0.36% 95% CI) to 0.04% (0.03% – 0.04% 95% CI), showing that the availability of RBC units does increase significantly. This is mainly caused by the fact that the availability of RBC units for the *"SCD"* patient group increases by 91.15% (90.23% – 92.07% 95% CI) as over 40% of all SCD patients receive RBC units matched from the distribution centre.

To assess the benefits of a multi-hospital setup in terms of minor antigen matching quality we have constructed Figures 20 and 21. These figures show the decrease in the number of antigen mismatches for each patient group when using the multi-hospital setup compared to the single-hospital setup for the regular and academic hospitals respectively.

Figure 20 shows that all the patient groups requiring extensive matching (except for *"Women <45"*) improve on the minor antigen matching quality. Especially the *"SCD"* patient group benefits by removing more than 70% of the mismatches on S and more than 60% of the mismatches on s. Furthermore, we see that the *"Patients without extended matching"* and *"Women <45"* patient groups pay for the benefits of the other patient groups. Especially the increase in mismatches for the antigens E and K is large for the regular patients. However, these values are relative and thus the absolute percentage of mismatches for these patient groups might still be low (we will show the average percentage of mismatches for each patient group in Figure 22). The same figure is constructed to see the improvements at the academic hospital (Figure 21). Here we can still see an overall improvement in the minor antigen mismatching, although the number of mismatches on $Fy(b)$ and $Jk(a)$ actually increases for patient groups *"With Antibodies"* and *"MDS"*. Also, the *"Thalassemia"* patient group shows a slight increase in the number of mismatches on $Jk(a)$. Again these increases are relative, and we would like to refer back to Table 29 where it can be seen that the percentage of mismatches on $Jk(a)$ in the single-hospital setup for patient groups *"With Antibodies"*, *"MDS"* and *"Thalassemia"* were 0.23% 0.37% and 0.29% and therefore already very small. This means that the increases shown in the figure do not correspond to many patients at all.

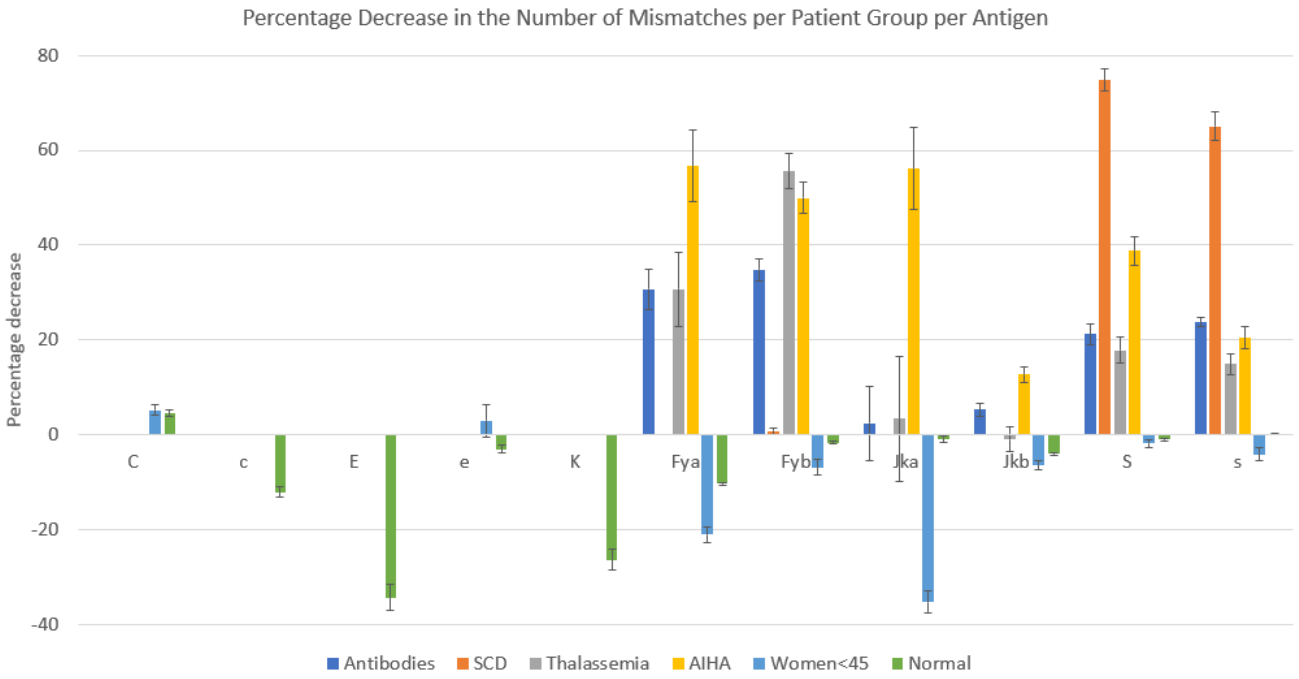

Figure 20: Percentage decrease in mismatches (95% CI) per patient group for the various antigens in the regular hospital (OLVG Oost) using the multi-hospital setup compared to the single-hospital setup.

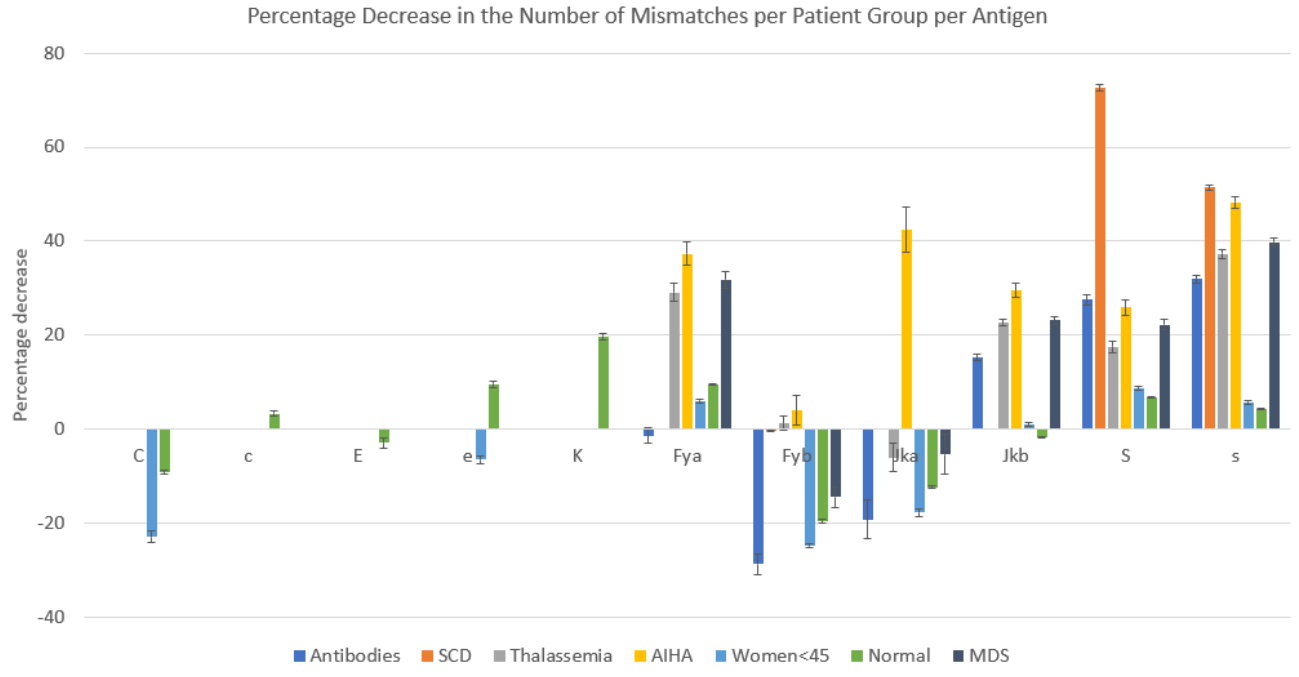

Figure 21: Percentage decrease in mismatches (95% CI) per patient group for the various antigens in the academic hospital (AMC) using the multi-hospital setup compared to the single-hospital setup.

