



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2012-01

A novel allocation strategy for blood transfusions: investigating the tradeoff between the age and availability of transfused blood

Atkinson, Michael P.; Fontaine, Magali J.; Goodnough, Lawrence T.; Wein, Lawrence M.

TRANSFUSION

TRANSFUSION Volume 52, January 2012
<https://hdl.handle.net/10945/48003>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

A novel allocation strategy for blood transfusions: investigating the tradeoff between the age and availability of transfused blood

Michael P. Atkinson, Magali J. Fontaine, Lawrence T. Goodnough, and Lawrence M. Wein

BACKGROUND: Recent studies show that transfusing older blood may lead to increased mortality. This raises the issue of whether transfusing fresher blood can be achieved without jeopardizing blood availability.

STUDY DESIGN AND METHODS: We propose a simple family of policies that is defined by a single threshold: rather than transfusing the oldest available blood that is younger than 42 days, we transfuse the oldest blood that is younger than the threshold, and if there is no blood younger than the threshold then we transfuse the youngest blood that is older than the threshold. To assess this policy, we build a simulation model using data from Stanford University Medical Center. We focus on the tradeoff between the mean age of transfused blood and the fraction of transfused blood that is imported.

RESULTS: For hospitals in which the local supply is greater than demand, our policy with a threshold of 14 days leads to a decrease of 10 to 20 days in the mean age of transfused blood while increasing the fraction of imported blood to less than 0.005 (i.e., 0.5%). If the health benefits from transfusing fresher blood can be confirmed by randomized clinical trials, then conservative assumptions suggest that this policy could reduce the annual number of transfused patients who die within 1 year by 20,000.

CONCLUSION: The proposed allocation policy with a threshold of 14 days could allow many US hospitals to significantly reduce the age of transfused blood, thereby possibly reducing morbidity and mortality, while having a negligible impact on supply chain operations.

The allocation of available red blood cells (RBCs) for transfusion to patients is a difficult problem: compatibility between donor and recipient must be maintained, and the age and quantity of available RBCs need to be accounted for. More specifically, while there is a tendency to use older available blood to prevent outdated, stored RBCs undergo biochemical changes that impact RBC function, and the age of transfused blood has been reported to be associated with an increased risk of infection, postoperative complications, and/or short-term and long-term mortality.¹⁻⁶ However, other results have not confirmed this association, and a recent review of 24 studies concluded that it is difficult to determine if there is an association, except possibly in trauma patients.⁷ In contrast to similar resource allocation problems for organ transplants, there is not a centrally administered or widely accepted allocation policy for RBC transfusions, nor are there many detailed studies that propose or assess allocation strategies.⁸

The aim of our study was to assess whether there is an RBC allocation policy that can increase the health benefits (via transfusing fresher blood) without straining the blood supply, in the event that the results in the studies by Koch and colleagues,¹ Zallen and colleagues,² Mynster and Nielsen,³ Offner and colleagues,⁴ Leal-Noval and colleagues,⁵ and Purdy and colleagues⁶ can be confirmed by

ABBREVIATIONS: FIFO = first in, first out; LIFO = last in, first out; SUMC = Stanford University Medical Center.

From the Operations Research Department, Naval Postgraduate School, Monterey, California; and the Department of Pathology and Medicine and the Graduate School of Business, Stanford University, Stanford, California.

Address reprint requests to: Lawrence M. Wein, Graduate School of Business, 655 Knight Way, Stanford University, Stanford, CA 94305-5015; e-mail: lwein@stanford.edu.

Supported by the Center for Social Innovation, Graduate School of Business, Stanford University (LMW).

Received for publication February 1, 2011; revision received May 3, 2011, and accepted May 4, 2011.

doi: 10.1111/j.1537-2995.2011.03239.x

TRANSFUSION 2012;52:108-117.

randomized clinical trials. We constructed a simulation model based on data from 18,915 nonirradiated RBC transfusions over a 1-year period at the Stanford University Medical Center (SUMC). While there are exceptions due to, for example, neonatal surgeries, the nominal policy used at SUMC is first in, first out (FIFO) among RBCs that are an exact match (i.e., among the eight blood types) and then FIFO among all compatible blood types if there are no RBCs that are an exact match. We introduce a novel family of allocation policies, which is indexed by a single variable and which includes SUMC's nominal policy as a special case. Our simulation model is used to assess this family of policies in terms of the tradeoff between blood availability and the age of transfused blood.

MATERIALS AND METHODS

The key tradeoff

The output of our simulation model quantifies the tradeoff between two performance measures. To motivate the first of our two performance measures, we performed some simple accounting (Fig. 1) that led to an equation for the conservation of blood. Locally (i.e., at the community blood center) donated units are either transfused or outdated (i.e., discarded if not transfused before the expiration date). Although all demanded transfusions are performed in our simulation model (canceled surgeries due to lack of RBCs are exceedingly rare in practice, with fewer than 500 nationwide in 2006,⁹ and SUMC has not canceled or delayed a surgery in years due to lack of blood), if compatible blood for a transfusion is not locally available then blood is imported in our simulation model. In practice, this blood would be obtained from, for example, an other-than-usual blood bank or the National Blood Exchange,¹⁰ although our simulation model, as explained below, does not specify the origin of the imported blood. Because there are no canceled surgeries, the amount of outdated blood

equals the amount of locally donated blood plus the amount of imported blood, minus the total amount of transfused blood; this equation describes the conservation of RBCs, and Fig. 1 provides the actual values of these quantities over a 1-year period at SUMC. As explained later, the amount of locally donated blood and the amount of transfused blood over any time period (e.g., 1 day or 1 year) are random variables that are not under direct managerial control in our model. It follows that, for any given amounts of locally donated blood and transfused blood over any time period, minimizing the amount of imported blood over that time period is equivalent to minimizing the amount of outdated blood over that time period; that is, if we minimize the amount of imported blood over 1 year then we automatically also minimize the amount of outdated blood over 1 year and vice versa. Hence, we need to consider only one of these two performance measures in our tradeoff analysis, and we choose the amount of imported blood because its cost is more quantifiable (indeed, the outdate cost is essentially sunk). The other performance measure in our key tradeoff is the mean age of transfused blood, and our goal is to simultaneously minimize both performance measures (i.e., the mean import rate and the mean age of transfused blood). Because the blood allocation policy can reduce one of these measures only at the expense of increasing the other, we investigate this problem with the use of tradeoff curves, as explained below.

