Optimisation and control of the supply of blood bags in hemotherapic centres via Markov Decision Process with discounted arrival rate

Henrique L. F. Soares[a], Edilson F. Arruda[a,d,*], Laura Bahiense[b], Daniel Gartner[d], Luiz Amorim Filho[c]

[a]*Industrial Engineering Program, Transportation Egineering Program, Alberto Luiz Coimbra Institute - Graduate School and Research in Engineering, Federal University of Rio de Janeiro, Caixa Postal 68507, Rio de Janeiro, RJ 21941-972, Brazil*

[b]*Systems Engineering and Computer Science Program, Transportation Egineering Program, Alberto Luiz Coimbra Institute - Graduate School and Research in Engineering, Federal University of Rio de Janeiro, Caixa Postal 68511, Rio de Janeiro, RJ 21941-972, Brazil*

[c]*Hemorio - Arthur de Siqueira Cavalcanti State Institute of Hematology, Rua Frei Caneca 8, Centro, Rio de Janeiro RJ 20.211-030, Brazil*

[d]*School of Mathematics, Cardiff University, Senghennydd Rd, Cardiff CF24 4AG, UK.*

Abstract

Running a cost-effective human blood transfusion supply chain challenges decision makers in blood services world-wide. In this paper, we develop a Markov decision process with the objective of minimising the overall costs of internal and external collections, storing, producing and disposing of blood bags, whilst explicitly considering the probability that a donated blog bag will perish before demanded. The model finds an optimal policy to collect additional bags based on the number of bags in stock rather than using information about the age of the oldest item. Using data from the literature, we validate our model and carry out a case study based on data from a large blood supplier in South America. The study helped achieve an overall increase of 4.5% in blood donations in one year.

Keywords: OR in health services; Perishable inventory; Blood; Stochastic modeling, Blood management.

1. Introduction

Blood transfusion is a fundamental part of contemporary medicine (e.g., World Health Organization, 2010). It is estimated that 85 million blood transfusions are carried out annually across the globe, which translates to approximately 3 blood transfusions per second

---

*Corresponding author. Tel.: (+44) 029 2087-0617.

Email addresses: henriquesoares@poli.ufrj.br (Henrique L. F. Soares), efarruda@po.coppe.ufrj.br (Edilson F. Arruda), laura@cos.ufrj.br (Laura Bahiense), GartnerD@cardiff.ac.uk (Daniel Gartner), luizamorimfilho@gmail.com (Luiz Amorim Filho)

(Carson et al., 2012). In order to carry out such a large service volume, reliable supply chains must be designed (e.g., Hamdan and Diabat, 2019; Nagurney and Dutta, 2019b; Attari et al., 2019a).

This paper addresses two particular issues connected to the management of blood bank inventories. The first issue is the need to secure self sufficiency, which is important because donation is voluntary in most countries and is often insufficient (Rock et al., 2000; Azizoglu et al., 2018). One way to compensate for the lack of blood is by deploying transshipment strategies (Wang and Ma, 2015), which incur logistics costs. Another strategy is to stimulate donation (Özener et al., 2019), which may include sending external collection vehicles or houses (Gunpinar and Centeno, 2016). The latter strategy, which is believed to encourage voluntary spirit and decrease variability (Yew-Kwang et al., 2016), is studied in this paper as way to complement insufficient donation at the main blood bank facilities.

The second issue that we investigate is the inventory management of blood supplies. The decision maker is faced with the complex trade-off of maintaining supply while also avoiding large expiration rates due to limited shelf life. Let us assume, for the sake of argument, that the decision is to keep a very large inventory in the blood distribution centres. This gives rise to heavy collection, disposal and inventory costs, not least because of the limited shelf life of donated blood packs (Clay et al., 2018). Very low inventories, however, tend to produce increased probabilities of stock-outs, that should be avoided due to the sensitive nature of the supply.

The aim of our paper is to model this tradeoff, to provide a full supply of blood to the hospitals, while keeping losses to a minimum (Stanger et al., 2012b). And the flexibility available to the decision maker is the choice of the number of external collection teams at every decision period. We develop a Markov decision process (MDP) with the objective of minimising the overall costs of internal and external collections, storing, producing and disposing of blood bags. The model finds an optimal policy to collect additional bags based on the number of bags in stock rather than using information about the age of the oldest item.

The main contributions of the paper are as follows. Firstly, it bridges a gap in the literature by proposing a novel MDP model that explicitly considers perishability for arbitrary shelf lives. Indeed, inventory models that explicitly consider perishability are comparatively rare and, to the best of our knowledge, computationally tractable models are limited to shelf lives of at most three days, see for example (Osorio et al., 2015) and (Puranam et al., 2017). Another contribution of the model is that it produces an optimal policy in infinite horizon, thus extending the results in the literature that are generally limited to single-period models or myopic multi-period models with fixed time-horizons (Attari et al., 2019a). Finally, the proposed model is simple and easy to use, while also powerful enough to provide an optimal long-term policy whilst imposing no constraint on its structure.

This research is motivated by the case-study of HEMORIO, the main hemotherapy institute in the city of Rio de Janeiro, Brazil. To compensate for insufficient donations, the organisation complements its inventory with two types of mobile collections: an itinerant facility built within a bus, and mobile blood collection stations that can be assembled in external locations. Both types of collection are generally implemented in remote regions of

the city, away from the city centre. HEMORIO is currently undertaking economic viability studies considering the expansion of these kinds of external collections.

The remainder of this paper is structured as follows. Section 2 describes the studied problem. Section 3 provides an overview of related works on blood logistics. Section 4 presents the discrete-time stochastic model and discusses its particularities. Section 4 validates the model by applying it to a case study in the literature. The case study in HEMORIO is discussed in the numerical results of section 6, followed by concluding remarks in Section 7.

2. Problem Statement

This work studies the inventory management of blood supplies by centralised blood banks that may face insufficient donation in the main facilities. We assume that the blood banks have access to external collection infrastructures, which they may send to remote locations in order to complement their inventory and prevent shortages.

We assume that the organisation has access to at most $M \geq 1$ external collection teams and that these collection teams can be deployed simultaneously at each decision period. We assume that each external collection brings back a random number of additional blood bags, which is described by a Poisson variable whose parameter is inferred from historical data. In addition, each external collection introduces logistics costs to the organisation.

The organisation incurs a number of costs that are directly related to collecting donations. One such cost is the inventory holding cost, which involves keeping the items refrigerated and in good condition. The infrastructure available in the inventory facility effectively sets up an upper bound in the number of blood supplies that can be kept in storage at any given time. In addition, each blood bag has a fixed shelf life of L days, after which the blood expires and is no longer suitable for use. In that case, the institution incurs a disposal cost and the blood bag has to be discarded.

In order to collect a blood bag, the blood bank incurs a fixed *production* cost. We assume that this cost is the same for internal and external collections. However, each external collection implies an additional cost per collected blood bag, which includes transportation, setup and staff costs. Finally, a shortage (depletion) cost is also assigned for low inventory states in order to avoid blood shortage. This cost decreases as the inventory increases, for the risk of shortages is inversely proportional to the number of blood bags in inventory.

The decision maker is then faced with a choice on the number of external collection units to deploy at each decision period. This choice depends on the current inventory and is guided by the inventory, shortage, disposal and external collection costs explained above. The rationale is to find the right balance between shortage risk and expiration/disposal risk, while accounting for the overall cost incurred by the system. The detailed model presented in section 4 makes use of a Markov decision process formulation to devise a long-term collection strategy that prescribes the number of external collection teams to be deployed for each possible inventory level.

3. Related Work

Recent literature reviews on blood supply chains can be found in (Pirabán et al., 2019), (Mansur et al., 2018), (Osorio et al., 2015) and (Beliën and Forcé, 2012). And a substantial number of articles study inventory management of blood components and perishable commodities. In general, the literature can be divided into process-based and quantitative models.

3.1. Process-based Models

Prastacos (1984) proposed management strategies that reach 100% of availability while also maintaining low levels of stock. Bedi et al. (2016) implemented a careful study of wastage in the blood supply chain of a hospital and proposed simple management rules that significantly improved resource utilisation. They found that reducing the blood preparation time from 72 to 48 hours, while also promoting a division of the storage into old and new blood bags, led to considerable waste prevention in their case study.

In another line of research, Clay et al. (2018) reported a whipping effect between demand and supply due to panic orders, while Stanger et al. (2012b) identified key procedures to improve processes, such as maintaining a target stock level and an order pattern. Another consensual issue (e.g., Stanger et al., 2012b,a) is the preference for small orders to the detriment of larger ones, with a view to inducing economies of scale.

3.2. Quantitative Models

Recently, quantitative models have been used to study aspects of the blood supply chain in disaster relief situations which make use of multi-objective programming (Samani et al., 2018), robust optimisation (Jabbarzadeh et al., 2014) and Markov chains (Hosseinifard et al., 2019). Other supply chain approaches include mathematical optimisation models for finding appropriate delivery routes for blood supplies (Ezugwu et al., 2019; Kazemi et al., 2017). In contrast, Nagurney and Dutta (2019a,b) focus on the competition among blood service organisations, whereas Rajendran and Ravindran (2019) investigate parametric ordering policies without considering perishability. Finally, Hamdan and Diabat (2019) employ a two-stage stochastic programming approach to produce location and scheduling decisions.

There is an ample variety of modelling choices for blood inventory problems. Among them, one can refer to linear regression models (e.g., Schreiber et al., 2005; Godin et al., 2007), Markov models (e.g., Pegels and Jelmert, 1970; Brodheim et al., 1975; Abubakar et al., 2014), discrete event simulation (e.g., Katsaliaki and Brailsford, 2016; Rytilä and Spens, 2006) and mathematical programming (e.g., Hemmelmayr et al., 2009; Dillon et al., 2017).

Among the main articles to be highlighted, Keilson and Seidmann (1990) compared FIFO and LIFO disciplines for perishables and addressed the trade-off between managerial and demand interests. Dehghani and Abbasi (2018) developed a heuristic solution using partial differential equations for transshipment between hospitals and blood banks and verified a benefit to the supply chain management and a reduced cost of mean age of blood transfusion.

Finally, Hosseinifard and Abbasi (2018) showed that the centralisation of stock at the end of the supply chain led to a reduction of up to 40% in the inventory costs of hospitals.

A common quantitative approach is to model blood inventory as a queue, possibly with controlled input and/or output rates. Typically, it is the level of control over these rates that defines inventory policies. While this modelling may not work for all applications due to some assumptions on supply and demand processes (Dillon et al., 2017; Abbasi et al., 2017), queuing models are general enough to accommodate different distributions. Moreover, common assumptions for the simplest models, such as Poisson input and output processes, are appropriate for many healthcare applications (Angelo et al., 2017; Goldwasser et al., 2016; Pearson et al., 2012). Additionally, viewing the system as a queue promotes analytic tractability and introduces flexibility in the design of sub-optimal policies.

Depending on the application, studies propose varying the input rate, the output rate, or both. Policies promoting a variation of the output rate were advocated in a series of works (e.g., Jo and Stidham, 1983; George and Harrison, 2001; Chan et al., 2011). George and Harrison (2001) studied analytical properties of the optimal policy of a controlled M/M/1 queue with fixed arrival rate and variable service rate. An M/G/1-based model was employed by Jo and Stidham (1983) to find optimal service rates for a model with a number of sub-stages whose duration was modelled by an Erlang distribution. A different focus was pursued by Chan et al. (2011), who searched for operational parameters that promote an adequate use of *speedup*, a myopic control used by medical teams to promote acceleration of service in ICU units.

