LI Simpósio Brasileiro de Pesquisa Operacional
Limeira–SP, 02 a 06 de setembro de 2019

Optimal inventory policies for surgical supplies through a decision making model using Markov Decision Process for a minimum service level

Nicholas B. Richers, Mateus G. S. Amorim, Edilson F. Arruda
Programa de Engenharia de Produção, Insituto Alberto Luiz Coimbra de Pós Graduação e Pesquisa de Engenharia, Universidade Federal do Rio de Janeiro
Av. Horácio de Macedo, 2030. Cidade Universitária, Rio de Janeiro RJ, 21.941-972
nicholasrichers@pep.ufrj.br, mateusamorim@poli.ufrj.br, efarruda@po.coppe.ufrj.br

Helena C. V. Rey
Instituto Nacional de Cardiologia
R. das Laranjeiras, 374 - Laranjeiras, Rio de Janeiro - RJ, 22240-002
hcramer@globo.com

RESUMO

A importância do desafio de gerenciar políticas de estoque para suprimentos de cirurgias em hospitais é agravado pelos altos custos dos materiais envolvidos. Logo o presente estudo, busca oferecer ao tomador de decisão um modelo que permite estabelecer uma política de reabastecimento ótima que minimiza os custos de longo prazo da operação. Como a demanda dos suprimentos é oriunda das cirurgias, a demanda é obtida de maneira indireta e em seguida o gestor deve optar entre modelar o problema como um Processo de Decisão de Markov, ou optar por utilizar um algoritmo de ponto de recompra, que possui uma implementação mais simples e tempo computacional mais baixo, mas que pode apresentar resultados sub ótimos. Para ambos os modelos discutiu-se os efeitos da adição do nível de serviço mínimo ao modelo sobre a respectiva política e os custos de longo prazo.

PALAVRAS CHAVE. Políticas de Inventário, Processo de Decisão de Markov, Nível de Serviço

ABSTRACT

The importance of a properly management of surgical supplies in hospitals is enhanced by high costs of the materials involved.Then, this study aims to offer to the decision maker a replenishment policy which minimizes the operation's long-run average cost. As the demand for supplies comes from surgeries, demand is indirectly obtained. Next, the model allows the manager to choose between a reorder point algorithm, which has a straight implementation and low computational time, although can present sub optimal results or model the problem as a Markov Decision Process and handle with the modeling task. For both models, the effects of adding the minimum service level to the model on the respective policy and the long-term costs were discussed.

KEYWORDS. Inventory Policy. Markov Decision Process. Service Level.

https://proceedings.science/p/106772?lang=pt-br

1. Introduction

A proper management of medical items, such as medicines and surgical supplies, is essential to ensuring hospital activity. Such a management is complicated due to the high costs involved and the low demand for very specialized supplies. Another distinguishing trait is that the demand is not direct. Instead, the demand is for surgical procedures and the actual need for specific surgical supplies is sometimes only unveiled at the time of surgery, depending on the conditions of the patient undertaking the procedure. Dallora and Forster [2013] report that 35% of the budget of a specialized research hospital is spent on medicines and materials. In Brazil, for example, cardiac surgery and hemodynamic supplies can cost as much $27,000 BRL a piece [Portal Comprasnet, 2018]. Moreover, as contracts with suppliers are occasionally built on minimum purchase quantities, an accurate knowledge of the stochastic demand is of the ultimate importance for hospital purchasing managers [Volland et al., 2017].

The increase in hospital inventory management literature is probably due, in part, to the large expenses on medical supplies. According to Volland et al. [2017], from 2012 to 2014 the literature on hospital management logistics was twice as large as in the preceding three years. Nevertheless, the majority of academic publications are somewhat related to perishable supplies such as blood [Puranam et al., 2017; Soares et al., 2017; Kouki et al., 2015] and medicines [Mubiru, 2018].

In terms of surgical supplies, each item may be considered critical and there is a perceived need to provide very high levels of service [Beier, 1995]. This work studies inventory policies for surgical supplies. The problem is modeled as a Markov decision process (MDP) whose output is an optimal weekly inventory policy for each supply. Additionally, an alternative method to derive a weekly inventory policy by means of a simpler reorder point algorithm is investigated. The latter method features a reduced computational level, however it may result in sub-optimal strategies.

The proposed model can seamless address the highly stochastic nature of the demand for surgical supplies. Indeed, such a demand may exhibit an intermittent nature, and possibly large variances. In such a case, compound distributions provide an adequate framework for modeling purposes [Turrini and Meissner, 2018]. So, in order to improve the inventory control of medical supplies, a discrete-time Markov decision model were formulated so that orders are issued according to a parametric policy which minimizes the long-run average cost per review period for a given service level.