Family of allocation policies

We introduce a new family of allocation policies that is simple to use, allows for varying importance to be placed on the age of transfused blood, and never discards blood before it is outdated. These policies will be tested in the simulation model that is described later. This family of policies is characterized by a single threshold, which is

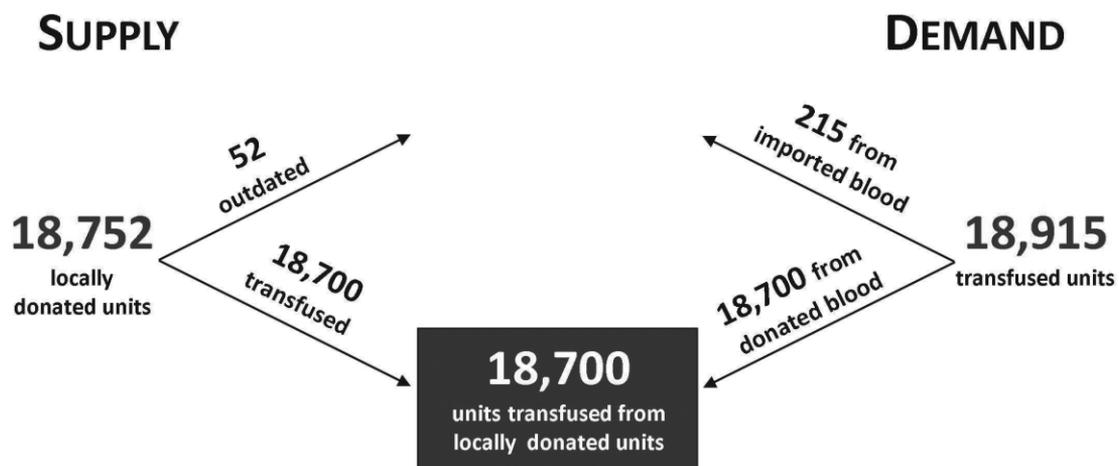


Fig. 1. Supply and demand for units of RBCs at SUMC during April 1, 2008, to March 31, 2009.

measured in days; all blood in our model takes 2 days to test and process before being available for transfusion, and outdates at 42 days of age, and so the threshold is a managerial variable that takes on an integer value between 2 and 42. When a patient with a certain blood type (i.e., O+, O-, A+, A-, B+, B-, AB+, or AB-) requires a transfusion of 1 unit of blood in our model, he goes through the following four-step procedure to determine what type of blood he or she will receive: 1) if there are RBCs that are an exact match and are younger than the threshold, then he or she receives the oldest unit younger than the threshold; 2) otherwise, if there are RBCs that are an exact match and are older than the threshold, then he or she receives the youngest unit older than the threshold; 3) otherwise, if there are RBCs of any compatible blood type (e.g., if the patient is type A+ then he or she could receive type O+, O-, or A-) that are younger than the threshold, then he or she receives the oldest (over all compatible types) of these RBCs; 4) otherwise, if there are RBCs of any compatible blood type that are older than the threshold, then he receives the youngest (over all compatible types) of these RBCs. If the patient does not receive any RBCs from this four-step procedure then there are no exact or compatible units available from local donations, and this transfusion requires imported blood. Note that setting the threshold equal to 2 or 42 days corresponds to the last in, first out (LIFO) and FIFO policies (within exact and compatible matches), respectively. Referring back to our performance measures, by varying the threshold from 2 to 42 days, we sweep out a tradeoff curve of the mean age of transfused blood versus the fraction of transfused blood that is imported. A variant of this policy that applies to only a known subset of patients (e.g., cardiac patients) is considered (and discussed in detail) under Results.

Simulation model

We develop a simulation model of a queuing system in which locally donated units of RBCs arrive to the queue, wait for service, and leave the queue at the end of service. The server in this queuing system transfuses blood into patients. More specifically, a service completion occurs at the time that a unit of RBCs is transfused into a patient, and a service time in the simulation model corresponds to the time between consecutive units of blood that need to be transfused. RBCs reaching the age of 42 days before completing service leave the queue due to outdating. If, at the end of a service, the four-step procedure does not generate an allocation (i.e., there was no exact or compatible blood in the queue), then the service is recorded as using 1 unit of imported blood. The simulation model only tallies the amount of imported blood and does not explicitly track the origin, type, or age of imported blood. Hence, the mean age of transfused blood is calculated using only the transfusions from local donations. It is possible that

imported blood might be slightly older, on average, than locally donated blood, at the time of transfusion (e.g., if blood is obtained from the National Blood Exchange, which in turn receives blood just before it outdates). However, the import rate is sufficiently low that this omission does not affect our qualitative results.

With the allocation policy described earlier, the model is fully specified by the exact timing of each donation and each transfusion and the blood type of each donation and each patient. We first describe the mix of blood types and then the timing. When a donated unit arrives to the queue, it is randomly assigned a blood type according to probabilities calculated from 18,752 nonirradiated units transfused at SUMC during April 2008 to March 2009 (Table 1). Although we do not know the blood types of all donated (as opposed to all transfused) units at SUMC, the blood type probabilities in Table 1 should be very close to the true blood type probabilities among all donated units because (as shown below) the outdate rate at SUMC was very small. Blood types for the recipients of these transfusions were not available. Consequently, when a unit of blood is transfused in our simulation model, the blood type of the patient receiving the transfused unit is randomly assigned according to probabilities calculated from 26,053 patients receiving (irradiated or nonirradiated) RBC transfusions at SUMC during July 1, 2007, to June 30, 2008 (Table 1). However, the bias due to irradiation occurs on the donor side; for example, SUMC tends to irradiate more group O blood for neonatal cardiac surgeries. Blood type mix for patients should also be representative of the blood type mix for patients receiving only nonirradiated blood, as long as, for example, the neonatal blood type mix does not differ substantially from the adult blood type mix.