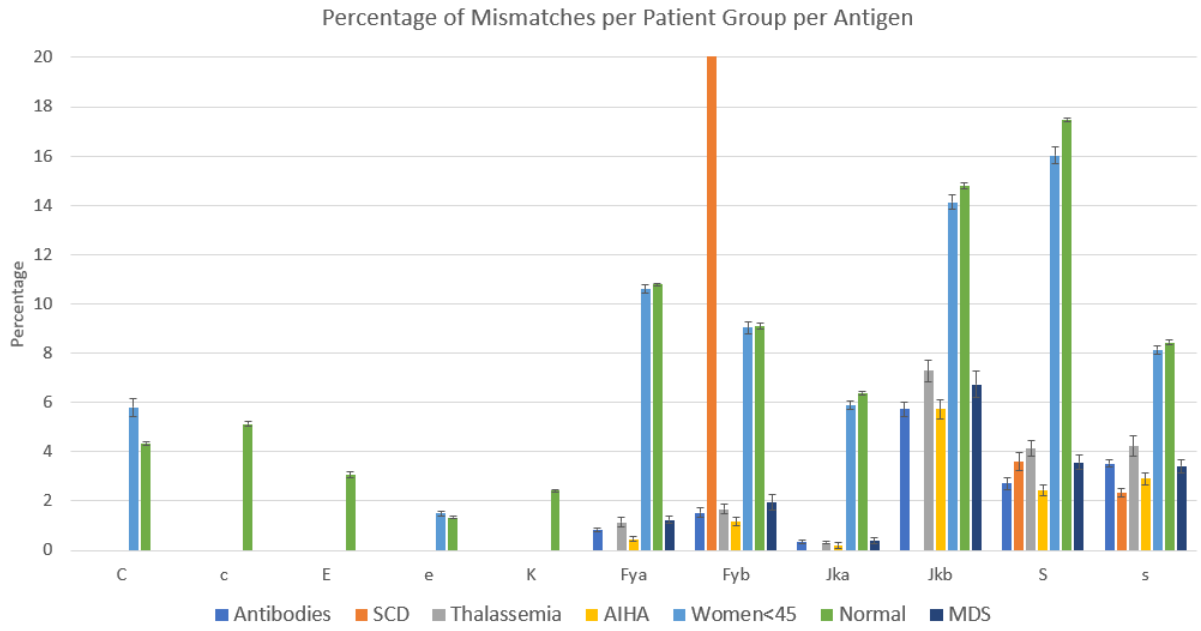

Figure 22: Percentage mismatches per antigen per patient group for the multi-hospital setup. Values shown are % (95% CI). Fy(b) mismatches for the SCD patient group are 68.24% (67.76% – 68.72% 95% CI) and not fully shown in the figure.

Finally, we have constructed Figure 22, which shows the percentage of patients from each group that is mismatched per antigen. We have limited the figure to 20%, as all values are below this, except for the *"SCD"* patient group, that is mismatched for Fy(b) 68.24% (67.76% – 68.72% 95% CI) of the time, for reasons previously explained. This figure shows the overall performance of minor antigen matching for the different patient groups in a multi-hospital setup. It shows that for all special patient groups (*"With Antibodies"*, *"SCD"*, *"Thalassemia"*, *"AIHA"* and *"MDS"*) the number of mismatches on the minor antigens can be kept small, while having almost no shortages, as shown in Tables 31 and 31. Furthermore, for these patient groups, mismatches on Fy(a) are kept below 1%, mismatches on Fy(b) and Fy(a) are kept below 2%, mismatches on S and s are (almost) kept below 4% and mismatches on Fy(b) are kept below 8%. Lastly, mismatches for the *"Patients without extended matching"* patients on the most immunogenic antigens K and E are kept below 3% and 4% respectively.

8.4 Discussion

In this chapter we have discussed how the MINRAR-Online ILP formulation can be used in combination with patient group specific mismatch costs to minimize the risk of alloimmunization, especially for those patient groups for which alloimmunization is more severe. We used empirical data on the relative immunogenicity, together with estimates of pathogenicity and relative patient group importance weights to reconstruct a matrix constructed by experts in immunology. This matrix contained a categorization of the desirability of the different antigen-patient group combinations in four levels: 4 (must), 3 (important) 2 (preferred) and 1 (if possible). We discussed why it was not possibly to directly use this matrix with desirability levels in the optimization process. Instead, by

reconstructing it using the relative immunogenicity values we have added new information to the matrix that will likely increase the proportion of prevented alloimmunization among patient when used in optimization. In Section 8.2 we investigated how many more mismatches can be prevented for the patient groups that require extensive matching when using the proposed matrix and how this also leads to a small increase in the number of mismatches for the regular patients. However, as we already noted in Chapter 3, extended matching is many times more important for patient groups at risk than it is for regular patients. The simulations run to test the performance of the proposed matrix versus the "normal" relative immunogenicity weights relied on many assumptions. First, we made very crude estimations of the distribution of lead times for the different patient groups considered. These estimates were mostly based on the fact that the demand for the patients who require extensive matching is usually known in advance. We chose one week of lead time between the request becoming known and the moment that one or more units have to be assigned for the patient groups *"Thalassemia"*, *"MDS"*, *"SCD"* and *"AIHA"*. We are aware that in practice there may be much variation in this, making the planning sometimes more difficult or easier. Secondly, we assumed that some patient groups always have demand for two units, whereas other groups have more variation in the number of units requested. The reason for this assumption was that in practice it is unlikely that SCD or Thalassemia patients have requests for more than two units. These patient groups are usually transfused on a regular basis as part of their treatment. Patients who demand more than two units are usually in surgery where heavy bleeding is a possibility. As this is not the case for the patient groups who receive transfusions as part of their treatment, we have therefore fixed their demand on two units per request.

