

OPTIMIZATION OF EXTENDED RED BLOOD CELL MATCHING IN TRANSFUSION DEPENDENT SICKLE CELL PATIENTS

Folarin B. Oyebolu¹, Marie Chion¹, Merel L. Wemelsfelder²,
Sara Trompeter^{3,4}, Nicholas Gleadall^{4,5}, and William J. Astle^{1,4}

¹MRC Biostatistics Unit, University of Cambridge, Cambridge, UK

²Donor Medicine Research Dept., Sanquin Research, Amsterdam, NETHERLANDS

³Dept. of Haematology, University College Hospital NHS Foundation Trust, London, UK

⁴NHS Blood and Transplant, UK

⁵Dept. of Haematology, University of Cambridge, Cambridge, UK

ABSTRACT

Blood transfusion is a life-saving treatment for people with sickle cell disorder (SCD). Presently, blood is matched manually for transfusion using incomplete red cell blood type information, to minimize the immunological incompatibility between donor and patient. We are investigating alternative approaches to blood allocation that exploit extended blood type information measured by a new genetic test shortly to be introduced by the National Health Service in England. We formulate sequential allocation decisions as a Markov decision process and study penalty-based policies for matching, which consider up to 17 blood group antigens, including policies that look ahead to future patient appointments. We tune the policy parameters of a matching rule to minimize formation of antibodies (alloimmunization) in SCD patients and estimate that a 98% reduction in alloimmunization can be achieved compared to current policies. Finally, we show that the tuned policy parameters are robust to major supply shocks.

1 INTRODUCTION

NHS Blood & Transplant (NHSBT) issues 10,000 blood units monthly for sickle cell disorder (SCD) patients in England. SCD is a severe inherited disorder affecting the haemoglobin protein in red blood cells (RBCs). Symptoms of SCD include acute painful vaso-occlusive crises, chronic organ damage and a high risk of stroke. Transfusion of RBCs is used to treat the acute symptoms and, prophylactically, to reduce the risk of crisis and stroke. Many people with SCD in England receive regular transfusions, mostly exchange transfusions, which replace the RBCs of the patient with RBCs from multiple donors.

RBC transfusion is complicated by genetic variation which creates blood types — differences in the antigens expressed on the surfaces of RBCs between individuals. An individual who has an antigen on their RBCs is *positive* for that antigen, otherwise they are *negative* for that antigen. If a person receives a foreign antigen in a transfusion (i.e., the donated unit is positive for an antigen for which the recipient is negative) there is a mismatch: the donated unit is incompatible with the recipient. A mismatch may lead to the formation of an alloantibody to the foreign antigen (this is called alloimmunization) which, upon a further mismatch, may trigger the immune system to reject the transfused blood. There are 362 known antigens, so blood that matches all the antigens — or even all the clinically relevant antigens — for which a patient is negative is rarely available. Current UK guidelines require that SCD patients receive RBCs matched for the ‘major’ antigens A, B and D and for the ‘minor’ antigens C, c, E, e and K, which pose a high risk of alloimmunization in patients receiving regular mismatched transfusions.

For various reasons it can be difficult to identify well matched blood for SCD patients in England. Firstly, many patients have pre-existing antibodies. When a patient forms an alloantibody, the pool of compatible units shrinks. Secondly, over 90% of patients but only 1% of NHSBT blood donors have recent

African ancestry, and there are significant differences in the frequencies of RBC antigens between African and European populations. Thirdly, because traditional antigen testing (phenotyping) is expensive, only a few of the antigens of most donors are known. Finally, the risk of alloimmunization in SCD is greater than in other conditions treated by transfusion because an exchange transfusion can expose a patient to up to 10 units of blood and many patients require regular transfusion for life.

Recently, a cheap genotyping test has been developed that can determine almost all the known RBC antigens simultaneously (Gleadall et al. 2020). Mass genetic blood typing of donors and patients can potentially streamline the procurement of RBCs for SCD patients and enable anticipatory matching to reduce alloimmunization. Our contributions in this paper include extending recently developed frameworks for blood matching in the context of mass genotyping (van Sambeek et al. 2022; van de Weem et al. 2022; Wemelsfelder et al. 2024) to examine the implications for the provision of blood for SCD patients in England. We investigate how the parameters of matching policies affect the composition of the national RBC stock and the age of the RBCs allocated to SCD patients. We then optimize those parameters to determine any trade-offs between minimizing alloimmunizations, shortages, and expiries at different inventory sizes.

2 LITERATURE REVIEW

Many aspects of blood supply chains (BSCs) have been studied previously by simulation, including donation and collection (Williams et al. 2020; McElfresh et al. 2023), inventory volatility (Clay et al. 2018), and supply and demand modeling (Ejohwomu et al. 2021). In a comprehensive review on BSC management, Beliën and Forcé (2012) classified publications by a number of features, noting the widespread use of simulation. A more recent survey pointed to the promise of approaches combining simulation and optimization (Pirabán et al. 2019).

Onggo (2014) presented a hybrid simulation model comprising three interacting modules: the donor population, simulated using system dynamics, and the blood center and the hospital, each simulated using discrete-event simulation (DES). Arani et al. (2020) use scenario analysis and a DES model to compare the current practice of a BSC with alternative operational decisions. Katsaliaki and Brailsford (2007) created a DES model of a vertical section of a blood supply chain from donation to transfusion, capturing blood type substitutions, and investigated several policies to improve system performance.

Osorio et al. (2018) heuristically solved a two-stage stochastic optimization problem for designing a BSC network and compared network configurations with different degrees of centralization. In a hybrid approach, Osorio et al. (2017) modeled a blood center issuing multiple products, analyzing shortages, expiries, the number of donors required and costs. They incorporated stochasticity of donations and requests by simulation and determined the number of donors of each blood type required and the products to be manufactured by optimization. More recently, Silva Magalhães et al. (2023) used simulation-optimization to determine a minimum size for an inventory of RBC units and replenishment points for each blood type.