The timing of the actual donations and transfusions at SUMC during April 2008 to March 2009 exhibits several complexities: predictable seasonality with weekly and annual cycles, multiple units of RBCs being donated by the same donor or transfused into the same patient on the same day, and statistical dependence between the donor and transfusion processes (e.g., expedited blood drives when RBC inventory is low). To allow for the flexibility to extrapolate beyond the specific SUMC setting (e.g., by

TABLE 1. Mix of blood types for donated units and patients

Blood type	Donations	Patients
O+	0.4323	0.4399
O-	0.0741	0.0417
A+	0.3024	0.3007
A-	0.0510	0.0337
B+	0.0952	0.1312
B-	0.0172	0.0115
AB+	0.0209	0.0371
AB-	0.0069	0.0042

varying the supply-to-demand ratio) while mimicking the impact of these realistic aspects of the SUMC system, we assume that the arrival process of donated units and the service process of transfused units are independent renewal processes (i.e., the time between consecutive events are independent and identically distributed random variables¹¹), each having the same coefficient of variation (CV), which is the standard deviation divided by the mean, of the time between consecutive events (i.e., donations or transfusions, respectively). Queuing theory suggests that the key performance measures can be accurately and succinctly derived from the mean donation rate, the mean transfusion rate, and the CV (Sections 9.6, 9.9, and 10.2 in Whitt¹¹). That is, the CV is a surrogate measure that accounts for the complexities mentioned above.

We calculate the CV, which is an input to our model, from data related to two performance measures at SUMC: the fraction of transfused units that were imported and the fraction of donated blood that was outdated. The Stanford Blood Center's estimate of the import rate during April 2008 to March 2009 was between 0.005 and 0.015, but they lacked explicit data to derive a more precise value. Consequently, we assume that a unit was imported during April 2008 to March 2009 if it was transfused on the same day it was delivered to the hospital, which yields an import rate of $215/18,915 = 0.0114$. Because this estimate is very crude, we consider a wide range of import rates (indeed, much wider than 0.005-0.015) in the sensitivity analyses. We have outdate data for SUMC only during July 2009 to June 2010, and we assume that the number of outdates during this time period is the same as during our time period of interest. Outdates can occur at either the SUMC or the Stanford Blood Center; we include the outdates at the Stanford Blood Center because SUMC is by far their largest customer. There were 52 units with a shelf life of 42 days that outdated at the blood center during July 2009 to June 2010. There were an additional 30 units that outdated at SUMC but we were unable to get information about their shelf life, and we did not include these 30 units because 94% of the outdates at the Stanford Blood Center had shelf lives of less than 42 days (due to irradiation). Thus, we estimate the outdate rate at SUMC to be $52/18,752 = 0.0028$. Note that $18,752 = 18,915 - 215 + 52$ because the RBCs for the transfusions requiring imported blood were not from local donations, but from elsewhere (Fig. 1). Consequently, the donation rate is $18,752/365 = 51.38/\text{day}$ and the transfusion rate is $18,915/365 = 51.82/\text{day}$. We refer to the donation rate as the supply and the transfusion rate as the demand, giving a supply-to-demand ratio of $18,752/18,915 = 0.9914$. We simulate the queuing system to find the value of the CV that minimizes the sum of the squared deviations between the exact and simulated import rate plus the exact and simulated outdate rate. This approach yields a CV of 1.32 (resulting in an import rate of 0.0108 and an

outdate rate of 0.0037), which is slightly larger than the canonical value of 1.0 that corresponds to the case in which donations and transfusions are independent Poisson processes. Recall that the supply-to-demand ratio of 0.9914 is a ratio of annual rates; in our simulation model, the daily ratio of donated units divided by transfused units will vary considerably due to the large amount of statistical uncertainty.

In summary, the timing and blood type of each donated unit and each transfused unit are random in the simulation model and do not correspond in a one-to-one manner to the timing and blood type of each donated unit and each transfused unit at SUMC during April 2008 to March 2009. However, the input variables of the simulation model (donor rate, transfusion rate, blood type mixes in Table 1, and CV) are chosen so that the aggregate output measures of the model (e.g., outdate rate, import rate) are roughly consistent with the aggregate performance measures experienced at SUMC. In particular, the high value of the CV leads to significant day-to-day variation in the number of donations and transfusions.

RESULTS

Main results

The tradeoff curves are generated by increasing the threshold from 2 to 42 days, with 42 corresponding to the upper left point on each curve and 2 corresponding to the lower right point (Fig. 2). The nature of the tradeoff between the mean age and the import rate of transfused blood is highly dependent on the supply-to-demand ratio (Fig. 2). When the supply-to-demand ratio is less than or equal to 0.98, the tradeoff curve is nearly horizontal (on the linear scale) and the mean age of transfused blood is only several days, regardless of the value of the threshold. In contrast, when the supply-to-demand ratio is greater than or equal to 1.06, the tradeoff curve is almost vertical and barely any imports occur, irrespective of the value of the threshold. In the latter case, the status quo policy (i.e., the threshold equals 42 days) leads to a mean age of approximately 35 days when the supply-to-demand ratio is between 1.04 and 1.10.

Sensitivity analyses

We assess the robustness of our results with respect to two variables: the amount of variability in the system (as measured by the CV of the arrival and service processes) and the size of the hospital (as measured by the donation rate and transfusion rate). We also consider a change in the blood type mix of donations and recipients.