In another context, Crabill (1974) utilised a semi-Markov based queuing model with control of both input and output rates to search for optimal machine maintenance policies. Both rates were also allowed to vary in the model proposed by Nahmias et al. (2004) for perishable items, whereas Perry and Posner (1990) propose one or two switch-over levels by varying demand and/or production rates.

There are many ways to vary the input rate. A dynamic pricing policy, for example, may stimulate or discourage demand (e.g., Low, 1974; Weber, 2015; Paschalidis and Tsitsiklis, 2000; Yoon and Lewis, 2004). Advertising can also promote effective demand control, which comes at a cost and is applied to boost the demand for a prescribed period (e.g., Weber, 2015). Regardless of the mechanism employed, varying demand rates gives rise to optimisation problems whereby optimal demand-switch levels should depend on the current inventory of the system. We argue that such problems are appealing for blood inventory management, especially when voluntary donations are not sufficient to consistently meet the demand.

The perishable nature of the supply is instrumental in quantitative models. One can deal with it implicitly, such as in airline overbooking studies (e.g., Subramanian et al., 1999) and contractual and logistical selective orders (Chao, 2013); or explicitly, by means of age-based inventory control (Graves, 1982; Perry and Posner, 1990). Using the latter approach, Wang and Ma (2015) projected an age-based transshipment model studying stock features by simulation. The rationale is to manage the system with incomplete knowledge of the age of the items in storage. Graves (1982) modelled the system by keeping count solely of the oldest item in stock, which can be seen as summarising the system's behaviour. Kaspi and

Perry (1983), on the other hand, proposed and explored the equivalence between shelf-life expiration and stock-outs and busy and idle periods of a corresponding queuing model with impatient customers, the Inventory Systems of Perishable Commodities (ISPC). That led to switch-over policies with one or two policy-change points. This model is relevant to the studied problem because it was applied, among several case studies, to optimise blood banks in the Netherlands (Gunpinar and Centeno, 2015) and Canada (Kopach et al., 2008). These works employed the idea of finding switch-over points that define an optimal policy that promotes cost minimisation.

Despite their great relevance, the models above generally present significant analytical complexity. In addition, the model in (Perry and Posner, 1990) can only deal with one or two switch-over levels, which limits the generalisation to more complex scenarios with a large number of input or output rates to choose from, each with its own characteristics in terms of costs incurred. In contrast, blood inventory optimisation models typically consider two dichotomous situations: normal and emergency (e.g., Kopach et al., 2008; Hosseinifard et al., 2019). Closed-form solutions are usually obtained with a view to maintaining the blood supply at an operationally satisfactory level. However, the multiplicity of endogenous factors makes that approach limited. Indeed, an analytical treatment of the expressions derived from models with two or more switch-overs is close to impracticable. More recently, stochastic optimisation models have also been applied that consider fixed shelf lives (Puranam et al., 2017; Chen et al., 2019), but the former work only considers shelf lives of at most two days, whereas the latter gives rise to a complex model that becomes rapidly intractable for all but very small shelf lives and hence seeks approximate solutions using heuristic procedures. Furthermore, both models generate myopic open-loop controls for fixed time horizons. These are reactive models that should be re-evaluated at each realisation of a random variable.

Therefore, a novel contribution of this paper with respect to analytic models is a simple yet general blood inventory model that remains tractable for multiple switch-over points. Another contribution is that, in contrast to stochastic optimisation models, it produces a long-term stationary optimal policy. In addition, the proposed model explicitly considers perishability and remains computationally tractable regardless of the shelf life. Such a powerful result is made possible because perishability is addressed by differentiating between useful and potentially expired blood bags, instead of keeping track of the expiration dates of all items in stock, which would render the model intractable (e.g., Chen et al., 2019). The approach, which was inspired by the analogy between impatient clients on a M/M/1 queue and the blood bag expiration in (Kaspi and Perry, 1983), discounts the probability that an incoming blood bag will expire before being distributed, as detailed in section 4. Finally, an important by-product of the model is that it is easier to use and implement because it proposes policy changes based on the number of bags in stock rather than the age of the oldest item. It is worth pointing out that the latter can be much more difficult to measure in practice.

4. Mathematical Formulation

Generally, blood collection centres separate their blood donations into their constituent components: platelets (PLT), plasma (FFP) and red blood cells (RBC) (Williamson and Cardigan, 2005). The latter often pose more problems in terms of management due to their significantly larger demand (Murphy and McSweeney, 2009). Therefore, most of the literature focuses on red blood cells. This paper follows that trend and studies optimal inventory policies for RBCs, hereby referred to as blood bags.

In this paper, we model blood inventory as a Markov decision process - MDP. MDPs are designed for sequential decision making under uncertainty (Puterman, 2014), hence their frequent use in medical applications (e.g., Andersen et al., 2019; Arruda et al., 2019; Steimle and Denton, 2017; Bennett and Hauser, 2013; Alagoz et al., 2010) and in blood inventory problems (e.g., Attari et al., 2019b; Hosseinifard et al., 2019). The use of an MDP ensures more flexibility as it allows an optimal decision for each available state, while also ensuring optimality in the long run.

In this article, demand and supply are modelled as Poisson processes. Such a modelling is adequate because both demand and supply can be seen as rare, independent events, in a very large population, which characterise Poisson processes, see also (e.g., Baron et al., 2010; Kim et al., 2015; Abbasi and Hosseinifard, 2014). For a better understanding of the parameters and equations, the adopted notations are explained in Table 1.

Consider a Markov decision process with state space $S = \{0, 1, 2, \ldots\}$ whose controlled dynamics are described by process $\{X_k\}$, $k \geq 0$. Each state $i \in S$ represents a possible number of useful blood bags available in stock. Here, the word *useful* confers a particularity to the model that will be explained later. At each state $i \in S$, an action is selected from the set of feasible actions for that state, denoted by A_i. The set of admissible actions is defined as $A = \cup_{i \in S} A_i$. For the present study, each action $a \in A_i$ prescribes a given number of mobile collection teams and itinerant buses that should be sent to complement the blood donation of the main facility. We assume that each external collection yields a random number of donated bags, which is given by a Poisson random variable whose parameter is inferred from data obtained in previous collections.

Suppose that $X_k = i$ at some period $k \geq 0$. The decision maker must choose an action $a \in A_i$ to be applied. A one-step transition results from this action, and the system jumps to $X_{k+1} = j \in S$ with probability $p_{ij}^a \geq 0$. To properly define the system, we must have $\sum_{j \in S} p_{ij}^a = 1$, $\forall i \in S$ and $a \in A_i$. Upon entering a state $i \in S$ and applying action $a \in A_i$, the system incurs a one-step cost $c(i, a)$, where $c : S \times A \rightarrow \mathbb{R}_+$ is a convex positive cost function. Figure 1 depicts the transitions for a system with 3 control actions.

Assume that each time process $\{X_k\}$, $k \geq 0$ visits state $i \in S$, the same action $a \in A_i$ is selected. We can define a stationary control policy $\pi : S \rightarrow A$ that prescribes a single action to be applied at each state in S; we also let Π be the set of all feasible stationary control

Table 1: Nomenclature.

| Parameter	Meaning
S	Set of MDP state space, *useful* blood bags available in stock
A_i	Set of possible actions for each state i
X_k	Markov process
p_{ij}^a	Probability of going from state i to state j by action a
Π	Set of all feasible stationary control policies
π	Each stationary control policy
η^π	Long-term average cost associated with policy π
η^*	Optimal stationary average cost
L	Shelf-life of a blood bag
μ	Average demand rate
λ^a	Average arrival rate depending on the action a
$F_T(t,i)$	Probability that a i^{st} item will be consumed at a time t
$c(i,a)$	Cost of stay in state i and take action a
$c_f(i)$	Lack cost
c_s	Storage cost
c_c	Collection cost
c_p	Production cost
e_c	External costs
c_d	Disposal costs
c_1	Input constant of lack cost
d_1	Decay constant of lack cost
c_2	Unit constant of disposal cost
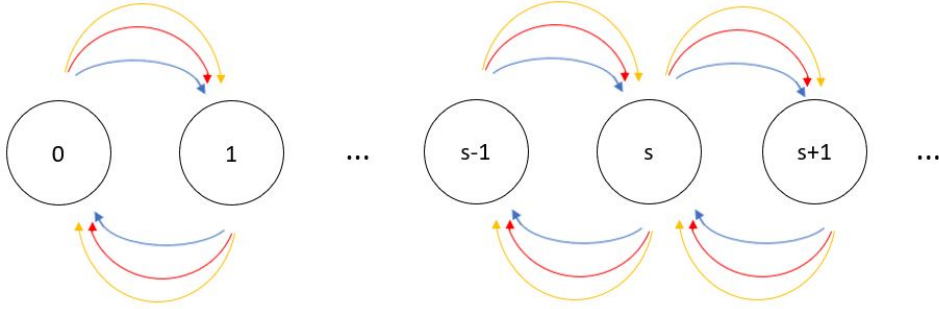
Figure 1: Transition graph for an MDP with 3 control actions. Yellow corresponds to having no external collection, red to sending one collection team and blue to sending two teams.

policies. For each policy $\pi \in \Pi$, there exists an associated long-term average cost

$$\eta^\pi = \lim_{N \to \infty} E_\pi \left(\frac{1}{N} \sum_{k=0}^{N} c(X_k, \pi(X_k)) \right). \tag{1}$$

We seek an optimal stationary policy π^* that minimises the long-term average cost, i.e. a policy that satisfies:

$$\eta^* = \eta^{\pi^*} \leq \eta^{\pi}, \ \forall \pi \in \Pi. \tag{2}$$

It is well-known that, under mild conditions, there exists an optimal stationary policy π^* which solves (2), for example see (Puterman, 2014).

4.1. Iterative Solution Procedure

To solve (2), one can apply the classical relative value iteration algorithm (Puterman, 2014). Let $V : S \to \mathbb{R}$ be a real-valued function and let $\mathcal{V}$ be the space of real-valued functions in S. The relative value iteration algorithm starts with an arbitrary initial solution $V_0 \in \mathcal{V}$ and iteratively updates it up to convergence. The procedure is detailed in Algorithm 1, where $\|x\|_s$ is the span semi-norm, defined as the difference between the maximum and the minimum values in vector x. Algorithm 1 has guaranteed convergence to the optimal long-term average cost η^* and always finds an optimal stationary policy π^* which solves (2) (e.g., Puterman, 2014).

Algorithm 1: Relative value iteration

Input: An arbitrary initial solution $V_0 \in \mathcal{V}$, a pivot state $i^* \in S$, and an arbitrary tolerance *tol*

Output: Optimal solution V^*, optimal long-term average cost η^*, and optimal stationary policy π^*

1   $k \leftarrow 1$;
2   $V_k \leftarrow \infty$;
3   **while** $\|V_k - V_{k-1}\|_s \geq tol$ **do**
4   **for** *each $i \in S$* **do**
5

$$TV_k(i) := \min_{a \in A_i} \left\{ c(i,a) + \sum_{j \in S} p_{ij}^a V_k(j) \right\} \tag{3}$$

6   **for** *each $i \in S$* **do**
7

$$V_{k+1}(i) := TV_k(i) - TV_k(i^*) \tag{4}$$

8     $k \leftarrow k + 1$;
9   $V^* \leftarrow V_k$;
10   $\eta^* \leftarrow TV^*(i^*)$;
11   $\pi^*(i) \leftarrow \arg\min_{a \in A_i} [TV^*(i)], \forall i \in S$;
12   **return** V^*, η^*, π^*.