Inventory management in a complex and dynamic BSC system often requires repeated sequential decision-making in the presence of uncertainty and may be modeled as a Markov decision process (MDP) (Civelek et al. 2015). For example, Soares et al. (2020) used an MDP formulation to find an optimal policy for collecting additional blood bags using external collection teams. Abdulwahab and Wahab (2014) developed a stochastic multi-period model for a platelet blood bank and analyzed several inventory control policies. Notably, they accounted for and penalized substitutions among the major blood groups. However, like most previous work, they only included the major antigens.

Recently, van Sambeek et al. (2022) proposed an MDP formulation and a mathematical framework which can be applied when an arbitrary set of antigens are measured on all patients and all units in inventory. This was the first study to incorporate clinically relevant minor RBC antigens. They model stochasticity due to donations and requests by simulation and optimize the allocation of units by solving a minimum-cost flow problem. Using historical donation and request data from the Netherlands they showed that more than 90% of requests can be fulfilled with RBC units that are a perfect match for 14 clinically relevant antigens. They note that their approach can be categorized as a simulation with optimization-based

iterations approach (Figueira and Almada-Lobo 2014). Subsequently, van de Weem et al. (2022) developed MINRAR, a simulation-optimization method which relaxes the requirement for strict matching of minor antigens, instead discouraging mismatches using penalties. Wemelsfelder et al. (2024) have generalized MINRAR to capture the heterogeneity in the clinical consequences of a mismatch across patient groups.

3 PROBLEM DESCRIPTION

In this paper, we tackle a stochastic sequential decision-making problem in blood inventory allocation. We address an MDP with daily decision epochs, in which each decision is to optimally allocate perishable inventory of multiple partially-substitutable products to patients with individually specific requests. Each decision can be formulated as a Hitchcock transportation problem.

We consider a finite sequence of daily RBC allocations $X_{t=1}, \dots, X_{t=T}$ over a time horizon T . The allocation decision X_t , made at time t , depends solely on the state S_t of the system at time t . S_t is an ordered pair (I_t, P_t) , where I_t is the state of the RBC inventory and P_t is the set of patient requests existing at time t . P_t contains all the requests that need to be met at time t and may contain, for every $t' > t$, some of those requests that need to be met at time t' . I_t and P_t are stochastic because each day newly donated RBC units are added to the inventory and new requests for patients are made. The blood types of new donations and new patients are randomly sampled according to the joint distributions of antigen frequencies in the donor and patient populations respectively. I_t also depends on I_{t-1} and P_{t-1} because any inventory that is not allocated to a request that must be met at time $t - 1$ and does not expire at time $t - 1$ is carried over to time t . A negative reward accumulates after the decision in each epoch, quantifying the clinical costs of the blood allocation.

3.1 Daily Matching Optimization Model

We now describe a procedure for making the allocation decisions. We consider the problem for a single epoch, in places dropping the dependence of notation on t for convenience. Let m denote $|I_t|$, the size of the inventory and let n denote $|P_t|$ the number of existing patient requests. Let g_i be the age in days of RBC unit $i \in I_t$, let $a_i = 1$ (see below), and let h_j be the maximum age at which a unit can be allocated to request $j \in P_t$. h_j will depend on the reason for the patient's treatment. Let b_j be the number of RBC units required to satisfy request $j \in P_t$. For each $(i, j) \in I_t \times P_t$, let c_{ij} denote the cost associated with allocating unit i to request j and let $d_{ij} \in \{0, 1\}$ indicate that unit i is compatible with the patient of request j . $d_{ij} = 0$ if unit i does not or will not exist in inventory at the fulfillment date of request j ; or if the patient corresponding to request j has an antibody against an antigen carried by unit i . Also, $d_{ij} = 0$ to enforce transfusion guidelines for a particular patient group. We wish to

$$\text{minimize } \sum_i^m \sum_j^n c_{ij} x_{ij}, \tag{1}$$

$$\text{subject to: } \sum_j^n x_{ij} \leq a_i \quad \forall i \in \{1, \dots, m\}, \tag{2}$$

$$\sum_i^m x_{ij} = b_j \quad \forall j \in \{1, \dots, n\}, \tag{3}$$

$$x_{ij} \in \{0, 1\} \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}, \tag{4}$$

$$x_{ij} \leq d_{ij} \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}, \tag{5}$$

$$x_{ij} g_i \leq h_j \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}. \tag{6}$$

Equations (1)–(4) are a variation of the classical Hitchcock transportation problem in which each RBC unit i in the inventory is a supply node able to offer a_i units of product and each patient request j is a demand node seeking b_j units of product. In the conventional formulation (Gass 1990), the constraint in

Equation (2) is an equality, but we relax this to an inequality since supply is greater than demand. The integer decision variable x_{ij} is 1 if unit i is assigned to request j and 0 otherwise. The model is infeasible in the case of shortages. This can be addressed in implementation by introducing dummy units representing shortages with large allocation costs, or by reducing the problem to one of minimum-cost maximum flow following van Sambeek et al. (2022). The solution to the problem determines X_t of the MDP, an allocation from I_t to satisfy those requests in P_t that need to be met on day t . Units assigned by the solution to requests that need to be met on days subsequent to t remain in the inventory to be reallocated in the next epoch.

3.2 Policy and Penalty Functions

Given S_t , the system state at time t , and a cost-matrix $c \in \mathbb{R}^{m \times n}$ the solution to the optimization problem described in Section 3.1 determines the allocation X_t . It therefore remains to specify a policy function $\pi : S_t \mapsto c$, equivalently a method for computing c_{ij} for each $(i, j) \in I_t \times P_t$ given S_t . We adopt a cost function similar to that used in the MINRAR model of van de Weem et al. (2022) — a linear combination of five penalty components in which each component π_f is given weight λ_f :

$$\pi(S_t)(i, j) = \sum_{f=1}^5 \lambda_f \pi_f(i, j). \quad (7)$$

The **immunogenicity** w_k of an antigen k is the probability that a recipient negative for k alloimmunizes against it when exposed by transfusion. π_1 penalizes for this risk. Let \mathbb{I}_j be the antigens for which penalization is required for recipient j . For each $k \in \mathbb{I}_j$, let $\phi_k(i, j) \in \{0, 1\}$ indicate a mismatch — i.e., that RBC unit i is positive for k and recipient j is negative for k — then,

$$\pi_1(i, j) = \sum_{k \in \mathbb{I}_j} w_k \phi_k(i, j) \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}.$$

Note that π_1 contributes to the objective linearly in the number of units mismatched for any given antigen, whereas the MINRAR model just penalizes the first mismatched unit in a request.