Recall that the CV was derived from our estimates of the SUMC import rate and outdate rate, both of which may be inaccurate. Before showing the results of the sensitivity analysis of the CV, we consider how much the CV

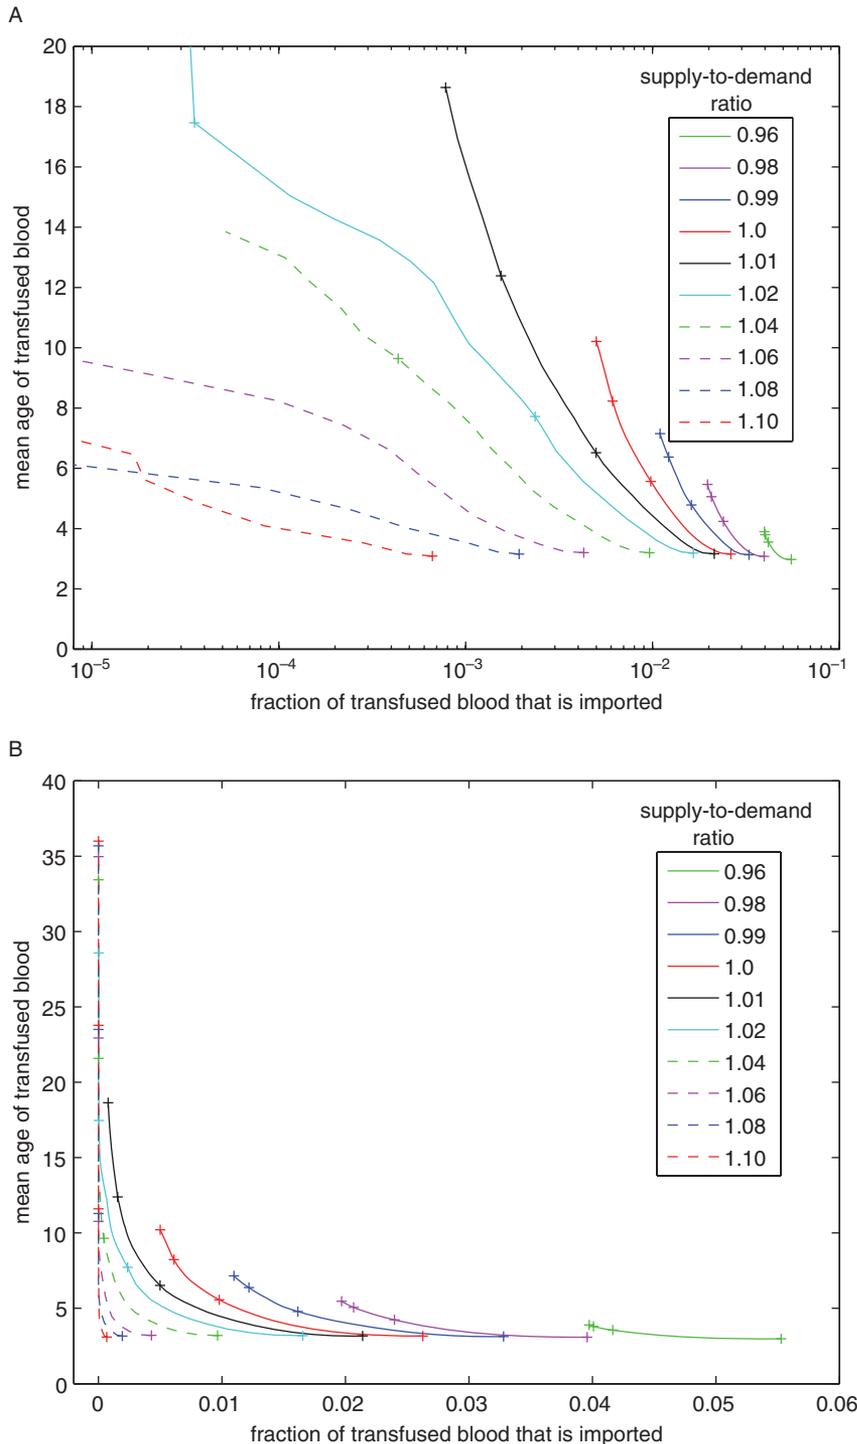


Fig. 2. The tradeoff curve for various values of the supply-to-demand ratio, with the import rate on a (A) logarithmic and (B) linear scale. The four tick marks on each curve correspond to threshold values of 2, 14, 28, and 42 days. Because 500,000 transfusions were simulated to generate each point on these curves, we are not able to accurately estimate import rates of less than 10⁻⁵.

would change if we changed our estimates of the SUMC import rate or outdate rate. Changing the import rate from 0.0114 to 0.002 or 0.05 changes the CV (via least squares, as described under Materials and Methods) from 1.32 to approximately 0 and 2.28, respectively. The simulation did not fit the data well when both the import rate and the outdate rate were less than 0.003. Changing the outdate rate from 0.0028 to its upper bound of 82/18,752 by including the 30 units outdated at the Stanford Blood Center, yields a CV of 1.55. Consequently, we choose to vary the CV between 0.01 and 2.64 (i.e., double the base value) in our sensitivity analysis, which comfortably subsumes these changes in the import rate and outdate rate.

Analogs of Fig. 2 when the CV equals 0.01 and 2.64 (rather than the base case of 1.32) show that the tradeoff curves shift slightly to the left when the CV is 0.01 and slightly to the right when the CV is 2.64 (data not shown). Similarly, multiplying the donation and transfusion rates by 0.5 shifts the tradeoff curves slightly to the right, and multiplying the rates by 1.5 shifts the curves slightly to the left (data not shown). Overall, our qualitative conclusions are unaffected by large changes in the amount of variability or the size of hospitals.