To make this paper possible, the presented data were collected in an elite public hospital specialized in cardiac surgeries in the city of Rio de Janeiro. It is composed of weekly demand for 235 items over the period of 2016–2018 and comprises the supplies demanded by patients subject to hemodynamic procedures. These procedures were selected because they tend to involve the most expensive supplies. Therefore, an objective of this study is to help improve the inventory management of surgical supplies at the partner hospital.

The remainder of the paper is organized as follows. In Section 2, it presents the method that was used followed by the literature review of each topic. Section 3 is dedicated to the model implementation, composed by the ordering policies, the stochastic demand estimation and the models implemented to optimize the safety stock. Section 4 presents the obtained results and lastly section 5, consists of conclusions and directions for further research.

2. Literature Review

This section features a brief literature review on the main aspects of this study. Section 2.1 introduces some approaches for estimating the demand distribution. In Section 2.2 some basic assumptions on inventory control models are briefly described, in addition an important relation

https://proceedings.science/p/106772?lang=pt-br

between shortage cost and service level is introduced. Section 2.3 outlines the research conducted in safety stock optimization, emphasizing Markov Decision Process approaches.

2.1. Demand estimation models

Several approaches have been proposed in the literature to estimate demand distributions. The most common framework assumes that the demand has a normal distribution [e.g., Syntetos and Boylan, 2005], and estimates the mean demand per period by applying exponential smoothing separately to the intervals between non-zero demands, also accounting for the size of these intervals. However, the use of the normal distribution to model discrete demands has little empirical support in the studied problem. Instead, the use of compound distributions is perfectly justified given the intermittent nature of the demand [Turrini and Meissner, 2018].

In this context Axsäter [2015] proposed the use of compound Poisson models when transactions are larger than a unit size. It is important to emphasize that the demand for supplies in the studied case is indirect. For a given patient, the actual demand is for a surgical procedure which, depending on the patient's condition, requires a set of supplies. The demand for the surgical procedure is modeled as a Poisson process, whereas the demand for each supply is given by an empirical distribution directly constructed from data [Porras and Dekker, 2008]. The modeling process for the indirect demand for surgical supplies is also a contribution of this paper.

2.2. Inventory Control System

Inventory models are generally based on certain assumptions. An inventory position can be reviewed *continuously* or *periodically*. In both cases an order is triggered as soon as the inventory position is sufficiently low, and will be delivered after a certain lead-time [Axsäter, 2015]. In addition, inventory systems sometimes follow a (r, Q) ordering policy, whereby an order of exactly Q units is placed whenever the inventory position is less than or equal to r at a given decision time [Kouki et al., 2015]. Other times an (s, S) policy is implemented whereby an order is placed that returns the inventory position to S items whenever the inventory is less than or equal to s units [O. Keizer et al., 2017]. Both strategies are called parametric or 'pure policies' [Johansen and Thorstenson, 2014].

It is well-known that the (s, S) policy is optimal under quite general conditions. For example, Scarf [1959] proved optimality under convex holding and shortage costs. More recently, Benkherouf and Suresh [2010] proved the optimality of (s, S) policies under a different assumption on the so-called lump-sum shortage cost. In some cases, however, parametric policies may fail to attain optimality. In such a situation, authors often resort to a heuristic procedure based on a parametric policy while aiming to improve performance measures such as the service level [e.g., Minner and Silver, 2007].

2.3. Safety Stock Optimization

In general the costs incurred in an inventory control problem are *holding*, *ordering* and *shortage costs* [Axsäter, 2015]. Shortage costs, however, may be very difficult to estimate because the effects may go far beyond just lost sales. To circumvent this issue, decision makers often classify inventory policies in terms of their resulting service level.

Some heuristic techniques may be employed to determine a safety stock level that corresponds to a given desired service level. Such an approach, however, may lead to sub-optimal inventory policies [Axsäter, 2015]. In fact, numerical studies in the literature confirm that Markov decision processes (MDP) generally outperform the best pure replenishment policy [e.g. Johansen and Thorstenson, 2014; O. Keizer et al., 2017]. Hence, an MDP model will be introduced later to provide a decision-aid framework for the studied application.

https://proceedings.science/p/106772?lang=pt-br

2.3.1. Markov Decision Process

Markov Decision Processes (MDP) are often used to solve inventory control problems with stochastic demand [Turrini and Meissner, 2018]. O. Keizer et al. [2017] formulated a joint policy to optimize inventory costs for spare parts of an equipment. The authors made use of an MDP model whose solution was found by means of a classical *value iteration algorithm* [e.g., Puranam et al., 2017] and compared to a parametric $(S-1, S)$ policy. Ahiska et al. [2013] consider a single-product periodic-review inventory system for a retailer who has adopted a dual sourcing strategy to cope with potential supply interruptions. The problem is modeled as an MDP whose action space is comprised of pairs of order quantities and service providers.