A **substitution** occurs for an antigen if a positive patient receives negative blood. Following MINRAR, we penalize substitutions differently for the major and minor antigens. The usability of a complete major blood group type is the fraction of the patient population compatible with transfusion of blood of that type. The penalty for the major antigens is the difference between the usability v of the major blood group types of i and j :

$$\pi_2(i, j) = v(i) - v(j) \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}.$$

Let \mathbb{S} be the set of minor antigens for which penalization is required and let $\zeta_k(i, j) \in \{0, 1\}$ indicate that j but not i carries antigen k . The penalty is the sum of the immunogenicities of the substituted antigens:

$$\pi_3(i, j) = \sum_{k \in \mathbb{S}} w_k \zeta_k(i, j) \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}.$$

An **expiry** occurs when a unit reaches the shelf-life G without being allocated. Following MINRAR we adopt a First-In First-Out (FIFO) penalty to control expiries. However we only apply it to requests for non-SCD patients. Let $\rho(j) \in \{0, 1\}$ indicate that j is a request for an SCD patient, then:

$$\pi_4(i, j) = -2^{-(G-g_i)/5} (1 - \rho(j)) \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}. \quad (8)$$

To penalize requests for SCD patients, we propose a new young blood (YB) penalty not in the MINRAR model. This was designed to meet an NHSBT guideline that SCD patients should be given blood not more than 14 days old and blood not more than 7 days old except in exceptional circumstances:

$$\pi_5(i, j) = (-g_i/\alpha_1 + \exp(g_i - \alpha_2))\alpha_3\rho(j) \quad \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}.$$

We set $\alpha_1 = 7.69, \alpha_2 = 9.58, \alpha_3 = 1.20$ and set h_j , which determines the constraint (6) on the maximum unit age that can be allocated to patient j , to 14 if $\rho(j) = 1$ and to G if $\rho(j) = 0$.

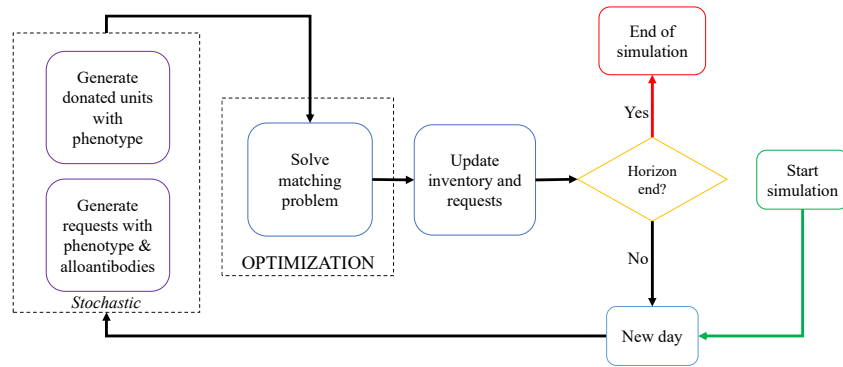


Figure 1: Flow diagram of the simulation procedure.

4 SIMULATION APPROACH

4.1 Conceptual Model

In this section we describe a method to simulate the MDP under various rules for blood allocation. The flow diagram in Figure 1 gives an overview of the procedure. Each iteration simulates an epoch corresponding to one day. In the stochastic part of each epoch we generate extensively typed RBC units from a mixed ancestry pool of donors. We also generate requests from extensively typed SCD patients with appointments for elective exchange transfusions and requests for patients without SCD for whom only the major blood types are known. In the optimization part of each epoch we allocate units from inventory to patients according to a given allocation policy, using the method described in Section 3.1. As each iteration completes we update statistics summarizing accumulated costs, including expected alloimmunization (x_A), the number of shortages, and the number of expiries of RBC units. We use the immunogenicity estimates of Evers et al. (2016) for the w_k and to calculate x_A .

We make various assumptions that simplify the NHSBT BSC. For example, we disregard transport costs, assuming that donations and transfusions all occur at a single stock holding location. We assume that any unit allocated for transfusion is used and cannot return to the inventory. We also assume that if, on any day, the supply of compatible units fails to meet demand patient requests will go unfulfilled, either partially or completely. Unmet demand does not pass to subsequent days — i.e., there is no backlogging.

4.2 Input Parameters

Table 1 summarizes the input parameters of the simulation. We chose a time horizon of six weeks: this is the approximate time between exchange transfusions for SCD patients and we do not model patient returns.

We consider only those antigens in $\mathbb{A}_{\text{MAJ}} \cup \mathbb{A}_{\text{MIN}}$, where $\mathbb{A}_{\text{MAJ}} = \{A, B, D\}$ is the set of major antigens and $\mathbb{A}_{\text{MIN}} = \{C, c, E, e, K, k, Fy^a, Fy^b, Jk^a, Jk^b, M, N, S, s\}$ is a subset of the minor antigens. We assume that all the SCD patients are of African ancestry, and that non-SCD patients and NHSBT donors come from a mixed population of which 1% is African ancestry and 99% is European ancestry. We assume the blood types determined by antigens in $\mathbb{A}_{\text{MAJ}} \cup \mathbb{A}_{\text{MIN}}$ are known for SCD patients and donors but only those determined by antigens in \mathbb{A}_{MAJ} are known for non-SCD patients. For donors, we assign blood types determined by the A, B, D, C, c, E, and e antigens by sampling units from historical NHSBT donation data. The blood types determined by the remaining antigens for donors and blood types determined by all antigens for patients are assigned independently according to their empirical population frequencies (Reid et al. 2012). We assign alloantibodies to SCD patients by sampling from historical data on patients at University College London Hospital (UCLH). We assume non-SCD patients have no alloantibodies and that they are positive for all minor antigens $k \in \mathbb{A}_{\text{MIN}}$ so that they never accrue immunogenicity penalties (i.e., $\varphi_k(i, j) = 0$ if $\rho(j) = 0$).