Finally, because we do not have precise blood type data for the patients at SUMC during April 2008 to March 2009, we recompute Fig. 2 under two different assumptions to examine how sensitive our results are to the blood type distributions of recipients and donations. Both of these scenarios use distributions that are likely to be farther from the true values for recipients or donations during April 2008 to March 2009 than the values used in our base-case analysis (Table 1). First, we assume that the blood type probabilities for patients (the last column of Table 1) are replaced by the blood type mix of donated units in Table 1. Second, we replace the blood type probabilities for donations by the blood type probabilities among all units donated to the Stan-

ford Blood Center during July 2009 to June 2010, which are 0.4037, 0.0918, 0.2797, 0.0624, 0.0995, 0.0192, 0.0354, and 0.0083. This mix differs from the mix of donated units in Table 1 because it includes both irradiated and nonirradiated units and includes units shipped to hospitals other than SUMC.

System performance degrades when the patient and donation distributions are the same, with the import rate nearly an order of magnitude larger than the base case in some instances (data not shown). However, when we use the donation mix from Stanford Blood Center during July 2009 to June 2010, the performance exceeds that of the base case and there are no imports for a supply-to-demand ratio greater than 1.02. Nonetheless, the qualitative behavior from Fig. 2B is maintained in both cases, with the tradeoff curves transitioning from horizontal to vertical as the supply-to-demand ratio increases. The discrepancy in performance is due to the fact that the system can more flexibly satisfy demand when the proportion of blood that is D- is larger among donations than among patients. Our three scenarios vary significantly in this regard: 15% of donations and only 9% of patients have D- blood in the base case (Table 1), and 18% of donations to the Stanford Blood Center during July 2009 to June 2010 were D-. Therefore, consideration should be given to the proportion of D- blood among donations and patients, which may vary geographically, when determining an appropriate threshold.

An illustrative assessment of costs and health benefits

Although there is essentially no tradeoff between our two performance measures when the supply-to-demand ratio is 1.06 or greater, there is a tradeoff between costs and benefits when this ratio is less than 1.06. For illustrative purposes, the quantification of this tradeoff is calculated using the tradeoff curve corresponding to a supply-to-demand ratio of 1.01 in Fig. 2 and a retrospective study on the impact of long storage times on 1-year death rates;¹ the veracity of the results in the study by Koch and colleagues¹ is addressed under Discussion. The points (0.0008,18.5) and (0.005,6.5) on the curve in Fig. 2, which are generated by threshold values of 42 and 14 days, respectively, imply that an increase in the fraction of imported blood by 0.0042 corresponds to a reduction in the mean age of transfused blood of 12 days. Koch and colleagues¹ observed an association for the 1-year death rate of 7.4% for 2872 cardiac patients receiving blood stored for 14 days or less (median, 11 days) and 11.0% for 3130 cardiac patients receiving blood older than 14 days (median, 20 days).¹ To obtain a rough estimate, we make two simplifying assumptions: we equate the mortality reduction due to the difference in mean age between 6.5 and 18.5 days for the two points on our curve to the mor-

tality reduction due to the difference in median age between 11 and 20 days in the study by Koch and colleagues¹ and we conservatively assume that the benefits from fresher blood apply only to 27% of the transfused population (27% of transfused units at SUMC were for cardiac patients,⁸ and the mean number of transfused units for cardiac patients and the general population are similar: 3.26¹ and 3.0,⁹ respectively). These two assumptions yield that the reduction in the 1-year death probability generated by an increase of 0.0042 in the fraction of imported blood is $0.27 (0.11 - 0.074) = 0.01$. The cost associated with imported blood equals the typical cost of blood collection, production, and distribution (approx. \$200/unit¹²) plus the additional cost to buy and transport blood from a nonlocal source (approx. \$50-\$100/unit, depending on the market value, the time of year, and whether there is a surplus of blood). At 3 units of blood per patient,⁹ the expected cost per transfused patient to reduce the 1-year death probability by 0.01 is $3 (0.005 - 0.0008) \$300 = \3.78 . However, we must temper this argument slightly because the imported blood is likely to be older and would possibly be associated with an adverse outcome. Finally, if we wanted to focus on serious adverse events (as defined by the Society of Thoracic Surgeons) rather than 1-year mortality, using the two endpoints in Fig. 2 of Koch and colleagues¹ yields a reduction in the serious adverse event probability (due to an increase in the fraction of imported blood of 0.0042) of $(0.27) (12 \text{ days}) (0.34 - 0.16)/(42 \text{ days}) = 0.0139$.

We can perform similar calculations—again using the results in the study by Koch and colleagues¹—to obtain rough estimates for the total health benefits achievable by the proposed policy. The number of patients receiving transfusions was 14.461×10^6 units/3.0 units/patient in 2006.⁹ We assume that health benefits from fresher blood are received only by cardiac patients (27% of total) who receive nonirradiated blood (92% of total, Whitaker et al., p. 28⁹), because we only study nonirradiated transfusions. Recalling that the nationwide supply-to-demand ratio is 1.085, we conservatively assume that only half of the nation's hospitals employ our policy, that these hospitals have a supply-to-demand ratio of 1.04, and that they use a threshold of 14 days rather than 42 days, thereby reducing the mean age of blood by 23.8 days, from 33.4 days to 9.6 days (Fig. 2). The proposed policy would then reduce the annual number of adverse events by $(14.461 \times 10^6/3.0) (0.92) (0.27) (0.5) (23.8 \text{ days}) (0.34 - 0.16)/(42 \text{ days}) = 61,066$. Conservatively assuming that the percentage reduction in 1-year death probability generated by a difference in mean age between 9.6 and 33.4 days in Fig. 2 is the same as the reduction generated by a difference in median age between 11 and 20 days in the study by Koch and colleagues,¹ the proposed policy reduces the annual number of patients who die within 1 year of transfusion by $(14.461 \times 10^6/3.0) (0.92) (0.27) (0.5) (0.11 - 0.074) = 21,552$.