4.2. Details of the blood inventory formulation

Assume that the blood demand follows a Poisson process with constant rate $\mu > 0$. The blood arrival process, which is a function of the donations, is a Poisson process with varying rate, depending on the control actions taken at each step. At period $k \geq 0$, with $X_k = i$ denoting the number of *useful* blood packs in storage, the blood arrival rate λ^a is a function of the control action $a \in A_i$ selected. As such, it depends on the number of itinerant buses and external collection teams prescribed by a.

To account for perishability, for each blood bag collected the model decides, upon arrival, whether this bag will actually be added to the inventory. Only bags that are expected to be used before their shelf life expires are added to the stock, hence the state of the system accounts for the number of *useful bags*, i.e. those that are not expected to reach their expiration date.

Inspired by the model in (Kaspi and Perry, 1983), the probability of expiration of a newly arrived bag works as a discount on the arrival rate. The magnitude of the discount depends, of course, on the current number of items in stock not expected to expire, i.e. the current state of the system. A newly arrived blood bag will expire if the demand for blood during its shelf-life does not surpass the current stock, as the current stock is obviously comprised of older items. Following Kopach et al. (2008), we set the shelf-life of blood bags to $L = 42$ days. Assume that, at a given period $k \geq 0$, $X_k = i$. Since the demand is a Poisson process with rate μ, the time to consume the i items in stock follows an Erlang distribution with parameters i and μ. Hence, letting T denote the time the system needs to deplete the current stock of i items, the expressions for its probability density and cumulative distribution are, respectively (e.g., Shiryayev, 1984):

$$f_T(t, i) = \frac{\mu^i t^{i-1} e^{-\mu t}}{(i-1)!},$$

and

$$F_T(t, i) = P(T < t) = 1 - \sum_{n=0}^{i-1} e^{-\mu x} \frac{(\mu x)^n}{n!}. \tag{5}$$

The probability that a newly arrived item will be consumed is $F_T(L, i)$, where L is the shelf life and i is the current state of the inventory system.

Now we are ready to quantify the transition probabilities for the proposed model. Suppose $X_k = i$, then upon the next event (arrival or departure of a blood bag from storage), the system jumps to $X_{k+1} = j$ according to the following probabilities:

$$p_{ij}^a = \begin{cases} \dfrac{\lambda^a F_T(L, i)}{\lambda^a F_T(L, i) + \mu}, & \text{if } j = i + 1; \\[3ex] \dfrac{\mu}{\lambda^a F_T(L, i) + \mu}, & \text{if } j = i \text{ - } 1; \\[3ex] 0, & \text{otherwise}. \end{cases} \tag{6}$$

The first expression in Eq. (6) accounts for the probability that the next event is an arrival. The factor $F_T(L, i)$ is the perishability discount. The second expression conveys the probability that the next event is a demand, in which case the non-expiring inventory is decremented by one unit.

Finally, a crucial feature of the formulation is the cost function. Recall that, at a given state $i \in S$ and under control action $a \in A_i$, the system incurs an instantaneous cost $c(i, a)$. Generically, one can think of a general cost function that encompasses production, inventory, deficit and disposal costs. In addition, we also need to identify the difference in costs between internal and external collection. Typically, there are many inherent characteristics of each application which cannot be easily generalised. However, some aspects are common and should be considered in all studies.

Deficit costs should be related to the probability of shortage at any state of the system, and therefore, should decrease as stocks increase. In this paper we propose an exponentially decreasing function for the deficit costs. As the cost is associated to positive stock levels, it is not actually a deficit cost, but a cost to account for the risk of future shortages when the stock is low. This cost is given by:

$$c_f(i) = c_1 e^{\frac{-i}{d_1}}, \tag{7}$$

where c_1 and d_1 are positive scalars.

Inventory costs tend to be increasing functions of the stock level. The shape of the function, however, is expected to change from blood centre to blood centre. In this study, we propose a generic piecewise linear function that considers economies of scale as the inventory increases. A similar approach was proposed by Besanko and Braeutigam (2005). The inventory cost is of the following form:

$$c_s(i) = s_1 i \mathbb{1}_{\{i \leq S_1\}} + s_2 i \mathbb{1}_{\{S_1 < i \leq S_2\}} + s_3 i \mathbb{1}_{\{S_2 < i \leq S_3\}},$$

where $\mathbb{1}_A$ denotes the indicator function of statement A, which equals one whenever A holds true and is nil otherwise. Additionally, s_1, s_2, s_3 and S_1, S_2, S_3 are modelling parameters. This function splits the inventory cost into three linear sub-functions. That is a modelling choice that can be seamlessly altered depending on the case study.

Production costs, i.e. collection costs, depend on the action. The fixed cost related to the operation of the blood centre is left out of the optimisation process, since it does not alter the ranking of the policies. As for the variable collection costs, we assume that collecting blood in the main facility is as costly as collecting outside. However, there is also a fixed cost per external collection team, namely the transport and fixed personnel costs. The collection cost is:

$$c_c(a) = c_p(a) + e_c(a),$$

where $c_p(a)$ is the expected cost of all donations under action a and $e_c(a)$ is the expected additional cost of the external collections prescribed by a. The latter also depends on the mix of itinerant buses and mobile collection stations deployed.

Disposal costs are directly associated with the probability that a blood bag expires before consumption. In our model, this is described by $(1 - F_T(L, i))$. In the equation below, c_2 is a unitary cost of discarding:

$$c_d(i) = c_2(1 - F_T(L, i)).$$

The total cost of visiting state $i \in S$ and applying action $a \in A_i$ is then given by:

$$c(i, a) = c_f(i) + c_s(i) + c_d(i) + c_c(a). \tag{8}$$

Remark 1. *An appealing trait of the proposed MDP model is that it works with any type of cost function. Hence, the formulation in (8) can be seamlessly altered depending on the application. The terms in (8) are included here as a suggestion based on a real-world application.*

Remark 2. *Low (1974) demonstrated that, for a queue with various rates, there exists an optimal stationary monotonous policy. In his work, the optimal price to advertise was a non-decreasing function of the number of customers in the system. In our model, we expect an analogous behaviour because the fundamental idea is that the more RBC's in stock, the less the need for complementing via external collection. In addition, since donation complements are discrete events (number of external collections), we expect the arrival process to be a non-increasing step function in which the variable $\Delta\lambda$ represents increases in demand. Figure 2 illustrates the basic form of the arrival rates, which, in fact is obtained, as will be seen in the numerical results.*
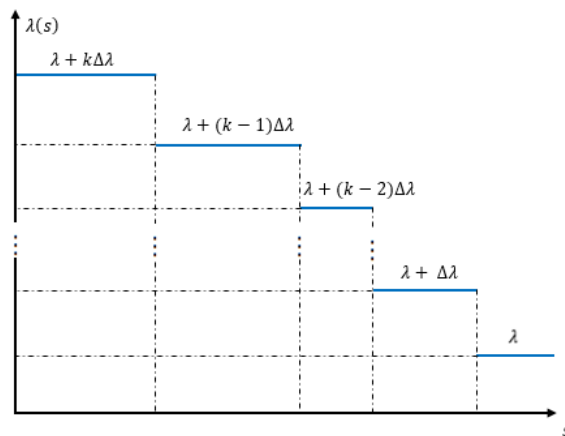

Figure 2: Evolution of λ and its decreases $\Delta\lambda$. The more RBC's in stock, the less the need for complementing via external collection. The figure does not consider the discount for perishability.

The next section features an example from the literature, with a view to validating the proposed model.

5. Model validation

To validate our approach, we make use of the results in (Kopach et al., 2008). For the sake of simplicity, the model is applied only to blood type O^+, but the algorithm could be separately applied to all blood types, see Remark 3. As for the parameters, we start with the input and output processes. We split the blood arrival rate in two parts: a constant part related to the donation in the main centre, and a variable part that depends on whether the single available external collection team is deployed. We assume that the team is responsible for an increase of $\Delta\lambda$ in the arrival rate. Table 2 presents the input and output rates, as well as the shelf life L. All rates are given in blood packs per day.

Table 2: Input/output parameters for the validating example - adapted from (Kopach et al., 2008).

Constant input rate (λ)	Increase via external team ($\Delta\lambda$)	Demand rate (μ)	Shelf life (days) (L)
154.7	42.3	123.7	42

The cost function is based on the structure in (8), and is defined as:

$$c(i,a) = 43,958e^{-\frac{1}{800}i} + 0.56i + \frac{134.49\lambda^a}{\lambda^a + \mu}, \tag{9}$$

where $A_i = \{0, 1\}$, with $a = 1$ if the external team is deployed and $a = 0$ otherwise. Then, the parameters in Algorithm 1 are those in Table 2, $\lambda^1 = \lambda + \Delta\lambda$ and $\lambda^0 = \lambda$. To solve the problem, it suffices to run Algorithm 1 with $c(i,a)$ obtained from Eq. (9), $F(L,i)$ obtained from Eq. (5) and the transition probabilities derived from Eq. (6).

Remark 3. *We decided to illustrate our model only with type O+ blood because it is the crucial type for a safe blood inventory. If 100% of blood inventory is of type O+, then we can transfuse everyone with this type of blood without any further safety concerns. In contrast, if we had a large inventory of A, B and AB blood types, but few units of blood type O, then there would be a great risk of blood shortage.*

In Rio de Janeiro, 45% of the population are O+ and 5% O negative (Luiz Amorim, personal communication). In the emergency rooms (ER), physicians do not have time to have the blood group results before transfusion – and they have to transfuse blood type O. This is not a problem, because A, B, and AB patients can be transfused with blood type O – but patients with blood type O can only receive blood type O.

Ideally, in ER as well as in other medical situations, the use of O negative blood would be the best option. However, only 5% of the Brazilian population is O negative; that is why this type of blood is reserved for females with childbearing potential, in order to avoid alloimmunization - development of anti-D antibody after receiving an O+ blood - bringing the risk of Hemolytic Disease of Fœtus and Neonates. A recent paper (Yazer et al., 2019) estimated that it is a very minor risk (0.3%), which in their opinion, warrants the transfusion of O+ blood even for the female population.

January 3, 2020

It is worth pointing out that most of the data, such as arrival and demand rate, increase by switch-over constant and production cost came directly from (Kopach et al., 2008). Other parameters - which are specific to our formulation - were properly adapted. Deficit cost parameters came from the penalty for operating at a lower demand level and costs of switching between donation profiles were assumed to be zero as in the reference.