Table 1: Deterministic simulation input parameters values provided by NHSBT to reflect their operations.

Parameter name	Value	Unit
RBC initial inventory size	30000	
Daily RBC donations	3500	
RBC shelf life (G)	35	days
Daily SCD transfusion appointments	33	per day
SCD RBC demand	10	units per patient
RBC antigens considered	17	
Non-SCD demand	3170	units per day
RBC age-limit for SCD patients	14	days
Simulation time horizon (T)	6	weeks

4.3 Matching Rules

We set d_{ij} (in Equation (5)) to enforce compatible matching on \mathbb{A}_{MAJ} for all patients and to enforce compatible matching of SCD patients on $\{A, B, D, C, c, E, e, K\}$. We then investigated three matching rules with different policy function configurations for SCD patients (i.e., for all j where $\rho(j) = 1$).

- **Limited Rule (L0)** reflects current matching guidelines for SCD patients, in which just the antigens $\{A, B, D, C, c, E, e, K\}$ must be matched compatibly: $\mathbb{I}_j = \emptyset$; $\mathbb{S} = \{C, c, E, e, K\}$.
- **Extended Rule (E0)** exploits genetically measured blood type information by penalizing for mismatches of SCD patients at extended antigens and for substitutions at extended antigens: $\mathbb{I}_j = \{k, Fy^a, Fy^b, Jk^a, Jk^b, M, N, S, s\}$; $\mathbb{S} = \mathbb{A}_{MIN}$.
- **Extended Rule with Forecasting (E1)** is like **E0** but anticipates five future days of SCD patient appointments and one future day of incoming stock when optimizing the match. We assume all patients keep their appointments but, based on NHSBT attrition data, that only 97% of the stock anticipated actually arrives.

4.4 Software Implementation and Hardware

The model and procedure were implemented in Python (version 3.8), using Google OR-Tools (version 9.3) to solve the transportation problem described in Section 3.1. The simulations were executed using 20 cores of an Intel® Xeon® Platinum 8368Q node with 2.60GHz CPUs and 250GiB of RAM. We verified that the simulation output and the distribution of system states were consistent with those of the independently implemented MINRAR model (see Wemelsfelder et al. (2024) for details).

5 RESULTS AND DISCUSSION

5.1 Comparing the Limited and Extended Rules

Initially, we ran 100 simulation replications of each rule, with $\lambda_f = 1$ for each f . Each replication had an 18 week warm-up period in order for inventory to reach a steady-state. We compared the rules by computing the mean x_A (over the replications) during the six week time horizon for each antigen (Figure 2). The mean x_A for each of the C, c, E, e, and K antigens is zero because incompatible transfusions of those antigens are prohibited. The mean x_A for the k and s antigens are also zero as we assumed that the immunogenicity of these antigens is zero. The mean x_A for the other antigens is substantially lower in **E0** than in **L0**. Most of the absolute reduction in x_A is explained by the antigens Fy^a ($\Delta(x_A) = 2.79$), Jk^a ($\Delta(x_A) = 1.26$), and M ($\Delta(x_A) = 1.70$). Alloimmunization has been virtually eliminated in **E0** for the Jk^a and M antigens.

Table 2 summarizes the distributions of x_A aggregated across antigens, shortages, and RBC expiries. We observed an 81% reduction in mean aggregated x_A from **L0** to **E0**, and an 86% reduction in mean aggregated x_A from **L0** to **E1**. We observed no shortages affecting SCD patients or non-SCD patients under

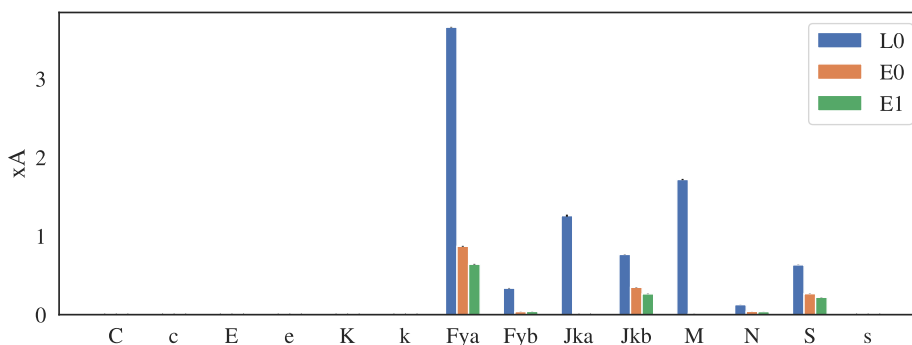


Figure 2: Mean alloimmunization rate in SCD patients (x_A) stratified by antigen (error bars: \pm std. error).

Table 2: Statistical summaries of the distributions (over simulation replications) of metrics for the three matching rules: mean \pm std. error of x_A , shortages to SCD patients, and RBC expiries; and mean \pm std. deviation of simulation runtime.

	L0	E0	E1
x_A	8.50 ± 0.016	1.58 ± 0.007	1.21 ± 0.005
SCD shortages	0	0	0
Expiries	0	0	0
Runtime (s)	111 ± 10	176 ± 18	736 ± 4

Table 3: Statistical summaries (mean \pm std. error) of the distributions of x_A , shortages to SCD patients, and expiries when the **E0** rule is applied and from sensitivity analysis with simplified penalty functions.

	E0	No IMM	No MAJ	No MIN	No SUB	No FIFO	Zero YB	No YB
x_A	1.58 ± 0.007	5.67 ± 0.011	0.50 ± 0.004	0.96 ± 0.005	0.29 ± 0.003	3.25 ± 0.007	1.03 ± 0.005	0.24 ± 0.002
SCD shortages	0	0	0	0	0	0.05 ± 0.04	0	0
Expiries	0	0	0	0	0	70.5 ± 2.4	0	0

any of the rules. Nor did we observe any expiration of RBC units. This suggests that the matching of RBC units to SCD patients using genetically measured extended blood types is unlikely to place a serious strain on the RBC inventory under the present NHSBT donor policy (assuming unrestricted RBC transportation).