Targeting a known subset of patients

Improvements in mortality and morbidity from receiving fresh blood have mainly focused on special subsets of patients, such as trauma patients⁷ or those undergoing cardiac surgery.¹ Therefore, we next examine the case in which the threshold policy applies to only a certain segment of the patient population (referred to as Type I patients). That is, we modify our policy so that Type I patients continue to use the proposed policy with a discretionary threshold, but the remaining (i.e., Type II) patients use a threshold of 42 days, which corresponds to using the oldest available exact match or the oldest available compatible match if no exact match is available. For concreteness, we define Type I patients to be cardiac patients. Based on data during April 2008 to March 2009, 27% of transfused units at SUMC are for cardiac surgery.⁸ Therefore, in our simulation model 27% of the transfused units are randomly assigned to be Type I and the remaining 73% are defined to be Type II.

As we vary the threshold for Type I patients, we sweep out two tradeoff curves: one for Type I patients and one for Type II patients (Fig. 3). A comparison of Fig. 2 (which corresponds to 100% of patients being Type I) and Fig. 3 allows for several observations. As expected, when Type I patients use a threshold of 42 days, both curves coincide in Fig. 3 (i.e., they intersect at the upper left part of these curves, which is also the same as the upper left point of the curves in Fig. 2). The import rates in Fig. 3 are nonzero for supply-to-demand ratios less than or equal to 1.02 and are less than the corresponding import rates in Fig. 2, which is not surprising given that in Fig. 2 the import rates are nonzero for supply-to-demand ratios up to 1.10 for patients with discretionary thresholds, but are nonzero only up to a ratio of 1.02 for patients with a threshold of 42 days (which comprises 73% of the transfused units in Fig. 3). The tradeoff curves for Type I patients in Fig. 3 dominate (i.e., lie to the southwest

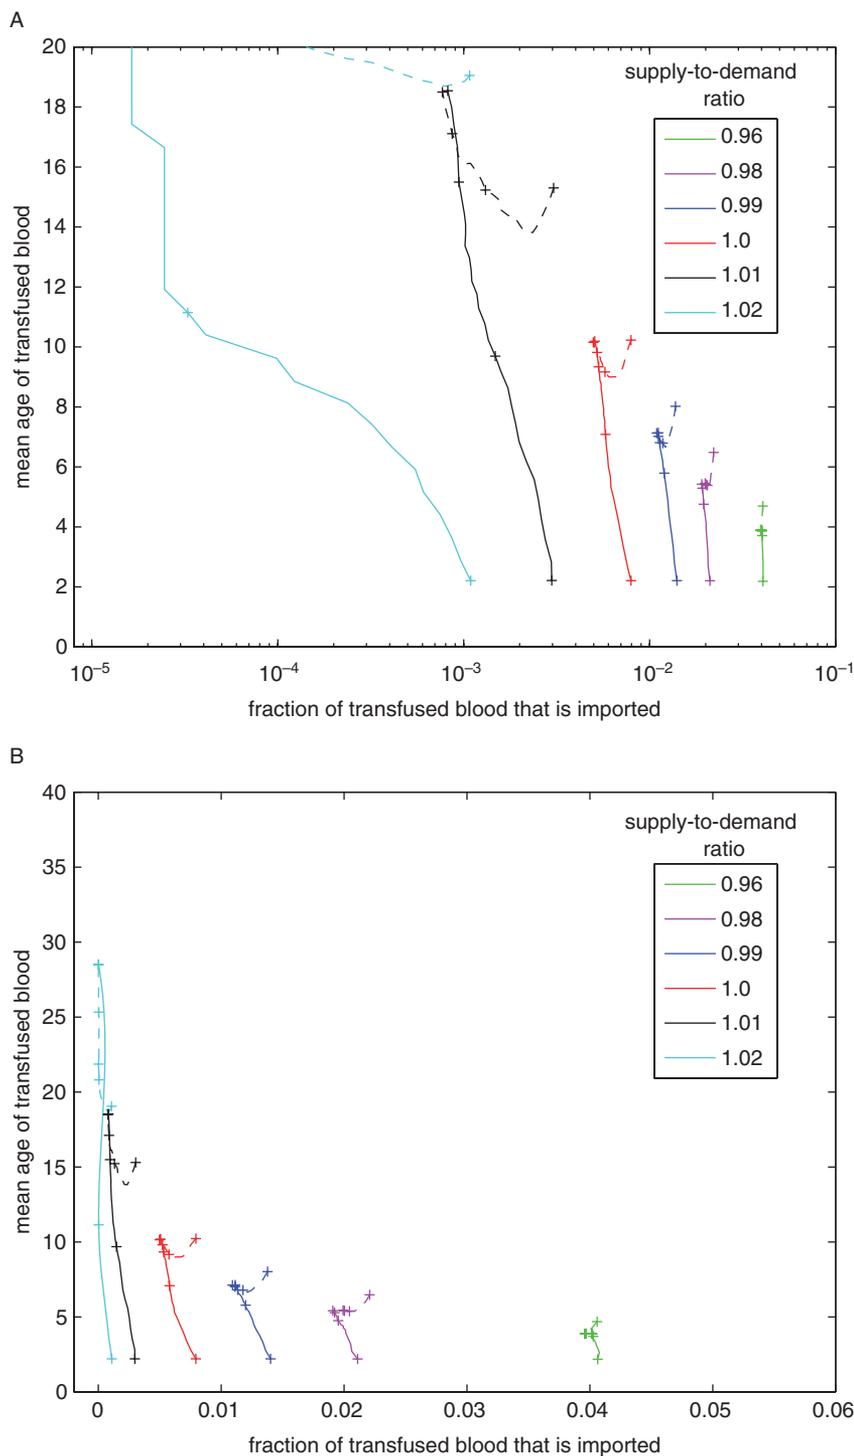


Fig. 3. Two tradeoff curves (a solid curve for the Type I patients that have a discretionary threshold and a dashed curve for the Type II patients that have a threshold of 42 days) for various values of the supply-to-demand ratio, with the import rate on a (A) logarithmic and (B) linear scale. The four tick marks on each curve correspond to threshold values (for the Type I patients with a discretionary threshold) of 2, 14, 28, and 42 days. Because 500,000 transfusions were simulated to generate each point on these curves, we are not able to accurately estimate import rates of less than 10⁻⁵.