The reference aims to find a switch-over point x^*, at which to alternate between two demand rates. The options concern two situations: normal and emergency, when the service level needs to be accelerated. In the reference, the output decision for blood O^+ is to switch when the age of the oldest item is greater than 0.21 days, therefore operating 27% of time in lower (normal) demand. Our output suggests changing policies when the number of RBC's in stock is 31. If we divide our switch-over point by the rate of increased demand, that is, $\frac{x^*}{\lambda+\Delta\lambda}$, we estimate an age of 0.16 days for the oldest item in stock, a result that is similar. Figure 3 shows the cumulative probability curve for the Markov chain associated to the optimal policy. One can see that the stock level stays between 24 and 40 items nearly 100% of the time, demonstrating the efficiency of the control policy.
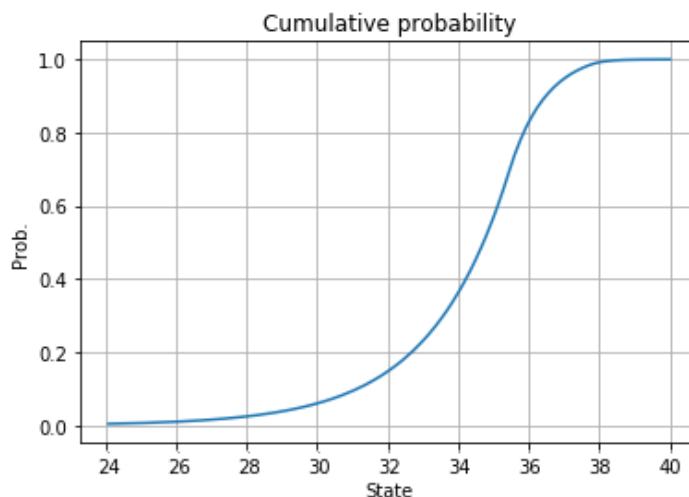
Figure 3: Behaviour of the accumulated probability of states of the stationary Markov chain around the switch-over point for the example of Kopach et al. (2008).

To perform a direct comparison with the results in (Kopach et al., 2008), we evaluate the control parameters introduced there and compare our results to their outcome, as shown in Table 3.

Table 3: Output parameters - proposed model versus (Kopach et al., 2008).

Parameters	(Kopach et al., 2008)	Our Model
E(Inventory)	34.45	34.56
E(shortage)	0.013	0.018
Cost/day	$23.835	$24.422

Table 3 reveals that both models yield very similar results. The cost per day of $24.422 is 2.46% higher than the original results, probably because the expected value of RBC's stocked is also slightly larger. At this point, it is worth pointing out that the expected number of items in inventory was derived in the original model analytically by a function $f_s(\omega)$, which describes the remaining shelf life of the oldest item, directly inspired by the results in (Perry and Posner, 1990). Therefore, it is natural that different assumptions reflect slight changes in the final results. In any case, it is important to remember that the methodology in this article is easier to implement, as it looks at the number of stored bags rather than the age of the oldest item. In addition, it is not constrained to a single switch-over level.

6. Numerical example

This section applies the proposed approach to our case study in HEMORIO. Section 6.1 features a brief analysis of the donations and external collections, starting in 2009. The numerical results for some selected scenarios are presented in Section 6.2. The data presented in this article comes from HEMORIO's information system and is detailed in Appendix A.

6.1. HEMORIO's Data

To complement the blood donations in the main facility, HEMORIO started to use itinerant buses about 15 years ago. Demountable collection structures were later incorporated due to logistic reasons. As the initiatives gained importance, inventory management became increasingly complex and assumed a more prominent role.

Figure 4 conveys the average number of blood bags collected per month due to both *internal collection* (IC) in the main facility and *external collection* (EC), by means of itinerant buses and demountable structures. Whereas the internal collection presents signs of stagnation and slight decline, the external collection is steadily increasing on a yearly basis.

Figure 4: Evolution of HEMORIO's internal (IC) and external (IE) collections. In 2017, the external collection accounted for more than 100,000 bags.

The evolution of the share of the external collection in the total number of collected bags is depicted in figure 5, starting in January 2009. The steady increase over the years suggests that the external collection tends to play an increasingly important role in the upcoming years.
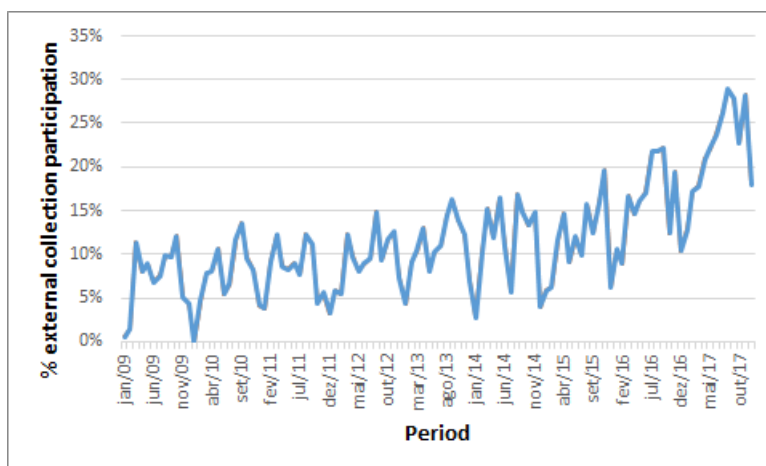

Figure 5: Evolution of the percentage of blood bags coming from the external collection.

In summary, the number of blood bags collected per month, either internally or externally, varies from 6,000 to 9,000 in the studied period. As figures 4 and 5 suggest, external collection is becoming more representative as time elapses. Not many years ago, it represented less than 10% of the total incoming blood, whereas in 2017 it reached around 22.7%. To convey a better picture of the reality, it is worth mentioning that the monthly average of blood bags collected in 2017 was around 8,602.9.

Another important analysis can be made regarding the type of external collection employed, either by demountable structure (type 1) or by bus (type 2). A thorough look at the data suggests that type 2 collections are more or less stable, whereas type 1 collections presented significant growth, especially in 2017, see figure 6 (b).

From 2009 to 2017, 1512 external collections took place, 1306 (86.4%) of type 1 and 206 (13.6%) of type 2. Each type 1 collection gathered 61.30 bags in average. The number of collections remained stable until 2015; thereafter there were two consecutive expressive growths, of 35.2% in 2016 and 92.3% in 2017. Figures 6 (a) and (b) summarise these results.

The performance in costs and the operability of type 1 collections are superior to those of type 2. However, for now, instead of comparing both methodologies, a weighted average is used to simulate the institution's current policy. This results in the following difference in the input rate:

$$\Delta\lambda \;=\; 61.30 \times 86.4\% \;+\; 46.43 \times 13.6\% \;=\; 59.3 \;.$$

In the next subsection, we will show the results of the described model applied to the case study and convey what would happen if HEMORIO decided to use only complementation by type 1 collection, which proved to be more efficient.

January 3, 2020
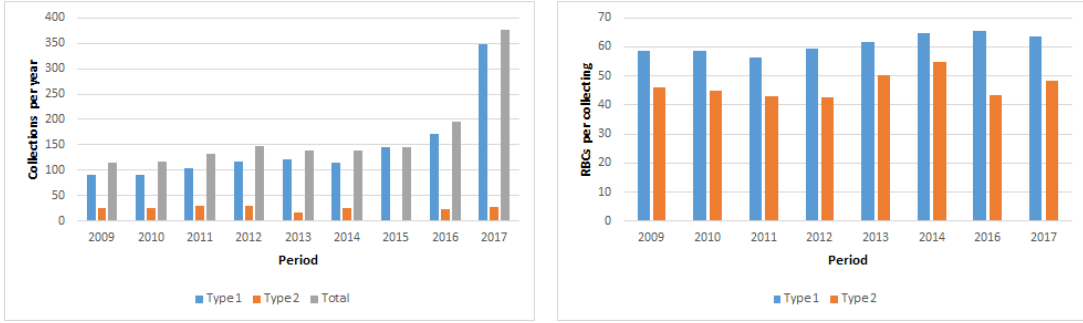
(a) Number of collections per year per type of collection. Type 1 (demountable structure) is predominant throughout the entire period.

(b) Number of bags collected per trip per type of collection. Type 1 is most effective in all years.

Figure 6: Number of collections per year per type of collection.

6.2. Numerical results

The model introduced in Section 4 was used in order to find an optimal policy for HEMORIO. The parameters we use in Table 4 were obtained from raw data and considering the characteristics of the organisation. Over the 9 years of data, an average of 206.1 bags were donated per day at the main facility. Input rate increased via external collections, each of which yield an average of 59.3 bags per day. The shelf life is 42 days and the average daily demand is 320 bags per day. The parameters, which serve as input to Algorithm 1, are summarised in table 4.

Table 4: Input/output parameters for HEMORIO case study.

Constant input rate (λ)	Increase via external team ($\Delta\lambda$)	Demand rate (μ)	Shelf life (days) (L)
206.1	59.3	320.0	42

The cost function is based on the structure in (8), and is defined as:

$$c(i,a) = 400,000\ e^{-\frac{1}{2000}\ i}\ +\ 73.2\ i\ \mathbb{1}_{\{i \leq 3333\}}\ +\ 75.8\ i\ \mathbb{1}_{\{3333 < i \leq 6666\}} \atop +\ 73.2\ i\ \mathbb{1}_{\{6666 < i \leq 10000\}}\ +\ 66.01 + 5.84\ k\ \mathbb{1}_{\{a=k\}}\ . \tag{10}$$

Daily input and output rates came from HEMORIO's data analysis. Cost function parameters were set empirically to simulate the behaviour illustrated in figure 7. The large costs for low inventory levels are due to the large expenses related to the deficit component, which fades as the inventory grows. As the inventory increases, the production-inventory component becomes dominant and the overall costs begin to increase. Observe that the increase is not exactly linear because of the economies of scale, simulated by a piecewise linear function, as detailed in section 4.2. The minimum cost is around inventory level $s^* = 2,100$ - where the green and red curves intersect - and represents a guarantee of approximately 7 days of blood self-sufficiency ($320 \times 7 = 2,240$), on average, for the institution. Production

January 3, 2020

and external collection costs were derived from bottom-up estimation (Snyder, 2013), considering factors such as human resources, materials and maintenance, which are undisclosed. The value of R$66.01 (R$ stands for Brazilian Reais) for internal collection represents the expected operating cost of collecting one bag, while R5.84k$ represents the additional cost of deploying k external collection teams.
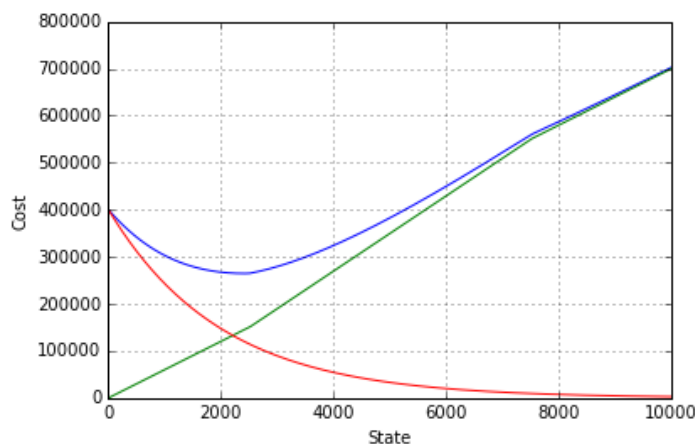

Figure 7: Evolution of costs through the states. Blue curve represents total costs, sum of production-inventory costs (green) and deficit costs (red).

In our example, there are 4 available control actions $A_i = \{0, 1, 2, 3\}$, namely to send 0, 1, 2, or 3 external collection teams. To solve the problem, it suffices to run Algorithm 1 with $c(i, a)$ obtained from Eq. (10), $F(L, i)$ obtained from Eq. (5) and the transition probabilities derived from Eq. (6), with $\lambda^a = 206.1 + 59.3a$. Recall that the parameters needed for these evaluations appear in table 4.