5.2 Components of the Penalty Function

Next, we investigated the sensitivity of the steady state of the MDP to modification of components of the overall penalty function (Equation (7)). As the mean runtime of **E1** was quadruple that of **E0**, and the two rules are similar in structure (bar the forecasting element), we chose to base our analysis in this (and subsequent) section(s) on **E0**. We removed selected subsets of the component penalties (i.e., set the corresponding $\lambda_f = 0$) and ran 100 replications under **E0** in each case. Table 3 summarizes the clinical and operational statistics of each scenario.

5.2.1 Immunogenicity Penalty (No IMM, $\lambda_1 = 0$)

As might be expected, mean x_A is greater (by over 250%) than it was with the original penalty. However, no shortages or expiries were recorded. This scenario is different to **L0**, because the minor antigen substitution penalty (π_3) is still active. This results in lower mean x_A : by penalizing the allocation of antigen-negative units to antigen-positive patients, more antigen-negative supply is available for antigen-negative patients.

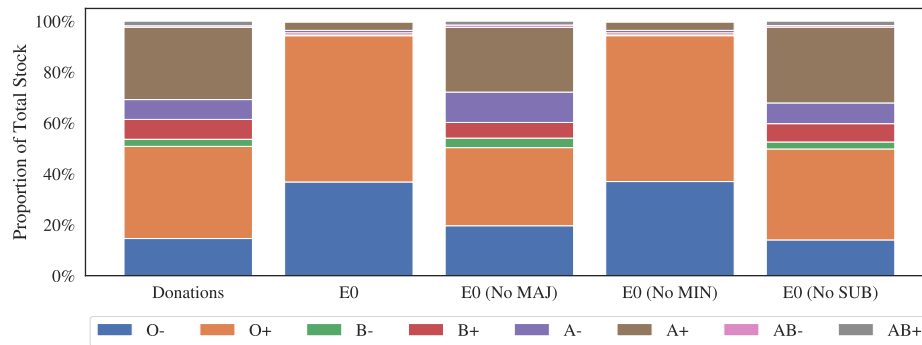


Figure 3: Distribution of the major blood group types of donated units compared with the mean distribution of the stock at steady state under **E0** and under **E0** with modifications of the substitution penalty. In every case, the std. error of the mean proportion corresponding to a blood type is less than 0.01%.

5.2.2 Substitution Penalties

Major substitution penalty removed (No MAJ, $\lambda_2 = 0$) When the major substitution penalty is removed the mean x_A is reduced by 68% from its value with the original penalty under **E0**.

Minor substitution penalty removed (No MIN, $\lambda_3 = 0$) When the minor substitution penalty is removed the mean x_A is reduced by 39% from its value with the original penalty under **E0**.

Both major and minor substitution penalties removed (No SUB, $\lambda_2 = 0$ and $\lambda_3 = 0$) We observe a reduction in mean x_A of 81% compared to the original penalty under **E0**. When compared to the No MIN and No MAJ scenarios there is a reduction in mean x_A of 70% and 42% respectively. Therefore, the approximate relative reduction in mean x_A caused by the removal of each of the two substitution penalties is independent of the order of their removal. Each penalty is responsible for non-trivial alloimmunization, contrary to the expectation that they would reduce alloimmunization by preserving antigen negative units in stock. The major substitution penalty has the larger effect.

Figure 3 illustrates the mean steady state (i.e., over replications and epochs) distribution of the major blood groups of the RBC units in stock during the 6 week simulation. It shows that the major substitution penalty causes substantial ‘hoarding’ of the valuable O- and O+ blood types. These blood types account for 95% of stock under **E0** when the major antigen substitution penalty is active. In contrast, they account for just half of stock when it is removed. In theory, an inventory dominated by type O units should be robust to demand or supply shocks. However, this comes at the cost of a higher rate of alloimmunization in SCD patients by reducing the pool of units available that are ‘good’ matches for patients on the minor antigens if those units are substitutions for a major antigen. Furthermore, the major antigen substitution penalty increases the mean age of units in stock (Figures 4a and 4b) and of RBC units allocated to patients.

5.2.3 FIFO Penalty (No FIFO, $\lambda_4 = 0$)

Without the FIFO penalty, mean x_A , shortages, and expiries increase. These changes are all likely due to the steady-state inventory shrinking from 30,000 to 5,000 units. Mean x_A more than doubles to 3.25, the mean number of expiries rises from zero to 70.5 units, and the mean number of shortages for SCD patients rises from zero to 0.05 units. The mean age of RBC units in stock is reduced to 2.0 days (4.3 days excluding new units).

5.2.4 Young Blood Penalty

YB penalty removed (Zero YB, $\lambda_5 = 0$) The mean x_A is reduced by 35% when the YB penalty is removed. There are still no shortages or expiries. The YB penalty ensures, as intended, that young blood is allocated to all SCD patients. When the penalty is complete the units allocated to SCD patients have a