The results of the simulations of this single-hospital setup show that for both the regular and academic hospitals the number of minor antigen mismatches can be reduced using the patient group specific mismatch weights. It further demonstrates the trade-off between inducing mismatches for regular patients and preventing them for the other patient groups. The main goal of these simulations was not to show specifically what is achievable, but instead how patient group specific mismatches can reduce alloimmunization risk for the patient groups that will benefit most from such a reduction. It may be clear that the estimated values used to reconstruct the patient group specific weights have a large impact on the final performance of the model. When large-scale extended antigen matching becomes available in the future, more work should be done to compute a suitable set of weights which provide a more accurate representation of the matching priorities for the different patient groups.

In Section 8.3 we combined two regular hospitals with a patient distribution according to the OLVG Oost hospital together with one academic hospital that uses the AMC patient distribution. These three hospitals are supplied from one distribution centre, which has an inventory size equal to five times the average daily demand of the three hospitals combined. A four-step allocation method was used to compute assignments of RBC units to patients within each hospital on each day, as well as improve matches for future requests by propagating these to the distribution centre. This propagation was restricted to patients from the groups *"Women <45"*, *"MDS"*, *"AIHA"*, *"SCD"*, *"Thalassemia"* and *"With Antibodies"*. The aim was to limit the number of requests propagated to the distribution centre in order to improve the quality of these requests. Furthermore, we used a simple ILP with a two-term objective function to replenish the inventories of the hospitals with units from the distribution centre. This two-term objective function aimed to maximize both the age of the units issued as well as the minor antigen usability of the units remaining in inventory. To determine this usability we used the antigen set $\mathcal{A}_{minor} = \{C, c, E, e, K, Fy(a), Fy(b), Jk(a), Jk(b)\}$. By deliberately not

including antigens A, B and D we aimed to prevent an undersupply of highly usable blood groups such as O^- and O^+ to the individual hospitals. We did not investigate the influence of using this antigen set as compared to other selections possible. However, we see that the total percentage shortage for both types of hospitals in the multi-hospital setup is $< 0.04\%$, which indicates that only a handful of shortages occurred during the entire one-year simulation. When large scale extended antigen matching becomes a reality in the future, a similar question will arise. There will be an incentive to save highly compatible blood in the distribution centres as much as possible to increase the number of patients to which is can possibly be matched. When exactly to save this blood depends on the usability of the blood. However, it is not only a function of the usability. This is because blood which has a high usability can still be positive for a few highly immunogenic antigens. Similarly, some AB^+ blood may still have high usability due to the absence of many minor antigens. Still it may not be useful for saving, as only a handful of patients will be able to receive the product. Deciding which blood should be stored separately and marked as highly usable is an interesting question which requires more thought than given in this Chapter. The basic ingredients for deciding whether a unit is useful are presented in this thesis as both the relative immunogenicity as the usability play a role in this process. Future, possibly simulation based, research is needed in order to create suitable guidelines which can effectively maximize the availability of highly compatible RBC units for as many patients as possible.

Finally, the simulations of a multi-hospital setup should not be thought of as an exact simulation of future RBC matching in the blood supply chain. Instead, it is a relatively simple yet effective model that should capture the essence of RBC matching for different patient groups in a more practical environment where hospitals are not limited to use only their own inventory. This model did not use order-up-to levels for the individual hospitals as this would substantially complicate the simulation as we first have to decide what suitable order-up-to levels would be for the individual hospitals, and secondly how the MINRAR-Online ILP should be altered in order to work well with these order-up-to levels. The aim of this research is to provide an indication of the feasibility of large-scale extended matching. Therefore, we argue that these assumptions are justified to be able to make a simplified model of the real-world, in order to show what is possible under those circumstances. When extended matching becomes available for more and more patients in the future, more thorough research is necessary to support the implementation of automated extended matching, as well as to assist the formulation of new guidelines for inventory management in hospitals and blood bank distribution centres.

9 Summary

The most common blood products transfused are red blood cell units. The Dutch blood bank Sanquin is responsible for managing a steady supply of RBC units and distribution of these units over all Dutch hospitals. The hospitals are responsible for matching units from inventory to their patients. The current matching policy for assigning RBC units to patients is to perform fully compatible issuing on major blood groups (antigens A, B and RhD). However, many more than these three antigens are found in human blood. Eleven of these minor antigens have enough clinical importance to be considered in practice. When a patient is transfused with an RBC unit that is incompatible on a minor antigen, the immune system of the patient may form antibodies against it, which is known as alloimmunization. The probability of alloimmunization given a mismatch is expressed in the immunogenicity of the antigen and is generally very low. However, once a patient is alloimmunized for a certain antigen, all subsequent blood transfusions cannot mismatch on this antigen as this leads to a transfusion reaction which can have severe consequences. Therefore, there is an incentive to prevent alloimmunization, especially for patients who require regular transfusions due to a genetic disorder or illness. Preventive matching on antigens c, E and K is also advisable for women within reproductive age to avoid complications during future pregnancies.

The aforementioned groups currently already receive extensively matched blood products. In order to perform extensive matching, the blood of both donor and patient has to be tested for the presence of minor antigens. This is done using serological tests which are time consuming and costly and therefore only used when necessary. However, advances in genotyping technology have allowed for the manufacturing of chips that can determine the full antigen phenotype of an individual with a single test. A future where all donors and patients can be typed for all relevant antigens is now foreseeable. When each patient could receive an RBC unit compatible on all relevant antigens, the risk of alloimmunization may be eliminated. However, large-scale extensive matching of RBC units to patients will also create new challenges due to an increase in the number of different blood groups, which grows exponentially with the number of antigens considered. This makes it highly unlikely that hospital inventories will contain units compatible on the three major and all eleven minor antigens for every patient. Furthermore, large-scale extended matching should not lead to an increase in shortages, as all matches compatible on A, B and RhD should still be considered valid. Lastly, as RBC units have a maximum shelf life of 35 days, an extensive matching approach should not lead to an increase in outdating. Thus, an ideal issuing strategy should minimize minor antigen mismatches while still allowing them if necessary, to avoid shortages or to avoid RBC units from outdating.

Currently, units are assigned to patients by manually selecting a suitable unit for each patient. When extensive matching becomes available for all patients such a manual sequential assignment strategy is likely no longer optimal and mathematical optimization is necessary. In previous work on this subject a model (FIFO/MROL) was proposed by Van Sambeeck *et al.* [6] which computes an assignment for a group of patients simultaneously by transforming the problem to a graph and computing an optimal assignment using a minimum cost maximum flow algorithm. This approach required a set of antigens to be chosen before optimization which must be compatibly matched. The remaining antigens were ignored in the optimization. This has the disadvantage of discarding some AB-RhD compatible matches as well as excluding the possibility of matching on the antigens ignored. In this thesis we have proposed a new Integer Linear Programming formulation called MINRAR which removes both these disadvantages. All AB-RhD compatible matches are considered valid and instead of discarding incompatible mismatches on minor antigens these are allowed however, at a cost. For this cost

we used the relative immunogenicity of each antigen, which is the normalized variant of the immunogenicity. This means that the issuing strategy has maximum flexibility while minimizing the overall risk of alloimmunization. Furthermore, our approach uses an improved definition of shortages as we explicitly consider patients with demands for multiple units, which concerns most patients. Later we have extended the MINRAR formulation to also contain terms which penalize the issuing of fresh units and the unnecessary issuing of rare blood units to improve performance in the long run. This extension was called MINRAR-Online.