Figures 8 (a) and (b) show that the optimal policy prescribes three external collection teams up to state 2,141; when the stock level is between 2,142 and 2,170, two collection teams are assigned and a single team should be sent whenever the stock level is in the interval $[2, 171, 2, 220]$. Finally, no external collection is needed when the stock level surpasses 2,220 blood packs. The optimal long-term average cost is R$264,476.73 and corresponds to keeping, on average, 2,197 blood bags in stock. Naturally, the policy can be adapted by the health institute, since the schedule of the external collection teams may need to be set up in advance. Nonetheless, it provides a useful guidance in the design of adequate sub-optimal policies, see section 6.4.

Concerning the optimal Markov chain long-term probabilities, we notice that the cumulative probability tends to 100% very quickly. States that are distant from the policy change points have probability almost nil, as illustrated in figures 9 (a) and (b). This suggests that the approach is very powerful in terms of regulating stock levels, keeping the long-term variation very small. Naturally, the policy change points depend upon the ratio between external collection costs and other costs, so that different data could generate a rather distinct configuration of policy change points.

January 3, 2020

(a) Evolution of λ and its three decreases.

(b) Evolution of λ and its three decreases (Zoom).

Figure 8: Evolution of λ.

(a) Steady state cumulative distributions.

(b) Steady state cumulative distribution around policy change points.
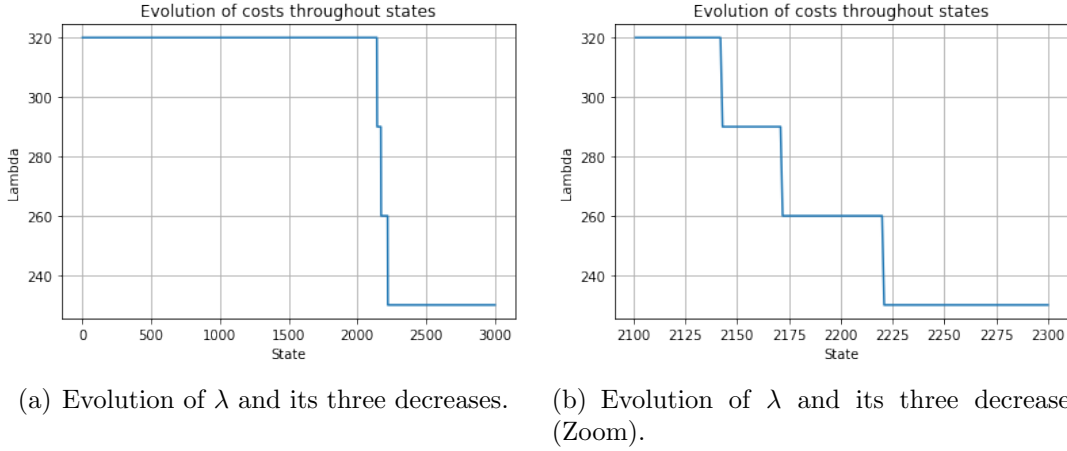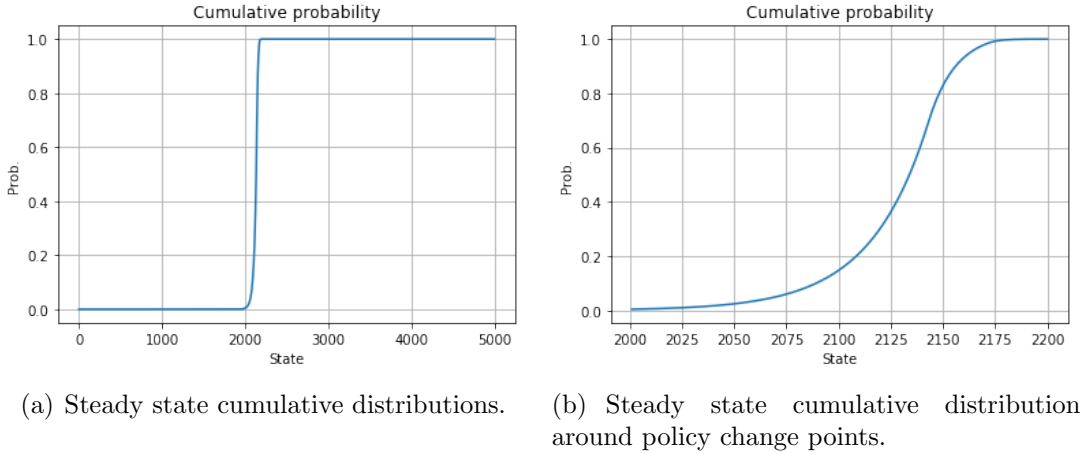
Figure 9: Steady state cumulative distributions.

6.3. Sensitivity analysis

Two important analyses are carried out in order to propose alternative solutions for the case study. Firstly, we investigate the effect on the optimal policy of a different adjustment of parameters in the cost function, setting up the lowest cost for 3 days of demand (instead of 7). Secondly, we evaluate the effect of permitting different combinations of type 1 and type 2 collections.

To establish a favourable setup for 3 days of demand ($3 \times 320 = 960$), deficit parameters should provide a dampened response. As a consequence, the optimal solution will keep lower inventory levels, given that the inventory and production costs will have a larger relative weight. Proposing an empirical change in the decay constant from 2000 to $2000 \times \frac{3}{7}$, we obtain a new cost curve, illustrated in figure 10, and defined in Eq. (11):

$$c(i, a) = 400,000 \ e^{-\frac{1}{857.14} \ i} \ + \ 73.2 \ i \ \mathbb{1}_{\{i \leq 3333\}} \ + \ 75.8 \ i \ \mathbb{1}_{\{3333 < i \leq 6666\}} \tag{11}$$
$$+ \ 73.2 \ i \ \mathbb{1}_{\{6666 < i \leq 10000\}} \ + \ 66.01 + 5.84 \ k \ \mathbb{1}_{\{a=k\}}.$$
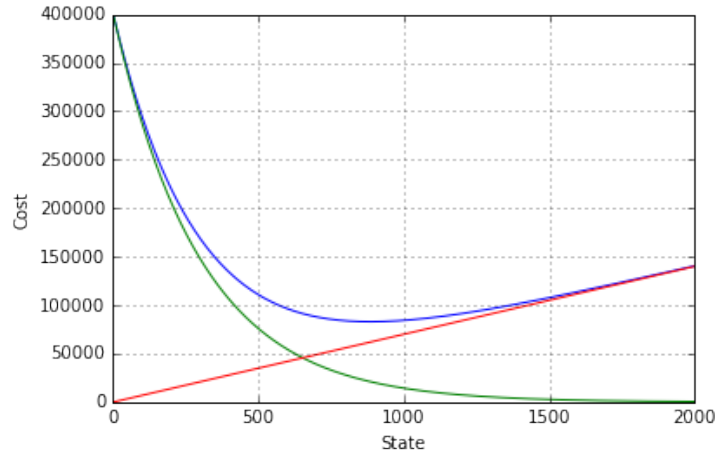
Figure 10: Evolution of costs through the states configured for a 3-day demand scenario. Blue curve represents total costs, sum of production-inventory costs (green) and deficit costs (red).

The optimal average cost for this new setup is R\$82,885.78 (Brazilian Reais), which corresponds to keeping an average stock of 884 bags. It is somehow astonishing that the overall cost reduces about 67%. However, this comes from a structural management change of the organisation that could have long-term consequences. In fact, one can notice that the model is quite sensitive to the adjustment of the deficit cost parameters.

In the second analysis, the $\Delta\lambda$ parameter - the number of useful bags per external collection - shown in table 4 and originally displayed as 59.3, will be varied. Considering external collection only by means of demountable structures (type 1), the collection rate via external teams increases from 59.3 to 61.3, as shown in section 6.1. In the short term, the effect is not as significant, although it will reduce the long-term cost by 3.26%. On the other hand, if HEMORIO was to use, hypothetically, two buses to make the collections, with an average $\Delta\lambda$ of 46.4, the long-term cost would increase by 27.80%. If we use types 1 and 2 in the same proportion, with $\Delta\lambda$ of 53.8, costs would also increase, but by around 10.22%. These results are summarised in Table 5.

Table 5: Sensitivity analysis for variation between collection types.

Policy	Cost	Variation
Original	R\$264,476.73	-
Only type 1	R\$255,854.79	-3.26%
Only type 2	R\$338,001.26	+27.8%
50% each type	R\$291,506.25	+10.22%

It is important to highlight that there are gains from the managerial point of view, such as reducing the complexity of the management variables. However, there may be a negative effect on the demand of potential customers who have preference for buses or restrictions on type 1 collection. Therefore, the selection of the types of external collection to be executed is a complex strategic decision to be made by the top management of the organisation.

January 3, 2020

6.4. Practical implementation

After analysing the results of this study, HEMORIO decided to change its strategy regarding donor recruitment and collection. The decision was to have at least two mobile collections six days a week - from Monday to Saturday. Specific campaigns to solve shortages were focused on recruiting donors for the fixed site.

Using this new strategy, HEMORIO increased by 4.5% the number of blood donations in 2018, compared to 2017. The proportion of donations coming from mobile collections increased from 22.1% to 31% in the same time span.

6.5. Experimental Replication

One of the strengths of the proposed method is the simplicity of the solution procedure. Hence, it can be easily reproduced for distinct case studies.

Algorithm 1 is a classical dynamic programming algorithm (Puterman, 2014). To implement it, one first needs to set up all the parameters for the case study under consideration. Firstly, it is necessary to enumerate all possible states of the system, i.e. all possible inventory levels in the set $S = \{0, 1, \ldots, N\}$, where N is the maximal inventory of blood bags allowed. The action set is $A = \{0, 1, \ldots, M\}$, where M is the maximum number of daily external collections. Then, the parameters in table 4 must be evaluated, namely the donation rate in the main facility λ, the extra rate per external collection team $\Delta\lambda$, the demand rate μ and the shelf life L. For each $a \in A$, we have $\lambda^a = \lambda + a \cdot \Delta\lambda$; and we are now able to evaluate the transition probabilities p_{ij}^a for all $i, j \in S$ and $a \in A$ by means of Eq. (6). Then, the deficit and inventory costs must be assigned for each inventory level $i \in S$. For the former, an exponentially decreasing function in the form of Eq. (7) could be employed, whereas the latter can be directly evaluated for the institution under consideration. Finally, estimates of $c_p(a)$, the expected cost per donation under action a; and $e_c(a)$, the expected additional cost per donation of implementing a external collections, complete the necessary costs. Hence, we are able to set up the cost function in Eq. (8) for all state-action pairs $i \in S$, $a \in A$.

After all the parameters are established, one can easily implement Algorithm 1 using any standard programming language, such as R, Matlab, C++, Python, etc. The output of the algorithm gives, for each possible inventory level $i \in S$, the number of external collection teams that should be sent - $\pi^*(i)$, as well as the long-term average cost η^*. The code in Python that we used to implement our case study is made available for the interested reader in Appendix B.

7. Concluding remarks

This article proposes a Markov decision process to support the decision on sending external collection teams to supplement the donation of blood centres. The model demonstrates that Markov decision processes are adequate because they provide a powerful, yet flexible tool to model the underlying characteristics of the problem. The proposed approach considers perishability and nonlinear cost functions.