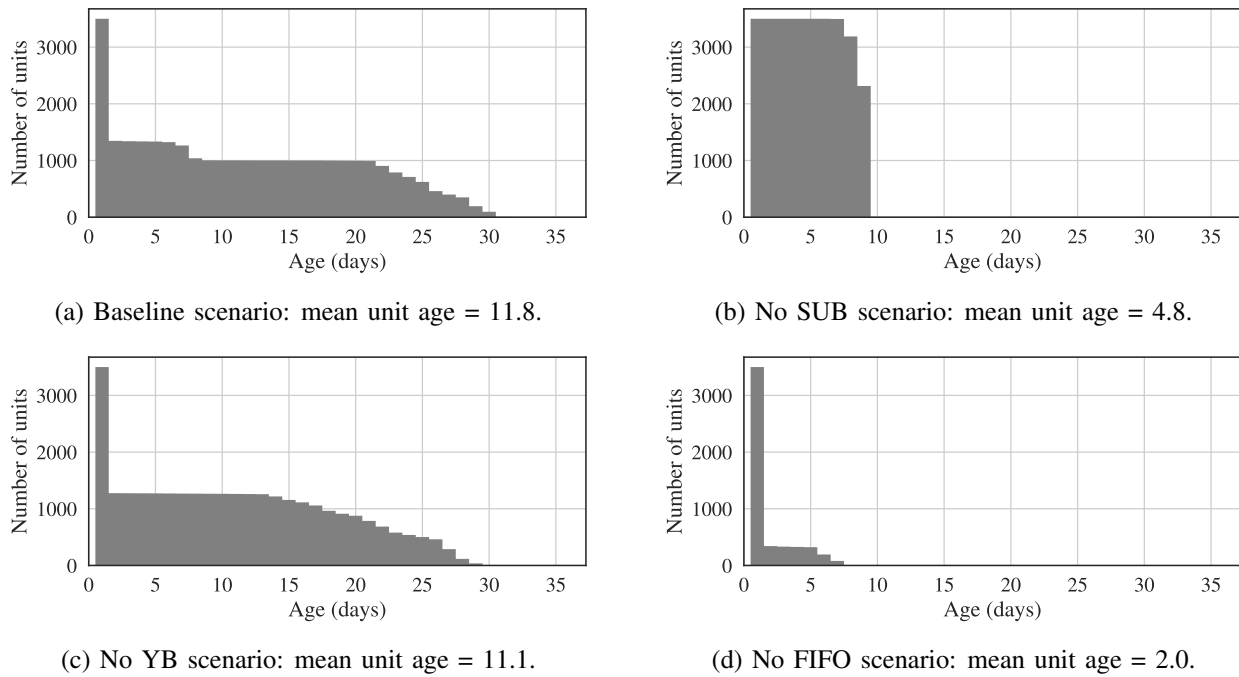


Figure 4: Mean steady-state (i.e., over replications and epochs) age distribution of RBCs in inventory. The standard error of the mean estimate for each bin is below one unit.

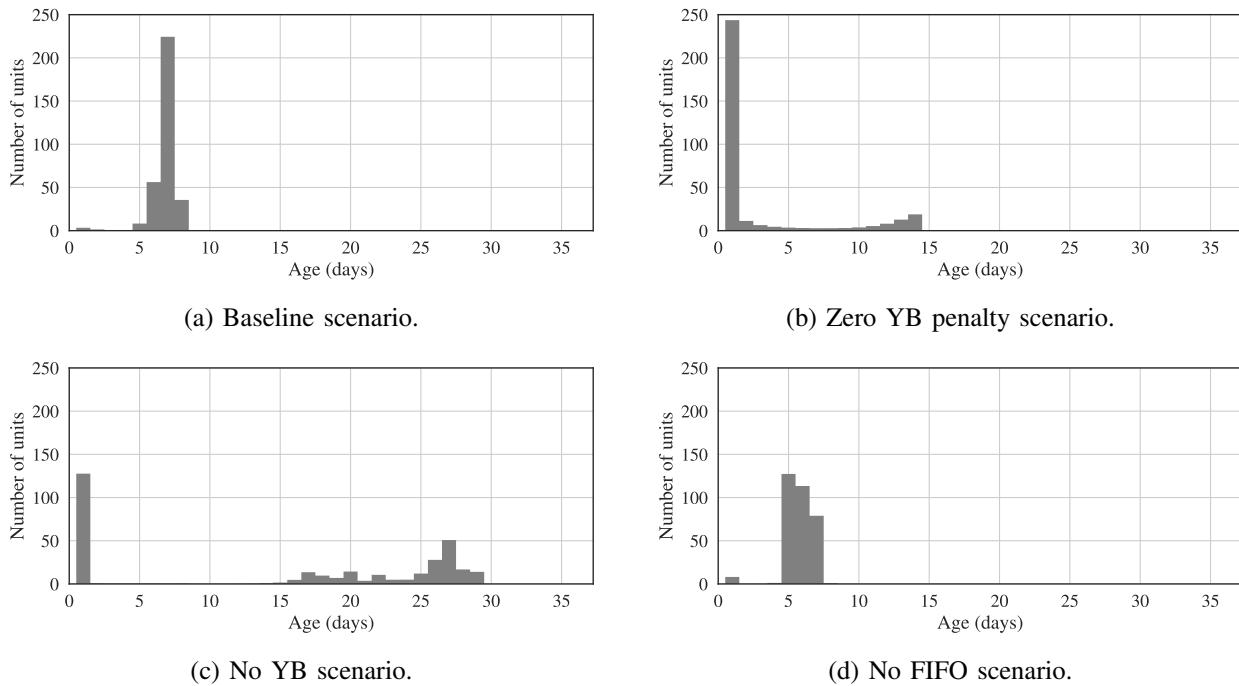


Figure 5: Mean steady-state age distribution of RBCs allocated to SCD patients. The standard error of the mean estimate for each bin is below one unit.

Table 4: Tuned policy parameters (λ^{30*} , λ^{15*} , λ^{10*} for 30, 15, and 10 thousand RBC units respectively) on three initial inventory sizes $|I|$.

$ I $ ('000)	λ_1	λ_2	λ_3	λ_4	λ_5
30	0.49	0.00	0.01	0.47	0.02
15	0.48	0.00	0.06	0.45	0.01
10	0.26	0.00	0.11	0.61	0.02

Table 5: Mean xA performance for each λ^* simulated for its respective initial inventory size and generalized to other $|I|$. All std. errors = 0.002.

$ I $ ('000)	λ^{30*}	λ^{15*}	λ^{10*}
30	0.133	0.156	0.187
15	0.174	0.183	0.195
10	0.205	0.206	0.213

modal age of 7 days and an age range of 1-9 days (Figure 5a). Whereas when the YB penalty is removed, the modal age is 1 day but the age range is 1-14 days (Figure 5b).