of) the original tradeoff curves in Fig. 2, but this is at the expense of Type II patients; indeed, the weighted means of the dashed and solid curves in Fig. 3 are dominated by the tradeoff curves in Fig. 2 for positive import rates. For supply-to-demand ratios of 1.02 or greater, there is a large difference in the mean age of transfused blood for the two types of patients (e.g., at a ratio of 1.04 and a threshold of 14 days, the mean age is 12 days for Type I patients and 31 days for Type II patients). The tradeoff curves for Type II patients are not always monotonic: at small supply-to-demand ratios, as the threshold rises from small values, Type I patients will receive relatively older blood, which gives Type II patients an opportunity to receive fresher blood.

DISCUSSION

Many physical¹³ and biologic¹⁴ systems undergo phase transitions, where system behavior changes dramatically as a key variable exceeds a certain level, and queuing models are no exception: their behavior depends greatly on whether supply is greater than, or less than, demand (Fig. 2). The nationwide supply-to-demand ratio of whole blood and RBCs was $15.688 \times 10^6 / 14.461 \times 10^6 = 1.085$ in 2006.⁹ Due to the difficulty in predicting supply (e.g., yield from a blood drive) and demand at the local level and the desire to hedge against shortages, it is optimal (based on standard results from inventory theory¹⁵) to hold a small amount of safety stock in excess of forecasted demand and hence to target a supply-to-demand ratio above 1.0. However, friction in the marketplace (e.g., the high relative cost of transporting blood long distances, limited power of small hospitals in their supply chains), and geographical variation in supply and demand cause the supply-to-demand ratio to vary significantly across hospitals and years, from below 1.0 (e.g., it was 0.9914 at SUHC in 2007-2008) to well above 1.0.

Consequently, many hospitals operate on each side of the phase transition, and rather than having a consistent nationwide allocation policy, a considerable reduction in the age of transfused blood can be achieved by having an allocation policy that varies according to whether supply is greater than, or less than, demand. More specifically, when supply is less than demand, it is not possible to significantly alter the age of transfused blood (see Fig. 2, where fresh blood is transfused regardless of the value of the threshold), and the status quo policy of allocating the oldest available blood (i.e., using a threshold of 42 days) should remain in use. However, when supply is greater than demand, which presumably is the state of affairs at the majority of hospitals, it is very unlikely that imports will occur due to unavailable local blood, and consequently the amount of outdated blood is not under managerial control (aside from the caveat that local surplus can be salvaged on the National Blood Exchange in the days

before expiration) and is equal to the supply minus the demand (e.g., in 2006, nationwide supply minus demand was $15.688 \times 10^6 - 14.461 \times 10^6 = 1.227 \times 10^6$ and the actual number of outdated units was 1.276×10^6 ,⁹ but some hospitals had supply less than demand). That is, if supply is greater than demand, then outdated units are unavoidable. Hence, if a reduction in the age of transfused blood is deemed to be beneficial (more on this below), then when supply is greater than demand, a cultural shift would be needed among system managers, where the traditional emphasis on minimizing outdated units needs to be replaced by an emphasis on transfusing fresher blood. Our results suggest that a very simple policy, using a threshold of 14 days, can significantly reduce the age of transfused blood while having a minimal impact on the amount of imported blood (Fig. 2). This approach can be used not only across hospitals, but also across the time of year within a hospital, for example, using the status quo policy during the summer months and winter holidays when supply is low and using lower thresholds during the remainder of the year.

The above argument relies on the assumption that—when supply is greater than demand—the benefits from a reduction in the age of transfused blood outweigh the costs of additional imported blood. Although there is essentially no tradeoff when the supply-to-demand ratio is 1.06 or greater, there is a tradeoff when the supply-to-demand ratio is less than 1.06. For illustrative purposes, the quantification of this tradeoff was calculated using a retrospective study on the impact of large storage times on 1-year death rates.¹ However, the health impacts reported by Koch and colleagues¹ are larger than those in most other studies that also found benefits from fresher blood²⁻⁶ and Koch and colleagues¹ have come under criticism.¹⁶ Moreover, a review of 24 studies—many of which were observational cohort studies—found contrasting evidence, with some studies generating no impact on outcomes from the age of transfused blood.⁷ Hence, until the results of Koch and colleagues¹ can be confirmed by randomized clinical trials, they should be viewed as an upper bound on the impact of blood age on health outcomes. However, our tradeoff calculations are conservative in every other way (i.e., aside from the use of the results in Koch et al.), and our tradeoff results—if the benefits reported by Koch and colleagues¹ can be confirmed—are striking: \$3.78 per transfused patient to reduce the 1-year death probability by 0.01 and an annual reduction in the number of patients who die within 1 year of transfusion by 21,552. In any case, ongoing clinical trials may shed more light on this topic,¹⁷ and our calculations provide a basis for quantifying the costs and benefits associated with a change from a FIFO-base policy to the proposed policy.