The model contributes to the literature by enabling the decision maker to find optimal policies with any number of policy change points without increasing the complexity of the model or the solution procedure. In addition, the model considers perishability, remains computationally tractable for arbitrarily large shelf lives and returns a long-term optimal policy that fully appreciates the stochastic variations of the system in the long-run.

In a first step, we validated the model using data from the literature. In a second step, we applied the model to an important regional blood bank in Rio de Janeiro, HEMORIO. The case study compiled data from HEMORIO's operation since 2009 and a sensitivity analysis was carried out to highlight important aspects from a managerial standpoint. The proposed approach provided invaluable insight for HEMORIO's tactical and strategic planning of external collection activities. As a result of the study, HEMORIO decided to deploy two external collections per business day and that resulted in a 4.5% increment in blood donations from 2017 to 2018. In addition, the percentage of mobile collections increased from 22.1% to 31% in the same period.

Future works could explore extensions to the proposed model including generalisations in the cost function. In addition, the model could be generalised to accommodate general distributions of demand and donation of blood products, giving rise to optimisation models based on GI/GI queues. Finally, we believe that further experimentation would be an interesting way forward. Future works could, for example, map the cost structure of a set of institutions and compare the resulting optimal policies from a managerial standpoint.

Acknowledgements

References

Abbasi, B., Hosseinifard, S.Z., 2014. On the issuing policies for perishable items such as red blood cells and platelets in blood service. Decision Sciences 45, 995–1020.

Abbasi, B., Vakili, G., Chesneau, S., 2017. Impacts of reducing the shelf life of red blood cells: A view from down under. Interfaces 47, 336–351.

Abubakar, U.Y., Hakimi, D., Ndanusa, A., Jevwia, V.A., 2014. Analysis of blood transfused in a city hospital with the principle of Markov-dependence. Applied Mathematical Sciences 8, 8611–8621.

Alagoz, O., Hsu, H., Schaefer, A.J., Roberts, M.S., 2010. Markov decision processes: A tool for sequential decision making under uncertainty. Medical Decision Making 30, 474–483.

Andersen, A.R., Vancroonenburg, W., Berghe, G.V., 2019. Strategic room type allocation for nursing wards through Markov chain modeling. Artificial Intelligence in Medicine 99, 101705.

Angelo, S.A., Arruda, E.F., Goldwasser, R.S., Lobo, M.S.C., Salles, A., Silva, J.R.L., 2017. Demand forecast and optimal planning of intensive care unit (icu) capacity. Pesquisa Operacional 37, 229 – 245.

Arruda, E.F., Pereira, B.B., Thiers, C.A., Tura, B.R., 2019. Optimal testing policies for diagnosing patients with intermediary probability of disease. Artificial Intelligence in Medicine 97, 89 – 97.

Attari, M.N.Y., Pasandideh, S.H.R., Niaki, S.T.A., 2019a. A hybrid robust stochastic programming for a bi-objective blood collection facilities problem (case study: Iranian blood transfusion network). Journal of industrial and production engineering. 36, 154–167.

Attari, M.Y.N., Shabbani, M., Lotfian, F., Jami, E.N., Torkayesh, A.E., 2019b. Red blood cells consumption: An optimization method, in: 2019 15th Iran International Industrial Engineering Conference (IIIEC), IEEE. pp. 274–282.

Azizoglu, M., Cetinkaya, F.C., Kurt, A., Çetinkaya, F.C., 2018. Solution approaches to the blood distribution problem of the turkish red crescent. European Journal of Industrial Engineering 12, 405–441.

Baron, O., Berman, O., Perry, D., 2010. Continuous review inventory models for perishable items ordered in batches. Mathematical Methods of Operations Research 72, 217–247.

Bedi, R.K., Mittal, K., Sood, T., Kaur, P., Kaur, G., 2016. Segregation of blood inventory: A key driver for optimum blood stock management in a resource-poor setting. International Journal of Applied and Basic Medical Research 6, 119.

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

Bennett, C.C., Hauser, K., 2013. Artificial intelligence framework for simulating clinical decision-making: A markov decision process approach. Artificial Intelligence in Medicine 57, 9 – 19.

Besanko, D., Braeutigam, R.R., 2005. Microeconomics. hoboken.

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

Carson, J.L., Grossman, B.J., Kleinman, S., Tinmouth, A.T., Marques, M.B., Fung, M.K., Holcomb, J.B., Illoh, O., Kaplan, L.J., Katz, L.M., Rao, S.V., Roback, J.D., Shander, A., Tobian, A.A.R., Weinstein, R., Swinton McLaughlin, L.G., Djulbegovic, B., 2012. Red Blood Cell Transfusion: A Clinical Practice Guideline From the AABB*. Annals of Internal Medicine 157, 49–58.

Chan, C.W., Yom-Tov, G., Escobar, G., 2011. When to use speedup: An examination of intensive care units with readmissions. Operations Research to appear .

Chao, G.H., 2013. Production and availability policies through the Markov decision process and myopic methods for contractual and selective orders. European Journal of Operational Research 225, 383–392.

Chen, S., Li, Y., Zhou, W., 2019. Joint decisions for blood collection and platelet inventory control. Production and Operations Management .

Clay, N.M., Abbasi, B., Eberhard, A., Hearne, J., 2018. On the volatility of blood inventories. International Transactions in Operational Research 25, 215–242.

Crabill, T.B., 1974. Optimal control of a maintenance system with variable service rates. Operations Research 22, 736–745.

Dehghani, M., Abbasi, B., 2018. An age-based lateral-transshipment policy for perishable items. International Journal of Production Economics 198, 93–103.

Dillon, M., Oliveira, F., Abbasi, B., 2017. A two-stage stochastic programming model for inventory management in the blood supply chain. International Journal of Production Economics 187, 27–41.

Ezugwu, A.E., Olusanya, M.O., Govender, P., 2019. Mathematical model formulation and hybrid meta-heuristic optimization approach for near-optimal blood assignment in a blood bank system. Expert Systems with Applications .

George, J.M., Harrison, J.M., 2001. Dynamic control of a queue with adjustable service rate. Operations research 49, 720–731.

Godin, G., Conner, M., Sheeran, P., Bélanger-Gravel, A., Germain, M., 2007. Determinants of repeated blood donation among new and experienced blood donors. Transfusion 47, 1607–1615.

Goldwasser, R.S., Lobo, M.S.C., Arruda, E.F., Angelo, S.A., Silva, J.R.L., amd C. M. David, A.S., 2016. Difficulties in access and estimates of public beds in intensive care units in the state of rio de janeiro.

Revista de Saúde Pública 50.

Graves, S.C., 1982. The application of queueing theory to continuous perishable inventory systems. Management Science 28, 400–406.

Gunpinar, S., Centeno, G., 2015. Stochastic integer programming models for reducing wastages and shortages of blood products at hospitals. Computers & Operations Research 54, 129–141.

Gunpinar, S., Centeno, G., 2016. An integer programming approach to the bloodmobile routing problem. Transportation Research Part E: Logistics and Transportation Review 86, 94 – 115.

Hamdan, B., Diabat, A., 2019. A two-stage multi-echelon stochastic blood supply chain problem. Computers & Operations Research 101, 130–143.

Hemmelmayr, V., Doerner, K.F., Hartl, R.F., Savelsbergh, M.W., 2009. Delivery strategies for blood products supplies. OR spectrum 31, 707–725.

Hosseinifard, Z., Abbasi, B., 2018. The inventory centralization impacts on sustainability of the blood supply chain. Computers & Operations Research 89, 206–212.

Hosseinifard, Z., Abbasi, B., Fadaki, M., Clay, N.M., 2019. Postdisaster volatility of blood donations in an unsteady blood supply chain. Decision Sciences .

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

Jo, K.Y., Stidham, S., 1983. Optimal service-rate control of m/g/1 queueing systems using phase methods. Advances in Applied Probability 15, 616–637.

Kaspi, H., Perry, D., 1983. Inventory systems of perishable commodities. Advances in Applied Probability 15, 674–685.

Katsaliaki, K., Brailsford, S., 2016. Using simulation to improve the blood supply chain, in: Operational Research for Emergency Planning in Healthcare: Volume 1. Springer, pp. 353–372.

Kazemi, S.M., Rabbani, M., Tavakkoli-Moghaddam, R., Shahreza, F.A., 2017. Blood inventory-routing problem under uncertainty. Journal of Intelligent & Fuzzy Systems 32, 467–481.

Keilson, J., Seidmann, A., 1990. Product selection policies for perishable inventory systems. Technical Report. University of Rochester. Rochester NY.

Kim, G., Wu, K., Huang, E., 2015. Optimal inventory control in a multi-period newsvendor problem with non-stationary demand. Advanced Engineering Informatics 29, 139–145.

Kopach, R., Balcıoğlu, B., Carter, M., 2008. Tutorial on constructing a red blood cell inventory management system with two demand rates. European Journal of Operational Research 185, 1051–1059.

Low, D.W., 1974. Optimal dynamic pricing policies for an M/M/s queue. Operations Research 22, 545–561.

Mansur, A., Vanany, I., Arvitrida, N.I., 2018. Challenge and opportunity research in blood supply chain management: a literature review, in: MATEC Web of Conferences, EDP Sciences. p. 01092.

Murphy, W., McSweeney, E., 2009. Donors and blood collection. Practical Transfusion Medicine, Third Edition , 190–199.

Nagurney, A., Dutta, P., 2019a. Competition for blood donations. Omega 85, 103–114.

Nagurney, A., Dutta, P., 2019b. Supply chain network competition among blood service organizations: a generalized nash equilibrium framework. Annals of Operations Research 275, 551–586.

Nahmias, S., Perry, D., Stadje, W., 2004. Perishable inventory systems with variable input and demand rates. Mathematical Methods of Operations Research 60, 155–162.

Osorio, A.F., Brailsford, S.C., Smith, H.K., 2015. A structured review of quantitative models in the blood supply chain: a taxonomic framework for decision-making. International Journal of Production Research 53, 7191–7212.

Paschalidis, I.C., Tsitsiklis, J.N., 2000. Congestion-dependent pricing of network services. IEEE/ACM Transactions on Networking (TON) 8, 171–184.

Pearson, G.A., Reynolds, F., Stickley, J., 2012. Calculating the need for intensive care beds. Archives of Disease in Childhood 97, 943–946.

Pegels, C.C., Jelmert, A.E., 1970. An evaluation of blood-inventory policies: A Markov chain application. Operations Research 18, 1087–1098.

Perry, D., Posner, M., 1990. Control of input and demand rates in inventory systems of perishable commodities. Naval Research Logistics (NRL) 37, 85–97.

Pirabán, A., Guerrero, W., Labadie, N., 2019. Survey on blood supply chain management: Models and methods. Computers & Operations Research 112, 104756.

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

Puranam, K., Novak, D.C., Lucas, M.T., Fung, M., 2017. Managing blood inventory with multiple independent sources of supply. European Journal of Operational Research 259, 500 – 511.

Puterman, M.L., 2014. Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons.

Rajendran, S., Ravindran, A.R., 2019. Inventory management of platelets along blood supply chain to minimize wastage and shortage. Computers & Industrial Engineering 130, 714–730.

Rock, G., Åkerblom, O., Berséus, O., Herve, P., Jacobs, P., Kelly, T., MacPherson, J., Nydegger, U., Segatchian, G., Urbaniak, S., et al., 2000. The supply of blood products in 10 different systems or countries. Transfusion science 22, 171–182.