The performance of the MINRAR-Online ILP was compared to three variants of FIFO/MROL issuing strategy using one-year simulations of various hospital sizes. The results show that the MINRAR-Online ILP is able to reduce the alloimmunization risk whilst avoiding shortages and outdating. We have constructed figures that show that our approach is able to match blood products for a large proportion of patients with low alloimmunization risk. These figures also show that a small proportion of patients is mismatched on the more immunogenic antigens, which indicates that these mismatches are sometimes unavoidable in order to prevent shortages. When larger demand and corresponding inventory sizes are considered, our approach is able to issue more and more patients with zero alloimmunization risk. When daily demand is 500 RBC units (inventory size 2500 units) more than 70% of patients are matched with 0% alloimmunization risk and more than 90% receive a match compatible on the five most immunogenic minor antigens (K, E, e, c and $Jk(a)$) without any shortages.

Next, we constructed a new ILP to compute the best possible performance of any issuing strategy over a given period in the past. The result was compared to the performance of the MINRAR-Online issuing strategy on the same input data. Simulations show that when the weight of the UAD term in the objective is tuned optimally, the performance of the MINRAR-Online strategy is likely no more than 1 percentage point off the optimal issuing policy for when the daily demand is 50 or 100 units. This difference roughly corresponds to a mismatch on antigens S or $Fy(b)$ which both only have low clinical relevance, thus showing that the MINRAR-Online strategy is able to produce near-optimal matching results for the majority of antigens with clinical relevance.

Lastly, we have discussed how to adapt the MINRAR-Online issuing strategy for more realistic use in hospitals where different patient groups have differing priorities for extended matching. Demand for a regular and an academic hospital was modelled by differentiating between the following patient groups: SCD patients, Thalassemia patients, patients with known antibodies, patients with AIHA, patients with MDS, women <45 and regular patients. We combined the relative immunogenicity with expert estimates on the desirability of compatible matching for each combination of antigen and patient group to compute a weight per combination that expresses the severeness of that specific mismatch. We showed that using these patient group specific weights, the percentage of patients mismatched for minor antigens for certain patient groups is reduced, although at the cost of an increase in mismatches for regular patients. Next, we investigated how a multi-hospital setup, where requests for some patient groups can be propagated to a distribution centre, can increase the availability of compatible RBC units for these patient groups. Simulations show that the possibility of reserving units in the distribution centre decreases the (already small) percentage of shortages for various patient groups. Furthermore, results show that overall, less than 1% of the patients who require extensive matching are mismatched on Jk(a). Mismatches on Fy(a) and Fy(b) are kept below 2% and mismatches on S and s are kept below 4% for nearly all patient groups. Mismatches on Jk(b) occur most frequently (about 5%-8%) of the time due to the low immunogenicity of Jk(b). Lastly, the overall availability of RBC units for each patient group is more than 99%, meaning that shortages occur only for less than 1% of all patients per group.

10 Discussion

Advances in genotyping technologies have opened up the possibility of large-scale extensive antigen matching in RBC transfusions. When all patients can receive extensively matched RBC units compatible on the fourteen clinically most relevant antigens then this would essentially prevent all risk of alloimmunization and thereby improve the quality and effectiveness of RBC transfusion practice. However, large-scale extensive matching also leads to more complexity in RBC inventory management. Manual matching of compatible RBC units to patient is likely no longer optimal. Furthermore, care should be taken to ensure that the availability of RBC units does not decrease when matches are incompatible on minor antigens. In this thesis we have explored how inventory management of extensively typed RBC units can be mathematically optimized to prevent shortages, outdating and alloimmunization risk in the long run.

By allowing all AB-RhD compatible matches, we consider all matches valid in line with current RBC matching guidelines. By assigning a cost to a mismatch between a unit and a patient on a particular minor antigen based on the immunogenicity of this antigen, we can effectively minimize the risk of antibody formation without an increase in the number of shortages or outdates. For this cost we have used the immunogenicity as was estimated by Evers *et al.* [3]. We have also used these values to assess the quality of an assignment and estimate the relative risk of antibody formation. Antigens that were assigned an immunogenicity of zero will have therefore also have zero risk of antibody formation in our model. This may be considered an oversimplification, as in practice there could still be some antibody formation against these antigens. However, we want to emphasize that our model can take any number of antigens with corresponding immunogenicity values as input. Therefore, it is not limited to the set of antigens and values used in this thesis.

We have explicitly modelled patients with demand for multiple RBC units. After discussion with experts, we have assumed that patients can only be mismatched once per antigen, irrespective of the number of mismatching RBC units transfused. The reason for this is that we consider the prevention of antigen mismatches more important than limiting the exposure to foreign antigens. Note that when mismatches can be prevented fully, there is no exposure at all. But as mentioned earlier, sometimes minor antigen mismatches are necessary to avoid shortages or outdating. In those cases, our method does not focus on limiting the exposure to foreign antigens for already mismatched patients, but instead minimizes the number of patients who are mismatched. The consequences of this are that we do not increase the mismatch penalty upon transfusion with multiple mismatching units as opposed to a single mismatching unit. However, we argue that preventing mismatches has more priority than limiting foreign antigen exposure, especially when it concerns patients who require extended matching in one form or another.

By explicitly modelling patients with demand for multiple units we have also added extra complexity to our model. Most importantly, the fact that we consider any patient who is denied RBC units an equal shortage, irrespective of the number of units demanded, makes it no longer trivial to compute an optimal assignment. However, we have also shown that when there are no shortages in the optimal solution, this problem is no longer prevalent. As in practice we expect essentially zero shortages, this means that the problem is mostly mitigated. In our simulations in Chapter 6 we have not included any knowledge about future demand. However, we have noted earlier that most RBC transfusions are elective, meaning that they are planned in advance. Including these requests will likely increase the minor antigen matching quality, but also increase the problem size. When one-week worth of demand is kept in inventory and all patient requests for the next three weeks are considered, it is unlikely that all this demand can be satisfied with the units in inventory. A naive implementation of our model in

which all known requests are included will therefore be harder to solve optimally as the solution is no longer free of shortages. To prevent this effect we have used an altered variant of our model in Chapter 8 in which the requests that must be satisfied on the current day receive an increased shortage penalty, such that they are always prioritized in computing an assignment. We further have not included demand for regular patients that was more than one day ahead in the future, to prevent an accumulation of requests and a worsened performance. This is a limitation of our model. Ideally, all information about the future and thereby all known patient requests should be included in the optimization to make the most informed decisions. Furthermore, when more future demand is considered ideally also future supply should be considered. This includes involving the distribution centre in the optimization. We have briefly experimented with this in Chapter 8. Simulations have shown that including the distribution centre in the optimization can effectively reduce the (already low) percentage of shortages for the different patient groups. We have consciously chosen to only allow special patient groups to be matched units directly from the distribution centre. The reason for this was twofold. Firstly, by limiting the number of requests for which we allow propagation, the optimization is faster as the solution space is reduced. Secondly, by denying regular patients matching on the distribution centre we intend to increase the availability of specific antigen negative RBC units for the patient groups who can benefit the most from extensively matched RBC unis.