YB penalty and constraint removed (No YB, $\lambda_5 = 0$, $\pi_4(i, j) = -2^{-(G-g_i)/5}$ and $h_j = G = 35 \forall i, j$)
 In this scenario the FIFO penalty (Equation (8)) is modified to remove the dependency on the type of patient. This results in the lowest mean xA of all the penalty modifications, an 85% reduction compared to the value under **E0** with the complete penalty. Again, there are no shortages or expiries. As in the Zero YB scenario, the modal age of RBC units allocated to SCD patients is 1 day (Figure 5c). However, with the age constraint removed 60% of units allocated to SCD patients are older than 14 days and some are as old as 29 days. This has a slight effect on the age distribution of the RBC units in stock (Figure 4c).

5.3 Optimized Penalty Weights

Next, we tuned λ , the vector of penalty weights, using a Bayesian optimization (BO) procedure in Emukit (version 0.4) to minimize xA under the **E0** rule. BO is a popular approach for optimizing time-consuming black-box functions. We adopted the procedure described in Jones et al. (1998), using a Gaussian process to model the function $\lambda \mapsto xA$, with a Matern-5/2 kernel modified so that the distance between any pair of vectors of penalty weights is the Euclidean distance between the corresponding pair of normalized vectors. We set the search space for the penalty weights to $\lambda \in [0, 1]^5 \setminus \{\mathbf{0}\}$, used Expected Improvement as the acquisition function, and used L-BFGS as the acquisition optimizer. We estimated the objective at each evaluation by taking the mean xA over 20 simulation replications. The computation budget included 25 initial samples (in a Latin hypercube design) and a further 75 function evaluations.

The optimized weights are shown in Table 4. The optimized policy achieves an xA of 0.133 (Table 5), approximately half that achieved by the best policy (No YB) so far investigated (Table 3). The immunogenicity penalty (π_1) and the FIFO penalty (π_4) have large and similar weights. The substitution penalties (π_2 , π_3) and YB penalty (π_5) have small weights. This is consistent with our sensitivity analysis, which suggested π_1 and π_4 are important components for minimizing alloimmunisation. No shortages or expiries were observed in the tuned simulation, indicating there is no trade-off between these two cost metrics and xA.

The size of NHSBT RBC stock sometimes falls far below the 30,000 unit target. To assess the robustness of the optimized weights to such reductions, we applied the BO procedure while simulating inventories of 15,000 and 10,000 units. The latter corresponds approximately to the lowest stock level experienced by NHSBT. The tuned penalty weights are presented in Table 4 and the mean xA achieved (evaluated over 100 replications) in Table 5. Predictably, xA increases as inventory size decreases. No expiries or shortages were observed. As the inventory size decreases the importance of the minor antigen substitution penalty seems to increase. The FIFO (π_4) and immunogenicity (π_1) penalties are respectively more and less important in the smallest case (λ^{10*}) compared to the larger cases (λ^{30*} and λ^{15*}). When we simulated an inventory of each size with each set of tuned policy parameters in turn, we found the weights tuned for 30,000 units generalize the best. This indicates that these policy parameters would be robust to sustained period of lower stock. They also produce the best performance on all inventory sizes while the weights tuned on 10,000 units consistently performed the worst, although the performance gap reduces with inventory size. This suggests incomplete convergence of the BO run for the smallest inventory size.

6 CONCLUSIONS

We address a blood inventory problem in the context of genetically measured extended blood types: allocating RBC units to SCD patients (some with pre-existing alloantibodies) using a simulation with optimization-based iterations approach. Allocation decisions were determined by solving a minimum-cost maximum flow problem in an extension of a previously proposed penalty-based policy framework. Our results demonstrated an 81% reduction in alloimmunization risk compared to current guidelines. We then analyzed the effects of policy parameters on the composition of the inventory, and the age of units allocated to SCD patients. We tuned the parameters of the matching policy using BO to minimize alloimmunization in SCD patients. We find that the optimal parameter settings for an inventory of 30,000 units put a low weight on the substitution and young blood penalties. Furthermore these optimal settings are robust to reductions in inventory size.

Extensions of the work could include the simulation of alloimmunization events and patient returns, enabling the estimation of the number of patients forming such a combination of antibodies that they become untransfusable. Our simplified model of the NHSBT blood supply chain, which assumes a single central inventory is likely to allow better matching than can be achieved in practice. Improving this model to incorporate the hierarchical structure of the supply chain presents another avenue for future research.

ACKNOWLEDGMENTS

This work was funded by NHSBT, by the NIHR under an Artificial Intelligence in Health and Care Award (AI Award 02331) and through the NIHR BTRU in Donor Health and Behaviour (NIHR203337) and the NIHR Cambridge BRC (BRC-1215-20014; NIHR203312). The research of Merel Wemelsfelder is supported by the Sanquin grant PPOC20-05/L2490. The views expressed are those of the authors and not necessarily those of the NIHR or the Department of Health and Social Care.

REFERENCES

- Abdulwahab, U. and M. I. Wahab. 2014. "Approximate Dynamic Programming Modeling for a Typical Blood Platelet Bank". *Computers and Industrial Engineering* 78:259–270.
- Arani, M., X. Liu, and S. Abdolmaleki. 2020. "Scenario-Based Simulation Approach for An Integrated Inventory Blood Supply Chain System". In *2020 Winter Simulation Conference (WSC)*, 1348–1359 <https://doi.org/10.1109/WSC48552.2020.9384018>.
- Beliën, J. and H. Forcé. 2012. "Supply Chain Management of Blood Products: A Literature Review". *European Journal of Operational Research* 217(1):1–16.
- Civelek, I., I. Karaesmen, and A. Scheller-Wolf. 2015. "Blood Platelet Inventory Management with Protection Levels". *European Journal of Operational Research* 243(3):826–838.
- Clay, N. M., B. Abbasi, A. Eberhard, and J. Hearne. 2018. "On the Volatility of Blood Inventories". *International Transactions in Operational Research* 25(1):215–242.
- Ejohwomu, O. A., J. Too, and D. J. Edwards. 2021. "A Resilient Approach to Modelling the Supply and Demand of Platelets in the United Kingdom Blood Supply Chain". *International Journal of Management Science and Engineering Management* 16(2):143–150.
- Evers, D., R. A. Middelburg, M. de Haas, S. Zalpuri, K. M. de Vooght, D. van de Kerkhof *et al.* 2016. "Red-Blood-Cell Alloimmunisation in Relation to Antigens' Exposure and their Immunogenicity: A Cohort Study". *The Lancet Haematology* 3(6):e284–e292.
- Figueira, G. and B. Almada-Lobo. 2014. "Hybrid Simulation-Optimization Methods: A Taxonomy and Discussion". *Simulation Modelling Practice and Theory* 46:118–134.
- Gass, S. I. 1990. "On Solving the Transportation Problem". *Journal of the Operational Research Society* 41(4):291–297.
- Gleadall, N. S., B. Veldhuisen, J. Gollub, A. S. Butterworth, J. Ord, C. J. Penkett *et al.* 2020. "Development and Validation of a Universal Blood Donor Genotyping Platform: A Multinational Prospective Study". *Blood Advances* 4(15):3495–3506.
- Jones, D. R., M. Schonlau, and W. J. Welch. 1998. "Efficient Global Optimization of Expensive Black-Box Functions". *Journal of Global Optimization* 13(4):455–492.
- Katsaliaki, K. and S. C. Brailsford. 2007. "Using Simulation to Improve the Blood Supply Chain". *Journal of the Operational Research Society* 58(2):219–227.
- McElfresh, D. C., C. Kroer, S. Pupyrev, E. Sodomka, K. Sankararaman, Z. Chauvin *et al.* 2023. "Matching Algorithms for Blood Donation". *Nature Machine Intelligence* 5(10):1108–1118.