Rytilä, J.S., Spens, K.M., 2006. Using simulation to increase efficiency in blood supply chains. Management Research News 29, 801–819.

Samani, M.R.G., Torabi, S.A., Hosseini-Motlagh, S.M., 2018. Integrated blood supply chain planning for disaster relief. International journal of disaster risk reduction 27, 168–188.

Schreiber, G.B., Sharma, U., Wright, D., Glynn, S., Ownby, H., Tu, Y., Garratty, G., Piliavin, J., Zuck, T., Gilcher, R., 2005. First year donation patterns predict long-term commitment for first-time donors. Vox sanguinis 88, 114–121.

Shiryayev, A., 1984. Elementary probability theory, in: Probability. Springer, pp. 5–128.

Snyder, C.S., 2013. A Guide to the Project Management Body of Knowledge: PMBOK (®) Guide. Project Management Institute. 5th edition.

Stanger, S.H., Wilding, R., Yates, N., Cotton, S., 2012a. What drives perishable inventory management performance? lessons learnt from the uk blood supply chain. Supply Chain Management: An International Journal 17, 107–123.

Stanger, S.H., Yates, N., Wilding, R., Cotton, S., 2012b. Blood inventory management: hospital best practice. Transfusion medicine reviews 26, 153–163.

Steimle, L.N., Denton, B.T., 2017. Markov decision processes for screening and treatment of chronic diseases, in: Boucherie, R.J., van Dijk, N.M. (Eds.), Markov Decision Processes in Practice. Springer International Publishing, Cham, pp. 189–222.

Subramanian, J., Stidham Jr, S., Lautenbacher, C.J., 1999. Airline yield management with overbooking, cancellations, and no-shows. Transportation science 33, 147–167.

Wang, K.M., Ma, Z.J., 2015. Age-based policy for blood transshipment during blood shortage. Transportation Research Part E: Logistics and Transportation Review 80, 166–183.

Weber, M., 2015. Optimal inventory control in the presence of dynamic pricing and dynamic advertising. Ph.D. thesis. Humboldt-Universität zu Berlin.

Williamson, L.M., Cardigan, R., 2005. Production and storage of blood components. Practical Transfusion Medicine, Second Edition , 259–273.

World Health Organization, 2010. Towards 100% Voluntary Blood Donation: A Global Framework for Action. WHO Press, Geneva.

Yazer, M.H., Delaney, M., Doughty, H., Dunbar, N.M., Al-Riyami, A.Z., Triulzi, D.J., Watchko, J.F., Wood, E.M., Yahalom, V., Emery, S.P., 2019. It is time to reconsider the risks of transfusing RhD negative females of childbearing potential with rhd positive red blood cells in bleeding emergencies. Transfusion .

Yew-Kwang, N., et al., 2016. Extending Economic Analysis to Analyze Policy Issues More Broadly. Technical Report. Nanyang Technological University, School of Humanities and Social Sciences, Economic Growth Centre.

Yoon, S., Lewis, M.E., 2004. Optimal pricing and admission control in a queueing system with periodically varying parameters. Queueing Systems 47, 177–199.

Özener, O.O., Ekici, A., Çoban, E., 2019. Improving blood products supply through donation tailoring. Computers & Operations Research 102, 10 – 21.

Appendix A. Hemorio's data

In this appendix, we present in more detail the data that allowed the parameter setting for HEMORIO's case study. The data was obtained directly from the institution in the context of the research project. The public information is listed below, however some of the data is very sensitive and remain undisclosed. The data in Table A.6 comprises a monthly series of donations that covers 108 consecutive months from 2009 to 2017. Unfortunately, the costs are not public and therefore cannot be exposed. However, they resulted in the cost function in Eq. (10), which suffices for the implementation of Algorithm 1. Hence, the experimental results in sections 6.2 and 6.3 can be easily reproduced.

Table A.6 features the consolidated series of monthly donations for both internal and external collections. It comprises the following data:

- **Attendance** (Attend.): number of people who volunteered for donation. It is worth pointing out that volunteers that clearly do not qualify for donation may be promptly disqualified. In that case, they are not referred to the screening phase;

- **Medical screening:** volunteers that were referred to the screening phase. Some volunteers may be disqualified in the screening phase;

- **Collection:** total number of donated blood bags.

- **External/Total:** proportion of bags collected externally.

Figures 4 and 5 in section 6 were elaborated making use of the data series presented in table A.6.

Table A.6: Donation Data

Period	INTERNAL COLLECTION			EXTERNAL COLLECTION		
	Attend.	Medical screening	Collection	Attend.	Collection	External/Total
jan/09	8537	8512	6838	50	40	1%
feb/09	7655	7620	6191	116	95	2%
mar/09	9366	9330	7524	1203	968	11%
apr/09	8440	8418	6809	751	588	8%
may/09	9314	9269	7202	970	702	9%
jun/09	8555	8519	6564	592	478	7%
jul/09	8693	8662	6784	705	550	7%
aug/09	9570	9503	7360	1067	799	10%
sep/09	8796	8758	6675	894	719	10%
oct/09	8867	8835	6680	1093	919	12%
nov/09	8440	8391	6248	462	330	5%
dec/09	8336	8273	6424	374	295	4%
jan/10	8855	8811	6740	0	0	0%
feb/10	7696	7636	5745	363	280	5%
mar/10	10108	10045	7870	867	665	8%
apr/10	9088	8989	7069	742	622	8%
may/10	8742	8673	6697	1083	800	11%
jun/10	8118	8059	6368	466	363	5%
jul/10	8542	8478	6604	572	458	6%
aug/10	8805	8747	6514	1153	857	12%
sep/10	8297	8249	6231	1292	976	14%
oct/10	8718	8680	6805	877	715	10%
nov/10	8437	8379	6526	746	580	8%
dec/10	8693	8649	6709	360	291	4%
jan/11	15399	14393	11139	591	439	4%
feb/11	8478	8413	6281	840	631	9%
mar/11	8817	8756	6461	1321	898	12%
apr/11	10312	9970	7279	900	691	9%
may/11	8356	8316	6057	757	543	8%
jun/11	8973	8926	6488	854	644	9%
jul/11	8814	8758	6499	758	542	8%
aug/11	10557	10435	7604	1597	1062	12%
sep/11	7599	7570	5628	932	709	11%
oct/11	7537	7514	5729	377	267	4%
nov/11	10896	10742	8130	648	486	6%
dec/11	8578	8515	6425	303	214	3%
jan/12	9212	9084	6722	551	422	6%
feb/12	8487	8401	6285	563	363	5%

January 3, 2020

mar/12	9446	9382	7004	1422	979	12%
apr/12	7555	7512	5624	943	604	10%
may/12	9729	9635	6992	919	615	8%
jun/12	8111	8072	5903	792	574	9%
jul/12	8444	8395	6229	902	657	10%
aug/12	9049	8992	6808	1601	1187	15%
sep/12	6891	6865	5219	731	537	9%
oct/12	9922	9867	7446	1350	990	12%
nov/12	8857	8781	6660	1287	955	13%
dec/12	7400	7366	5523	659	430	7%
jan/13	9413	9364	7126	438	320	4%
feb/13	7439	7374	5574	753	559	9%
mar/13	7779	7734	5837	929	674	10%
apr/13	8842	8745	6395	1320	959	13%
may/13	8810	8717	6433	807	562	8%
jun/13	7973	7917	5899	908	674	10%
jul/13	7821	7758	5726	937	708	11%
aug/13	8959	8920	6395	1415	1076	14%
sep/13	7677	7649	5536	1395	1070	16%
oct/13	7976	7940	5711	1246	928	14%
nov/13	7879	7813	5738	1031	805	12%
dec/13	7072	7030	5185	524	384	7%
jan/14	7648	7614	5568	255	160	3%
feb/14	8642	8573	6202	970	705	10%
mar/14	8349	8299	6217	1504	1116	15%
apr/14	6892	6855	5073	932	686	12%
may/14	8654	8511	6137	1704	1203	16%
jun/14	6416	6378	4608	766	562	11%
jul/14	7155	7121	5085	396	307	6%
aug/14	7977	7920	5663	1589	1148	17%
sep/14	7912	7842	5699	1374	989	15%
oct/14	7866	7813	5653	1185	867	13%
nov/14	7431	7305	5496	1259	955	15%
dec/14	6991	6928	5120	298	212	4%
jan/15	7146	7111	5315	486	331	6%
feb/15	7530	7480	5605	500	374	6%
mar/15	8335	8296	6174	1091	813	12%
apr/15	6664	6621	4931	1041	850	15%
may/15	6546	6507	4777	673	476	9%
jun/15	8154	8074	6028	1048	821	12%
jul/15	7846	7801	5640	862	618	10%
aug/15	8573	8524	6206	1586	1155	16%
sep/15	6841	6741	4895	926	698	12%

oct/15	7670	7608	5693	1372	1075	16%
nov/15	8986	8857	6447	2104	1571	20%
dec/15	6706	6654	4865	453	326	6%
jan/16	6895	6870	5085	790	601	11%
feb/16	6817	6773	4977	675	487	9%
mar/16	7505	7449	5510	1513	1101	17%
apr/16	5908	5872	4321	938	736	15%
may/16	6240	6201	4581	1184	881	16%
jun/16	8397	8320	6239	1647	1283	17%
jul/16	7816	7710	5813	2106	1619	22%
aug/16	6646	6585	4957	1844	1379	22%
sep/16	7925	7761	5558	2206	1582	22%
oct/16	6320	6126	4553	873	645	12%
nov/16	7688	7569	5783	1773	1398	19%
dec/16	6135	6032	4671	691	541	10%
jan/17	8016	7842	5891	1120	867	13%
feb/17	7217	7053	5259	1448	1099	17%
mar/17	12114	11727	8896	2562	1928	18%
apr/17	8479	8389	6249	2076	1650	21%
may/17	8975	8862	6500	2392	1842	22%
jun/17	9106	8995	6754	2711	2093	24%
jul/17	9629	9535	7214	3334	2558	26%
aug/17	8942	8866	6731	3470	2741	29%
sep/17	8737	8644	6448	3235	2489	28%
oct/17	8417	8374	6449	2396	1899	23%
nov/17	10170	9947	7720	3853	3039	28%
dec/17	7126	7092	5677	1571	1242	18%

Table A.7 consolidates the donations via internal and external collections per year. The data is summarised in figure 6(a). Table A.8 summarises mean number of donated bags per type of collection, which was represented in figure 6(b). It is worth recalling that Type 1 collections are made by demountable structures whereas Type 2 collections make use of an itinerant bus.