Future work is needed to determine how all known future requests can be included in the daily optimization to maximize the availability of RBC units for all patient groups without substantially worsened performance. When this can be effectively implemented, it may become possible to perform long term extensive matching. Ideally, this planning could also reach as far as the donor population. Up until now we have modelled supply as a stochastic process. In a future where all donors have been typed on all clinically relevant antigens it becomes possible to invite specific donors to meet future demand for certain antigen negative RBC units. Although this extensive planning could certainly benefit large groups of patients, there will always be unforeseen demand. Until now, we assumed that demand for urgent requests is not within the scope of extended matching. In practice it may however still be useful to have a regional or national inventory filled with (possibly frozen) highly usable units to meet this kind of demand. Currently, Sanquin does have a small frozen inventory of such RBC units. However, when the entire donor population is extensively typed, there will also be an increase in the number donors with highly compatible blood groups. An interesting logistical problem is then when blood should be classified as highly usable and become eligible for storing in such regional or national emergency inventories.

Lastly, we have not accounted for unused units in our model. Unused units are units which are requested for transfusion but end up unused. In practice this is not uncommon, especially for patient who undergo a surgery and may need a transfusion. In this case, it is no longer guaranteed that assigning a mismatching unit will lead to foreign antigen exposure. Furthermore, the risk of outdating may be increased when units are repeatedly not transfused. Another issue is what should be done with highly compatible units which are not transfused after issuing. Returning such units to the distribution centre in order to increase their availability is likely better than keeping them in the hospital after they are initially assigned. Such an approach requires more planning and coordination within the blood supply chain as currently no units are returned from the hospitals to the distribution centres. We briefly alluded on an extension to the MINRAR formulation that can account for unused units in Section 6.8.4. However, we have not investigated how unused units affect the logistics of the blood supply chain. Before large-scale extensive matching is implemented, we certainly recommend further investigation on how to handle unused extensively matched RBC units.

11 Conclusion

In this thesis we have investigated how comprehensive antigen matching can be effectively implemented for RBC inventories of various sizes. We have built on previous work by Van Sambeeck *et al.* [6] and adapted their issuing strategy to perform extensive matching in smaller inventories without an increase in shortages or outdating. By first identifying which priorities govern future large-scale extensive matching, we concluded that allowing shortages in order to perform extensive matching makes a model unrepresentative of any future scenario. The prevention of shortages is therefore given absolute priority over all other attributes of our model. Furthermore, our model uses the full valid solution space by allowing all AB-RhD compatible matches. Next, we noted that although in theory phenotype-identical RBC issuing is the best possible issuing strategy, in practice this is often very restrictive or infeasible due to a limited inventory size and discrepancies between the donor and patient populations. Instead, our MINRAR-Online issuing strategy aims to minimize the risk of alloimmunization in the long run by weighing every match between RBC unit and patient by both the minor antigen mismatches as well as the minor antigen substitutions. Simulations show that this issuing strategy, where we do not limit ourselves to a specific issuing policy such as antigen-identical or antigen-compatible, is able to outperform the previously proposed issuing strategies. This is mainly achieved by using the relative immunogenicity of the antigens, a measure for the likelihood of antibody formation given a mismatch, as a penalty for mismatches or substitutions on the minor antigens. This allows our model to effectively compute trade-offs between different antigen mismatches when fully compatible matching on all fourteen antigens is impossible. Furthermore, the near optimal matching ability of our strategy is demonstrated by a comparison to the performance of an offline model, which we use to compute the best possible assignment of RBC units to patients over an entire year in the past. This comparison shows that our model is mostly only 1 percentage point of the best possible minor antigen matching quality, thereby showing the ability to perform near optimal matching for most of the minor antigens.

By allowing all minor antigen mismatches at a cost, we have created not only a very flexible model but also opened up the possibility for several extensions in which more complex cost functions can be used to capture more aspects of RBC matching. One of these possibilities is to use specific costs for antigen mismatches for various patient groups. This approach was investigated in Chapter 8. Our implementation of this extension relied on many assumptions about the clinical relevance of antigen mismatches for the patient groups considered. However, despite these assumptions our simulations demonstrated how such an approach can prioritize certain patient groups in preventing minor antigen mismatches effectively. Furthermore, we have shown that the availability of extensively matched RBC units can be increased by allowing difficult to match patients to be matched units from the distribution centre.

Current RBC matching is standardized and regulated using guidelines, which Dutch hospitals follow such that antigen matching practices are similar throughout the Netherlands. This guideline driven antigen matching provides clarity and support for RBC inventory managers, as well as the Dutch blood bank Sanquin. The extensiveness of these guidelines, together with the fact that extensive antigen matching is currently only feasible for a limited patient group, makes it that RBC inventory management in hospitals is largely still a manual process. Furthermore, most hospitals have the luxury of AB-RhD specific ordering as well as multiple delivery moments per day, allowing them to sequentially allocate patients with blood using a greedy FIFO-based strategy without the risk of shortages or outdating. Large-scale extensive matching based on the availability of genotyping chips is likely to make such a manual issuing strategy no longer practically feasible. This means that future RBC inventory management will become

largely computer based. Previous work by Van Sambeeck *et al.* [6] has already shown how an assignment of RBC units to patients can be computed given a set of antigens for which all patients must be compatibly matched. Similarly, work by Van Sambeeck *et al.* [7] showed in which order antigens should be ignored for matching when no compatible unit is available for a given patient. Although computer driven, both approaches are still guideline based. Either a specific antigen set is determined on which all patients should receive compatible matching, or a specific order is determined in which antigens should be ignored.

In this thesis we have taken a step back and concluded that large-scale extensive antigen matching should not lead to a decrease in shortages or outdating. We have constructed an RBC allocation model which is able to perform high quality minor antigen matching without the risk of inducing shortages. Furthermore, our model does not rely on a predetermined set of antigens to be compatibly matched and neither does it need an order in which antigens should be discarded for compatible matching. This model is shown to outperform multiple variants of the issuing strategy by Van Sambeeck *et al.* [6], mainly because of the removal of this predetermined set of antigens to be compatibly matched. Our model shows that a large-scale extensive matching policy works best when it is not guideline driven. Restrictions such as "no mismatches allowed on antigen K" or "Rhesus compatible matching for all patients" will only limit the matching possibilities of this model and thus likely lead to a decrease the overall matching quality or a possible increase in the number of shortages.

This shows that future large-scale extended matching does not only give rise to computational challenges but also implementations challenges. RBC inventory managers, which currently rely on extensive guidelines, will need to trust decision support matching tools for computer aided extensive matching to become viable. A first step to this paradigm shift is to create awareness within the blood transfusion chain of extended matching and its challenges and possibilities. After all, large-scale extensive antigen matching has the potential to improve the effectiveness and quality of RBC transfusions by preventing nearly all alloimmunization, but the support of the entire blood transfusion chain is needed to make this possibility into a future reality.