Turning to the targeted policy, a comparison of Figs. 2 and 3 suggest that if indeed a known class of patients are sensitive to the age of transfused blood (e.g.,

27% of transfused units went to cardiac patients in our calculations, based on SUMC data) and the remaining patients' medical outcomes are independent of the age of transfused blood, then the targeted policy (where Type I patients have a discretionary threshold and Type II patients have a threshold of 42 days) outperforms the untargeted policy in which all patients have the same threshold. Given the equity issues generated by the targeted policy and the current uncertainty regarding the clinical impact of receiving different age blood for various classes of patients, coupled with the obvious benefits generated by the untargeted policy, we believe that the untargeted policy is preferable to the targeted policy. Moreover, one could investigate more complicated policies in which the threshold varies with the blood types of both the donated blood and the patient. While this approach may improve the average performance (as measured by the tradeoff curves) and could be used to reduce inequity of the mean age of transfused blood across the blood types of recipients, the increased complexity of such a policy makes it less attractive for implementation.

Our study has several limitations. Our model captures some of the complexities in the timing of donations and transfusions in a succinct but approximate manner. Every hospital differs on a variety of dimensions—the surgical volume, the exact mix of blood types for donors and patients, the supply-to-demand ratio, the mix of procedures it performs, the nature of seasonality, the correlation between demand and replenishment, the details of their supply chain (e.g., blood centers vs. American Red Cross, the fraction of their main supplier's RBCs that a hospital consumes)—that can affect our main tradeoff curve, and we have focused on data from a single hospital in a single year. In addition, our model considers only nonirradiated RBCs with a shelf life of 42 days (to incorporate the approximately 10% of the US blood supply that is currently irradiated, our model would need to be generalized to 16 blood types—eight with shelf lives of 42 days and eight with shorter shelf lives—and a policy with two thresholds: one for irradiated blood and one for nonirradiated blood); ignores the need for cytomegalovirus and other secondary compatibility considerations; and does not incorporate crossmatching, pediatric patients, or an array of rare situations. A few of these factors may cause the model to underestimate the mean age of transfused blood and overestimate the impact of the threshold on the mean age of transfused blood. Nevertheless, our sensitivity analyses—and a large body of queuing theory implying that the first-order effect in queuing systems is due solely to the supply-to-demand ratio¹¹—suggest that our main result is robust.

In conclusion, when the annual local supply of RBCs exceeds annual demand, as it does in many US hospitals, the excess supply will eventually outdate (perhaps at the

community blood centers rather than at the hospitals), regardless of the allocation policy. In this case, the easy-to-implement family of allocation policies introduced here, which provides a continuum of policies between FIFO and LIFO, has the potential to significantly reduce the mean age of transfused blood, possibly providing considerable health benefits, while having a negligible impact on blood availability.

ACKNOWLEDGMENT

We thank Maurene Viele for the patient data in Table 1 and Susan Galel and Harry Sussmann for estimates of the outdate rate and import rate.

CONFLICT OF INTEREST

The authors declare that they have no conflicts of interest relevant to the manuscript submitted to **TRANSFUSION**.

REFERENCES

1. Koch CG, Li L, Sessler DI, Figueroa P, Hoeltge GA, Mihalevic T, Blackstone EH. Duration of red-cell storage and complications after cardiac surgery. *New Engl J Med* 2008; 358:1229-39.
2. Zallen G, Offner PJ, Moore EE, Blackwell J, Ciesla DJ, Gabriel J, Denny C, Silliman CC. Age of transfused blood is an independent risk factor for postinjury multiple organ failure. *Am J Surg* 1999;178:570-2.
3. Mynster T, Nielsen HJ. Storage time of transfused blood and disease recurrence after colorectal cancer surgery. *Dis Colon Rectum* 2001;44:955-64.
4. Offner PJ, Moore EE, Biffi WL, Johnson JL, Silliman CC. Increased rate of infection associated with transfusion of old blood after severe injury. *Arch Surg* 2002;137:711-6.
5. Leal-Noval SR, Jara-López I, García-Garmendia JL, Marín-Niebla A, Herruzo-Avilés A, Camacho-Laraña P, Loscertales J. Influence of erythrocyte concentrate storage time on postsurgical morbidity in cardiac surgery patients. *Anesthesiology* 2003;98:815-22.
6. Purdy FR, Tweeddale MG, Merrick PM. Association of mortality with age of blood transfused in septic ICU patients. *Can J Anaesth* 1997;44:1256-61.
7. Lelubre C, Piagnerelli M, Vincent JL. Association between duration of storage of transfused red blood cells and morbidity and mortality in adult patients: myth or reality? *Transfusion* 2009;49:1384-94.
8. Fontaine MJ, Chung YT, Erhun F, Goodnough LT. Age of blood as a limitation for transfusion: potential impact on blood inventory and availability. *Transfusion* 2010;50: 2233-9.
9. Whitaker BI, Green J, King MR, Leiberger LL, Mathew SM, Schlump KS, Schreiber GB. The 2007 national blood

- collection and utilization survey report. Bethesda (MD): American Association of Blood Banks; 2008.
10. American Association of Blood Banks. National blood exchange. 2011. [cited 2011 Jun 10]. Available from: URL: <http://www.aabb.org/programs/nbe/Pages/default.aspx>
 11. Whitt W. Stochastic-process limits. New York: Springer-Verlag; 2002.
 12. Amin M, Fergusson D, Wilson K, Tinmouth A, Aziz A, Coyle D, Hebert P. The societal unit cost of allogenic red blood cells and red blood cell transfusion in Canada. *Transfusion* 2004;44:1479-86.
 13. Landau LD, Lifshitz EM. Statistical physics. Part 1, Vol. 5 of Course of theoretical physics. 3rd ed. Oxford: Butterworth-Heinemann; 1994.
 14. Murray JD. Mathematical biology. 2nd ed. Berlin: Springer; 1993.
 15. Cachon G, Terwiesch C. Matching supply with demand. 2nd ed. New York: McGraw-Hill Irwin; 2009.
 16. Dzik W. Fresh blood for everyone? Balancing availability and quality of stored RBCs. *Transfus Med* 2008;18:260-5.
 17. Steiner SE, Stowell C. Does red blood cell storage affect clinical outcome? When in doubt, do the experiment. *Transfusion* 2009;49:1286-90. 