Table A.7: Number of external collections by type

Year	Type 1	Type 2	Total	%T1	%T2
2009	90	25	115	78%	22%
2010	92	26	118	78%	22%
2011	104	29	133	78%	22%
2012	118	30	148	80%	20%
2013	122	18	140	87%	13%
2014	114	26	140	81%	19%
2015	145	0	145	100%	0%
2016	172	24	196	88%	12%
2017	349	28	377	93%	7%
Total	1306	206	1512	86%	14%

Table A.8: Number of bags per collection and type

YEAR	Type 1	Type 2	Total
2009	58,71	46,04	55,96
2010	58,74	45	55,71
2011	56,36	42,83	53,41
2012	59,54	42,53	56,09
2013	61,66	50,39	60,21
2014	64,65	54,77	62,81
2015	62,86	0	62,86
2016	65,66	43,42	62,93
2017	63,51	48,29	62,38
Mean	61,3	46,43	60,02

Appendix B. Implementation of the case study in Python

```
#DATA ENTRY MODEL BY USER

import math
import numpy as np
import time
import matplotlib.pyplot as plt
from numpy import linalg as LA
get_ipython().magic('matplotlib inline')

t_inicial = time.time()
# parameter definition

lamb = 182                              #stock increase rate associated with donation offer
mi = 258                                #inventory consumption rate associated with hospital demand
c = 70                                  #expected capacity of each external collection
estados = 1000                          #number of states
bus = 140                               #fixed cost of external collection
prod = 50                               #cost of production of a blood bag
cf = 1000                               #default cost parameter
df = 80                                 #concerning the exponential decay of the cost of missing a blood bag
ce1 = 0.7                               #parameter of growth rate of blood bags inventory cost
```

```
ce2 = 0.1                                               #relative to the economies of scale in inventory costs
ce3 = 0.7                                               #relative to the loss of economies of scale in inventory costs
n0 = 0
n_bus = 2



#parameters chosen by the user

lamb = input("input lambda:   ")
mi = input("input mi:  ")
c = input("input delta_lambda:  ")
estados = input("input s:  ")
bus = input("input cc:  ")
prod = input("input cp:  ")
cf = input("input c1:  ")
df = input("input d1:  ")
ce1 = input("input s1:  ")
ce2 = input("input s2:  ")
ce3 = input("input s3:  ")

lamb = int(lamb)
mi = int(mi)
c = int(c)
estados = int(estados)
bus = int(bus)
prod = int(prod)
cf = int(cf)
df = int(df)
ce1 = float(ce1)
ce2 = float(ce2)
ce3 = float(ce3)




# Initialization of the cost vector

C = [0]* (estados)                                      #Co,C1,C2,...,Cn-1,Cn -> Costs vector
Cf = [0]* (estados)                                     #Lack cost vector
Ce = [0]* (estados)                                     #Inventory cost vector
for i in range(estados):
    C[i] += cf*math.exp(-(i/df)) + ce1*i
#    Cf[i] += cf*math.exp(-(i/df))
#for i in range (0,round(estados/4)):
#    Ce[i] += ce1*i
#for i in range (round(estados/4),round(3*estados/4)):
#    Ce[i] += (Ce[round(estados/4)-1] + ce2*(i-round(estados/4)+1))
#for i in range (round(3*estados/4),estados):
#    Ce[i] += (Ce[round(3*estados/4)-1]+ ce3*(i-round(3*estados/4)+1))
#
#for i in range(estados):
#    C[i] += Cf[i] + Ce[i]


v_old = [0] * len(C)
v_new = [0] * len(C)
v_new_menos_old = [0] * len(C)

mem_acoes = ['0bus'] * len(C)
gap = 100
it = 0

lista_estados = []
for i in range(1,estados-1):
```

```python
        lista_estados.append(i)




v_p_sup = [0]*(n_bus+1)
v_p_inf = [0]*(n_bus+1)
v_custo_acoes = [0]*(n_bus+1)

for i in range(n_bus+1):
    v_p_sup[i] = round((lamb+i*c)/(lamb+i*c+mi),3)
    v_p_inf[i] = round(mi/(lamb+i*c+mi),3)
    v_custo_acoes[i] = v_p_sup[i]*(C[i+1]+v_old[i+1]+prod) + v_p_inf[0]*(C[i-1]+v_old[i-1])


p_sup_0bus = round(lamb/(lamb+mi),3)          #sup refers to the probability of going from an n state to an n + 1
p_inf_0bus = round(mi/(lamb+mi),3)              #inf refers to the probability of going from an n + 1 state to an n
p_sup_1bus = round((lamb+c)/(lamb+c+mi),3)
p_inf_1bus = round(mi/(lamb+c+mi),3)
p_sup_2bus = round((lamb+2*c)/(lamb+2*c+mi),3)
p_inf_2bus = round(mi/(lamb+2*c+mi),3)

#--------------------------------------------------------------------------------------------------------------
#algorithm

#while gap > 0.1:
for i in range(10000):

    it += 1
    for j in range(len(v_old)):
        v_old[j] = v_new[j]
    v_new[0] = min(p_sup_0bus*(C[1]+v_old[0]) + p_inf_0bus*(C[0]+v_old[0]),\
                p_sup_1bus*(C[1]+bus+v_old[0]) + p_inf_1bus*(C[0]+bus+v_old[0]), \
                p_sup_2bus*(C[1]+2*bus+v_old[0]) + p_inf_2bus*(C[0]+2*bus+v_old[0]))
    if v_new[0] == p_sup_0bus*(C[1]+v_old[0]) + p_inf_0bus*(C[0]+v_old[0]):
        mem_acoes[0] = '0bus'
    elif v_new[0] == p_sup_1bus*(C[1]+bus+v_old[0]) + p_inf_1bus*(C[0]+bus+v_old[0]):
        mem_acoes[0] = '1bus'
    else:
        mem_acoes[0] = '2bus'

    for i in lista_estados:
        v_new[i] = min(p_sup_0bus*(C[i+1]+v_old[i+1]) + p_inf_0bus*(C[i-1]+v_old[i-1]), \
                    p_sup_1bus*(C[i+1]+bus+v_old[i+1]) + p_inf_1bus*(C[i-1]+bus+v_old[i-1]), \
                    p_sup_2bus*(C[i+1]+2*bus+v_old[i+1]) + p_inf_2bus*(C[i-1]+2*bus+v_old[i-1]))
        if v_new[i] == p_sup_0bus*(C[i+1]+v_old[i+1]) + p_inf_0bus*(C[i-1]+v_old[i-1]):
            mem_acoes[5] = '0bus'
        elif v_new[i] == p_sup_1bus*(C[i+1]+bus+v_old[i+1]) + p_inf_1bus*(C[i-1]+bus+v_old[i-1]):
            mem_acoes[i] = '1bus'
        else:
            mem_acoes[i] = '2bus'

  v_new[estados-1] = min(p_sup_0bus*(C[estados-1]+v_old[estados-1]) + p_inf_0bus*(C[estados-2]+v_old[estados-estados-2]), \
                p_sup_1bus*(C[estados-1]+bus+v_old[estados-1]) + p_inf_1bus*(C[estados-2]+bus+v_old[estados-2]),   \
                p_sup_2bus*(C[estados-1]+2*bus+v_old[estados-1]) + p_inf_2bus*(C[estados-2]+2*bus+v_old[estados-2]))
    if v_new[estados-1] == p_sup_0bus*(C[estados-1]+v_old[estados-1]) + p_inf_0bus*(C[estados-2]+v_old[estados-2]):
        mem_acoes[estados-1] = '0bus'
  elif v_new[estados-1] == p_sup_1bus*(C[estados-1]+bus+v_old[estados-1]) + p_inf_1bus*(C[estados-2]+bus+v_old[estados-2]):
        mem_acoes[estados-1] = '1bus'
    else:
        mem_acoes[estados-1] = '2bus'


    for i in range(len(v_old)):
        v_new_menos_old[i] = v_new[i] - v_old[i]
```

```
        gap = max(v_new_menos_old) - min(v_new_menos_old)

t_final = time.time()




#Creating the probability matrix
P = np.zeros((estados,estados))
for i in range(1,estados-1):
    for j in [i-1]:
        if mem_acoes[i] == '0bus':
            P[i][j] = v_p_inf[0]
        else:
            P[i][j] = v_p_inf[1]
    for j in [i+1]:
        if mem_acoes[i] == '0bus':
            P[i][j] = v_p_sup[0]
        else:
            P[i][j] = v_p_sup[1]
for i in [0]:
    for j in [i]:
        if mem_acoes[i] == '0bus':
            P[i][j] = v_p_inf[0]
        else:
            P[i][j] = v_p_inf[1]
    for j in [i+1]:
        if mem_acoes[i] == '0bus':
            P[i][j] = v_p_sup[0]
        else:
            P[i][j] = v_p_sup[1]
for i in [estados-1]:
    for j in [i-1]:
        if mem_acoes[i] == '0bus':
            P[i][j] = v_p_inf[0]
        else:
            P[i][j] = v_p_inf[1]
    for j in [i]:
        if mem_acoes[i] == '0bus':
            P[i][j] = v_p_sup[0]
        else:
            P[i][j] = v_p_sup[1]

#system resolution P (Transposed). x = x
P_acum= LA.matrix_power(P, estados*100)
P_acum = np.add.accumulate(P_acum[0])
P_acum = P_acum.tolist()


#Calculating blood bags in stock K

fp = [0] * estados
for i in range(len(P_acum) -1):
    fp[i] += P_acum[i+1] - P_acum[i]

fp_n = [0] * estados                        #fp normalization predicting numerical errors
for i in range(estados):
    fp_n[i] += fp[i]/sum(fp)


Ex = [0] * estados
for i in range(estados):
    Ex[i] += list(range(estados))[i]*fp_n[i]

V_Ex = sum(Ex)                          #Expected number of blood bags is the expected value of the prob function.
```

```
#---------------------------------------------------------------------------------------------------
#output

#print("Actions: ", mem_acoes)
print('')
print("Optimal policy: Send 2 bus from state 0 to state ",mem_acoes.count('2bus')-1-n0)
print("send 1 bus from state",mem_acoes.count('2bus')-n0, "to state", mem_acoes.count('2bus')+mem_acoes.count('1bus')-1-n0)
#print("Expected value of Inventory E(n): ", V_Ex)
#print("Iterations: ", it)
#print("Execution Time ", round(t_final - t_inicial,1), " segundos")




#-------------------------------------------------------------------
# Figures

#graph of the cumulative probability of states across states
#plt.plot(range(1-n0,estados+1-n0),P_acum)
#plt.title("Accumulated probability across states")
#plt.ylabel("Prob.")
#plt.xlabel("State")
#plt.grid(True)

#cumulative probability inverse logarithm graph
#ln_P_acum = []
#for Pn in P_acum:
#    ln_P_acum.append(math.log(Pn**-1))

#plt.plot(range(1-n0,estados+1-n0),ln_P_acum)
#plt.title("Accumulated probability across states")
#plt.ylabel("Prob.")
#plt.xlabel("State")
#plt.grid(True)

#graph of C cost evolution
#plt.plot(range(1-n0,len(C)+1-n0),C)
#plt.title("graph of C cost evolution")
#plt.ylabel("Cost")
#plt.xlabel("State")
#plt.grid(True)

#graph of Cf cost evolution
#plt.plot(range(1-n0,len(Cf)+1-n0),Cf)
#plt.title("graph of Cf cost evolution")
#plt.ylabel("Lack cost")
#plt.xlabel("State")
#plt.grid(True)

#graph of Ce cost evolution
#plt.plot(range(1-n0,len(Ce)+1-n0),Ce)
#plt.title("graph of Ce cost evolution")
#plt.ylabel("inventory cost")
#plt.xlabel("State")
#plt.grid(True)


#graph of lambda arrival rate
v_lambda = [lamb]*estados
for i in range(mem_acoes.count('2bus')-1-n0):
    v_lambda[i] += c
for i in range(mem_acoes.count('2bus')- n0 +mem_acoes.count('1bus')-1):
    v_lambda[i] += c
```

```
plt.plot(range(1-n0,len(v_lambda)+1-n0),v_lambda)
plt.title("Policy")
plt.ylabel("Lambda")
plt.xlabel("States")
plt.grid(True)
plt.show()
```