References

[1] J. R. Storry, F. B. Clausen, L. Castilho, Q. Chen, G. Daniels, G. Denomme, W. A. Flegel, C. Gassner, M. De Haas, C. Hyland, *et al.*, "International society of blood transfusion working party on red cell immunogenetics and blood group terminology: Report of the dubai, copenhagen and toronto meetings", *Vox sanguinis*, vol. 114, no. 1, pp. 95–102, 2019 (cit. on p. 8).

[2] CBO, *Cbo blood transfusion guideline*, 2011 (accessed November 24, 2020). [Online]. Available: http://www.isbtweb.org/fileadmin/user_upload/blood-transfusion-guideline.pdf (cit. on pp. 10, 15).

[3] D. Evers, R. A. Middelburg, M. De Haas, S. Zalpuri, K. M. de Vooght, D. van de Kerkhof, O. Visser, N. C. Péquériaux, F. Hudig, H. Schonewille, *et al.*, "Red-blood-cell alloimmunisation in relation to antigens' exposure and their immunogenicity: A cohort study", *The Lancet Haematology*, vol. 3, no. 6, e284–e292, 2016 (cit. on pp. 11–13, 64, 65, 81, 101).

[4] Sanquin, *Sanquin annual report*, 2018 (accessed November 26, 2020). [Online]. Available: https://www.sanquin.nl/binaries/content/assets/sanquinen/about-sanquin/annual-reports/sanquin_annual_report_2018.pdf (cit. on p. 14).

[5] J. Jongerius, M. Bouman, H. Bos, and H. de Wit, "Typeren van donorerytrocyten", *Tijdschrift voor Bloedtransfusie*, vol. 6, pp. 62–64, 2013 (cit. on p. 14).

[6] J. Van Sambeeck, S. Van Brummelen, N. Van Dijk, and M. Janssen, "Blood group specific issuing policies to improve inventory management of red blood cells", *Submitted for publication*, 2020 (cit. on pp. 17, 27, 32, 40, 42, 47, 99, 103, 104).

[7] J. Van Sambeeck, C. Van der Schoot, N. Van Dijk, H. Schonewille, and M. Janssen, "Mathematical optimization for alloimmunization prevention", *Submitted for publication*, 2020 (cit. on pp. 17, 65, 104).

[8] R. Haijema, J. van der Wal, and N. M. van Dijk, "Blood platelet production: Optimization by dynamic programming and simulation", *Computers & Operations Research*, vol. 34, no. 3, pp. 760–779, 2007 (cit. on pp. 20, 21).

[9] R. Haijema, N. van Dijk, J. van der Wal, and C. S. Sibinga, "Blood platelet production with breaks: Optimization by sdp and simulation", *International Journal of Production Economics*, vol. 121, no. 2, pp. 464–473, 2009 (cit. on pp. 20, 21).

[10] Q. Duan and T. W. Liao, "A new age-based replenishment policy for supply chain inventory optimization of highly perishable products", *International journal of production economics*, vol. 145, no. 2, pp. 658–671, 2013 (cit. on pp. 20–23).

[11] U. Abdulwahab and M. Wahab, "Approximate dynamic programming modeling for a typical blood platelet bank", *Computers & Industrial Engineering*, vol. 78, pp. 259–270, 2014 (cit. on pp. 20–22).

[12] I. Civelek, I. Karaesmen, and A. Scheller-Wolf, "Blood platelet inventory management with protection levels", *European Journal of Operational Research*, vol. 243, no. 3, pp. 826–838, 2015 (cit. on pp. 20, 21).

[13] S. Rajendran and A. R. Ravindran, "Platelet ordering policies at hospitals using stochastic integer programming model and heuristic approaches to reduce wastage", *Computers & Industrial Engineering*, vol. 110, pp. 151–164, 2017 (cit. on p. 20).

[14] Q. Duan and T. W. Liao, "Optimization of blood supply chain with shortened shelf lives and abo compatibility", *International Journal of Production Economics*, vol. 153, pp. 113–129, 2014 (cit. on pp. 20, 22, 23).

[15] M. Dillon, F. Oliveira, and B. Abbasi, "A two-stage stochastic programming model for inventory management in the blood supply chain", *International Journal of Production Economics*, vol. 187, pp. 27–41, 2017 (cit. on pp. 20, 22).

[16] K. G. Pauls-Worm, E. M. Hendrix, R. Haijema, and J. G. van der Vorst, "An milp approximation for ordering perishable products with non-stationary demand and service level constraints", *International Journal of Production Economics*, vol. 157, pp. 133–146, 2014 (cit. on pp. 20, 21).

[17] M. Najafi, A. Ahmadi, and H. Zolfagharinia, "Blood inventory management in hospitals: Considering supply and demand uncertainty and blood transshipment possibility", *Operations Research for Health Care*, vol. 15, pp. 43–56, 2017 (cit. on pp. 20, 22, 23).

[18] S. Gunpinar and G. Centeno, "Stochastic integer programming models for reducing wastages and shortages of blood products at hospitals", *Computers & Operations Research*, vol. 54, pp. 129–141, 2015 (cit. on pp. 20, 21).

[19] B. Hamdan and A. Diabat, "A two-stage multi-echelon stochastic blood supply chain problem", *Computers & Operations Research*, vol. 101, pp. 130–143, 2019 (cit. on p. 20).

[20] C. Sapountzis, "Allocating blood to hospitals", *Journal of the Operational Research Society*, vol. 40, no. 5, pp. 443–449, 1989 (cit. on pp. 20, 23).

[21] M. P. Atkinson, M. J. Fontaine, L. T. Goodnough, and L. M. Wein, "A novel allocation strategy for blood transfusions: Investigating the tradeoff between the age and availability of transfused blood", *Transfusion*, vol. 52, no. 1, pp. 108–117, 2012 (cit. on pp. 20, 21).

[22] A. Simonetti, R. A. Forshee, S. A. Anderson, and M. Walderhaug, "A stock-and-flow simulation model of the us blood supply", *Transfusion*, vol. 54, no. 3pt2, pp. 828–838, 2014 (cit. on p. 20).

[23] M. J. Fontaine, Y. T. Chung, F. Erhun, and L. T. Goodnough, "Age of blood as a limitation for transfusion: Potential impact on blood inventory and availability", *Transfusion*, vol. 50, no. 10, pp. 2233–2239, 2010 (cit. on pp. 20, 23).

[24] A. Grasas, A. Pereira, M.-A. Bosch, P. Ortiz, and L. Puig, "Feasibility of reducing the maximum shelf life of red blood cells stored in additive solution: A dynamic simulation study involving a large regional blood system", *Vox sanguinis*, vol. 108, no. 3, pp. 233–242, 2015 (cit. on pp. 20, 23).

[25] G. P. Prastacos, "Blood inventory management: An overview of theory and practice", *Management science*, vol. 30, no. 7, pp. 777–800, 1984 (cit. on p. 20).