- Onggo, B. S. 2014. "Elements of a Hybrid Simulation Model: A Case Study of the Blood Supply Chain in Low- and Middle-Income Countries". In *2014 Winter Simulation Conference (WSC)*, 1597–1607 <https://doi.org/10.1109/WSC.2014.7020011>.
- Osorio, A. F., S. C. Brailsford, H. K. Smith, and J. Blake. 2018. "Designing the Blood Supply Chain: How Much, How and Where?". *Vox Sanguinis* 113(8):760–769.
- Osorio, A. F., S. C. Brailsford, H. K. Smith, S. P. Forero-Matiz and B. A. Camacho-Rodríguez. 2017. "Simulation-Optimization Model for Production Planning in the Blood Supply Chain". *Health Care Management Science* 20(4):548–564.
- Pirabán, A., W. J. Guerrero, and N. Labadie. 2019. "Survey on Blood Supply Chain Management: Models and Methods". *Computers and Operations Research* 112.
- Reid, M. E., C. Lomas-Francis, and M. L. Olsson. 2012. *The Blood Group Antigen Factsbook*. 3rd ed. Boston: Academic press.
- Silva Magalhães, V., L. R. Pinto, L. F. Rodrigues, and J. T. Blake. 2023. "Simulation-Optimisation Approach to Support Management of Blood Components Inventory". *Journal of Simulation* 18(4):671–686.
- Soares, H. L., E. F. Arruda, L. Bahiense, D. Gartner and L. Amorim Filho. 2020. "Optimisation and Control of the Supply of Blood Bags in Hemotherapeutic Centres via Markov Decision Process with Discounted Arrival Rate". *Artificial Intelligence in Medicine* 104:101791.
- van de Weem, R. H. G., M. L. Wemelsfelder, J. S. Luken, M. de Haas, R. W. Niessen, C. E. van der Schoot *et al.* 2022. "Preventing Alloimmunization Using a New Model for Matching Extensively Typed Red Blood Cells". *Vox Sanguinis* 117(4):580–586.
- van Sambeek, J. H., S. P. van Brummelen, N. M. van Dijk, and M. P. Janssen. 2022. "Optimal Blood Issuing by Comprehensive Matching". *European Journal of Operational Research* 296(1):240–253.
- Wemelsfelder, M. L., R. H. G. van de Weem, J. S. Luken, M. de Haas, R. W. L. M. Niessen, C. E. van der Schoot *et al.* 2024. "Extensive Red Blood Cell Matching Considering Patient Alloimmunization Risk". *Vox Sanguinis* 119(4):368–376.
- Williams, E. P., P. R. Harper, and D. Gartner. 2020. "Modeling of the Collections Process in the Blood Supply Chain: A Literature Review". *IISE Transactions on Healthcare Systems Engineering* 10(3):200–211.

AUTHOR BIOGRAPHIES

FOLARIN B. OYEBOLU is a Research Associate in the MRC Biostatistics Unit at the University of Cambridge. His research interests include meta-heuristics, decision-support tools, sustainable research software, and applications of simulation optimisation. His email address is folarin.oyebolu@mrc-bsu.cam.ac.uk.

MARIE CHION is a Research Associate in the MRC Biostatistics Unit at the University of Cambridge. She is interested in developing statistical methods for molecular-level data, including in haematology and transfusion medicine. Her email address is mc2411@cam.ac.uk.

MEREL L. WEMELSFELDER is a PhD Candidate in the Donor Medicine Research Department at Sanquin and in the Business Analytics department at the University of Amsterdam. She is interested in the interplay between mathematical optimization and machine learning methods, and their application in solving logistical challenges. Her email address is m.wemelsfelder@sanquin.nl.

SARA TROMPETER is a Consultant Haematologist and Paediatric Haematologist at University College London Hospitals and NHS Blood and Transplant. She co-leads the [HAEM-MATCH](#) consortium which aims to bring extended genotypically matched blood to the patient bedside to reduce the complications of transfusion. Her email address is sara.trompeter@nhs.net.

NICHOLAS GLEADALL is an Assistant Professor in Clinical Genomics of Transfusion and Transplantation in the Department of Haematology at the University of Cambridge and a Principal Investigator at NHS Blood and Transplant. He works at the interface between research and clinical application, aiming to bring the benefits of big data, advanced analytics, and genomics to the patient bedside. His email address is ng384@cam.ac.uk.

WILLIAM J. ASTLE is a University Senior Lecturer in the MRC Biostatistics Unit at the University of Cambridge and a Principal Investigator at NHS Blood and Transplant. He is interested in the development and application of statistical methods to answer questions in haematology, transfusion medicine and related fields. His e-mail address is wja24@cam.ac.uk.