[26] G. Zallen, P. J. Offner, E. E. Moore, J. Blackwell, D. J. Ciesla, J. Gabriel, C. Denny, and C. C. Silliman, "Age of transfused blood is an independent risk factor for postinjury multiple organ failure", *The American journal of surgery*, vol. 178, no. 6, pp. 570–572, 1999 (cit. on p. 20).

[27] P. J. Offner, E. E. Moore, W. L. Biffl, J. L. Johnson, and C. C. Silliman, "Increased rate of infection associated with transfusion of old blood after severe injury", *Archives of Surgery*, vol. 137, no. 6, pp. 711–717, 2002 (cit. on p. 20).

[28] F. R. Purdy, M. G. Tweeddale, and P. M. Merrick, "Association of mortality with age of blood transfused in septic icu patients", *Canadian Journal of Anaesthesia*, vol. 44, no. 12, pp. 1256–1261, 1997 (cit. on p. 20).

[29] J. T. Blake, "On the use of operational research for managing platelet inventory and ordering", *Transfusion*, vol. 49, no. 3, pp. 396–401, 2009 (cit. on p. 21).

[30] J. Charpin and A. Adewumi, "Optimal assignment of blood in a blood banking system", Tech. Rep., Mathematics in Industry Study Group (MISG), Tech. Rep., 2011 (cit. on p. 22).

[31] A. Adewumi, N. Budlender, and M. Olusanya, "Optimizing the assignment of blood in a blood banking system: Some initial results", in *2012 IEEE Congress on Evolutionary Computation*, IEEE, 2012, pp. 1–6 (cit. on p. 22).

[32] M. O. Olusanya, M. A. Arasomwan, and A. O. Adewumi, "Particle swarm optimization algorithm for optimizing assignment of blood in blood banking system", *Computational and mathematical methods in medicine*, vol. 2015, 2015 (cit. on p. 22).

[33] P. Govender and A. E. Ezugwu, "A symbiotic organisms search algorithm for optimal allocation of blood products", *IEEE Access*, vol. 7, pp. 2567–2588, 2018 (cit. on p. 22).

[34] K. Igwe, M. Olusanya, and A. Adewumi, "On the performance of grasp and dynamic programming for the blood assignment problem", in *2013 IEEE Global Humanitarian Technology Conference (GHTC)*, IEEE, 2013, pp. 221–225 (cit. on p. 22).

[35] C. G. Koch, L. Li, D. I. Sessler, P. Figueroa, G. A. Hoeltge, T. Mihaljevic, and E. H. Blackstone, "Duration of red-cell storage and complications after cardiac surgery", *New England Journal of Medicine*, vol. 358, no. 12, pp. 1229–1239, 2008 (cit. on p. 23).

[36] C. Sapountzis, "Allocating blood to hospitals from a central blood bank", *European Journal of Operational Research*, vol. 16, no. 2, pp. 157–162, 1984 (cit. on p. 23).

[37] K. E. Kendall and S. M. Lee, "Formulating blood rotation policies with multiple objectives", *Management Science*, vol. 26, no. 11, pp. 1145–1157, 1980 (cit. on p. 23).

[38] R. M. Fasano and S. T. Chou, "Red blood cell antigen genotyping for sickle cell disease, thalassemia, and other transfusion complications", *Transfusion medicine reviews*, vol. 30, no. 4, pp. 197–201, 2016 (cit. on p. 24).

[39] S. T. Chou, T. Jackson, S. Vege, K. Smith-Whitley, D. F. Friedman, and C. M. Westhoff, "High prevalence of red blood cell alloimmunization in sickle cell disease despite transfusion from rh-matched minority donors", *Blood, The Journal of the American Society of Hematology*, vol. 122, no. 6, pp. 1062–1071, 2013 (cit. on p. 24).

[40] C. Paccapelo, F. Truglio, M. Antonietta Villa, N. Revelli, and M. Marconi, "Hea beadchip technology in immunohematology", *Immunohematology*, vol. 31, no. 2, pp. 81–90, 2015 (cit. on p. 24).

[41] C. Bonet Bub and L. Castilho, "Id core xt as a tool for molecular red blood cell typing", *Expert review of molecular diagnostics*, vol. 19, no. 9, pp. 777–783, 2019 (cit. on p. 24).

[42] S. Sapatnekar and P. I. Figueroa, "How do we use molecular red blood cell antigen typing to supplement pretransfusion testing?", *Transfusion*, vol. 54, no. 6, pp. 1452–1458, 2014 (cit. on p. 24).

[43] G. Meny, "The duffy blood group system: A review", *Immunohematology*, vol. 26, no. 2, pp. 51–56, 2010 (cit. on pp. 24, 87).

[44] C. Tournamille, N. Meunier-Costes, B. Costes, J. Martret, A. Barrault, P. Gauthier, F. Galactéros, R. Nzouékou, P. Bierling, and F. Noizat-Pirenne, "Partial c antigen in sickle cell disease patients: Clinical relevance and prevention of alloimmunization", *Transfusion*, vol. 50, no. 1, pp. 13–19, 2010 (cit. on p. 25).

[45] K. Wilkinson, S. Harris, P. Gaur, A. Haile, R. Armour, G. Teramura, and M. Delaney, "Molecular blood typing augments serologic testing and allows for enhanced matching of red blood cells for transfusion in patients with sickle cell disease", *Transfusion*, vol. 52, no. 2, pp. 381–388, 2012 (cit. on p. 25).

[46] L. Castilho and C. L. Dinardo, "Optimized antigen-matched in sickle cell disease patients: Chances and challenges in molecular times-the brazilian way", *Transfusion Medicine and Hemotherapy*, vol. 45, no. 4, pp. 258–262, 2018 (cit. on p. 25).

[47] J. Perreault, J. Lavoie, P. Painchaud, M. Côté, J. Constanzo-Yanez, R. Côté, G. Delage, F. Gendron, S. Dubuc, B. Caron, *et al.*, "Set-up and routine use of a database of 10 555 genotyped blood donors to facilitate the screening of compatible blood components for alloimmunized patients", *Vox sanguinis*, vol. 97, no. 1, pp. 61–68, 2009 (cit. on p. 25).

[48] W. A. Flegel, J. L. Gottschall, and G. A. Denomme, "Integration of red cell genotyping into the blood supply chain: A population-based study", *The Lancet Haematology*, vol. 2, no. 7, e282–e288, 2015 (cit. on p. 25).

[49] M. R. Garey and D. S. Johnson, *Computers and intractability*. freeman San Francisco, 1979, vol. 174 (cit. on p. 30).

[50] R. M. Karp, "Reducibility among combinatorial problems", in *Complexity of computer computations*, Springer, 1972, pp. 85–103 (cit. on pp. 30, 36).

[51] D. Bertsimas and J. N. Tsitsiklis, *Introduction to linear optimization*. Athena Scientific Belmont, MA, 1997, vol. 6 (cit. on p. 31).

[52] K. R. Rebman, "Total unimodularity and the transportation problem: A generalization", *Linear Algebra and its Applications*, vol. 8, no. 1, pp. 11–24, 1974, ISSN: 0024-3795 (cit. on p. 31).

[53] H. W. Lenstra Jr, "Integer programming with a fixed number of variables", *Mathematics of operations research*, vol. 8, no. 4, pp. 538–548, 1983 (cit. on p. 38).

[54] I. J.-B. F. Adan, M. J. van Eenige, and J. A. C. Resing, *Fitting discrete distributions on the first two moments*. Eindhoven University of Technology, Department of Mathematics and Computing ..., 1994 (cit. on pp. 50, 111).

Appendices

A Antigen Frequencies

Blood group system	Phenotype	Relative frequency among Caucasians	Relative frequence among individuals of African descent
ABO	{}	0.43	0.27
	{A}	0.44	0.20
	{B}	0.09	0.49
	{A,B}	0.04	0.04
Rhesus	{C,c,E,e}	Rare	Rare
	{C,E}	Rare	Rare
	{C,e}	Rare	Rare
	{c,E}	Rare	Rare
	{c,e}	0.151	0.068
	{C,c,E}	Rare	Rare
	{C,c,e}	0.008	Rare
	{C,E,e}	Rare	Rare
	{c,E,e}	0.009	Rare
	{D,C,c,E,e}	0.133	0.056
	{D,C,E}	Rare	Rare
	{D,C,e}	0.185	0.20
	{D,c,E}	0.023	0.02
	{D,c,e}	0.021	0.458
	{D,C,c,E}	0.001	Rare
	{D,C,c,e}	0.349	0.210
	{D,C,E,e}	0.002	Rare
	{D,c,E,e}	0.118	0.186
Kell	{}	Rare	Rare
	{k}	0.91	0.98
	{K}	0.002	Rare
	{K,k}	0.088	0.02
Duffy	{}	Rare	0.68
	{Fy(a)}	0.17	0.09
	{Fy(b)}	0.34	0.22
	{Fy(a), Fy(b)}	0.49	0.01

continued on next page

continued from previous page

Blood group system	Phenotype	Relative frequency among Caucasians	Relative frequence among individuals of African descent
Kidd	{}	Rare	Rare
	{Jk(a)}	0.263	0.511
	{Jk(b)}	0.234	0.081
	{Jk(a), Jk(b)}	0.503	0.488
MNS	{M,S}	0.06	0.02
	{M,S,s}	0.14	0.07
	{M,s}	0.08	0.16
	{M,N,S}	0.04	0.02
	{M,N,S,s}	0.24	0.13
	{M,N,s}	0.22	0.325
	{N,S}	0.01	0.02
	{N,S,s}	0.06	0.05
	{N,s}	0.15	0.19
	{M}	Rare	0.004
	{M,N}	Rare	0.004
	{N}	Rare	0.007

Table 33: Relative frequencies of antigen combinations for the ABO, Rhesus, Kell, Duffy, Kidd and MNS blood group systems among Caucasians and individuals of African descent. Some combinations are not included, as they do not exist.

B Demand Distributions

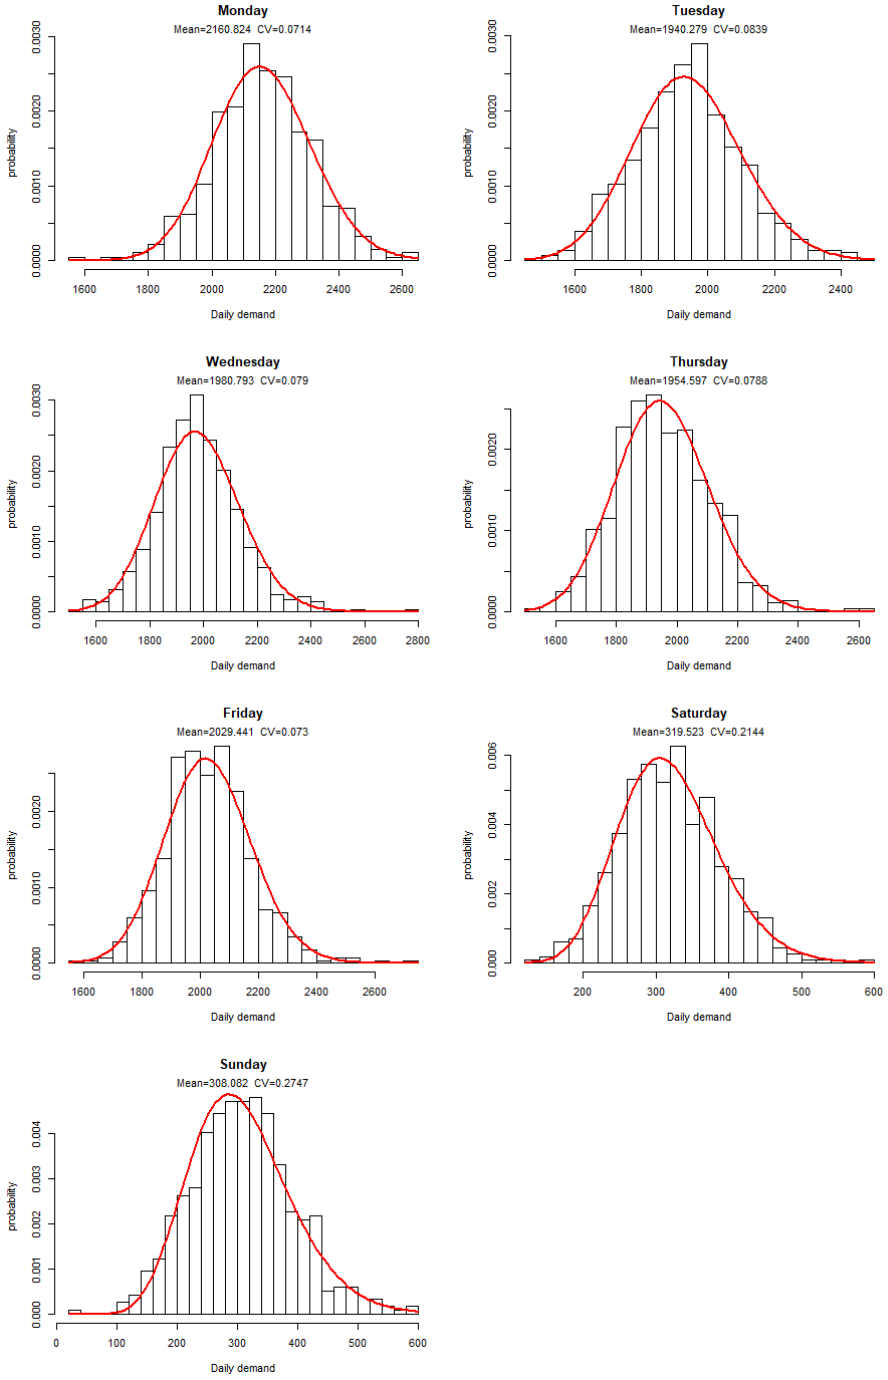

Figure 23: Empirical distributions for every day of the week for national demand. The red curve is the fitted mixture of two negative binomial distributions which preserves the mean and variance of the empirical distribution according to Adan *et al.* [54]. For each distribution the mean and coefficient of variation are shown.