Nevertheless, neither of these studies cover some peculiarities related to hospital supply management. In a related context, Ferreira et al. [2018] contribute with an MDP model for inventory management of perishable goods for humanitarian operations with stochastic demands and donations. The specific management of blood platelet inventory was addressed in [Civelek et al., 2015] by means of an MDP model aided by an allocation heuristic. Note that, albeit correlated, the problem studied here is more complex in that it involves intermittent and indirect demand for surgical supplies.

2.3.2. Research Contribution

This paper contributes with a novel decision making model for inventory management in hospitals. The novelty lies in how the demand for items is estimated. Generally, when planning for the future, hospital managers use the historical demand for individual supplies to estimate the future need. We argue that such an approach is biased because supplies do not have a demand of their own. Instead, they are demanded (or not) in connection with each surgery performed. As a result, an estimation based on historical demand would only be valid if the actual demand for surgeries did not vary over time. We argue, however, that this is rarely the case, given the population changes over time and the dynamic nature of the problem. To circumvent this issue, we model the demand for individual supplies for each type of surgical procedure. The proposed model makes use of the demand demand for supplies of each individual surgery, estimated from historical data. Then, given the distribution of the expected demand for the surgical procedure in question, we introduce a recursion that estimates the overall demand for a given supply considering all surgical procedures of that type. To obtain the overall demand for each supply, it suffices to add the demand distributions over the set of surgical procedures.

Another contribution of this work is in striving to uncover the relationship between service level and shortage costs. The latter is often employed in inventory models, particularly those that make use of MDPs. We explore a concept found in the literature to recursively find an optimal parametric policy by means of Markov decision processes.

3. Problem Description

This section introduces a Markov decision process which obtains the optimal parametric (s, S) and (r, Q) policies. The results are contrasted to those provided by a reorder point algorithm.

3.1. Model Characteristics and Assumptions

1. The demand for procedures is stochastic and follows a Poisson distribution with mean (λ);

2. The demand for a given supply during a procedure is described by a non-negative discrete random variable. More details in Section 3.2;

3. The cost parameters of the system are stationary and each unit carried in inventory during a period incurs a *holding cost*. Orders arrive instantaneously and incur a fixed *ordering*

cost. The demand is immediately deducted from the inventory when there is stock on-hand. Otherwise, it is backlogged and met as soon as sufficient inventory is made available. A shortage cost is incurred per period per unmet demand.

4. For each inventory item, an independent *ordering policy* is assigned.

3.2. Mathematical Formulation

In this section, the mathematical formulation of the model will be detailed. Let Y be a discrete non-negative random variable representing the demand for surgical procedures at a given period. By assumption, Y has a Poisson distribution with parameter λ, i.e.

$$P(Y = k) = e^{-\lambda}\frac{\lambda^k}{k!}.$$

The periodic demand for a given supply is represented by random variable D, and the probability that j items of this supply will be demanded, given that k surgeries are performed is given by:

$$f_j^k = P(D = j|Y = k). \tag{1}$$

Observe that, for ease of notation, we drop the index corresponding to the specific supply. The remaining parameters are described in Table 1.

Tabela 1: Model parameters and definitions

Notations	Definition
K	Ordering Cost per period
h	Holding Cost per period
b	Shortage Cost per period
η_π	Long-run Average Cost

It is easy to see from the definition in Eq. (1) that $f_0^0 = 1$. For a single surgery, one can estimate f_j^1, $j = 0, 1, \ldots$ from actual data on procedures of the same type previously performed. Assuming that the surgeries are independent an that the demand for supplies per surgery is independent and identically distributed ($i.i.d$), we can now build a recursion to determine the remaining values of f_j^k as follows:

$$f_j^k = \sum_{i=0}^{k} f_i^1 f_{j-1}^{k-i},\ j = 2, 3, 4 \ldots \tag{2}$$

Considering the previous results, one can see that:

$$P(D = j) = \sum_{k=0}^{j} f_j^k P(Y = k). \tag{3}$$

Algorithm 1 below illustrates the process of finding the distribution of the overall demand for a given supply, considering all possible procedures by means of a recursive routine.

https://proceedings.science/p/106772?lang=pt-br

Algorithm 1: Demand probability of a supply considering all possible procedures

Result: P_{supply}
$\rho_{prev} = \emptyset$;
forall *Procedures* **do**
  Estimate the distribution of D given the procedure, using Eq. (3);
  Make $P_{supply} = \rho_{prev} + D$ and find the distribution of P_{supply} ;
  $\rho_{prev} \leftarrow P_{supply}$;
end

3.2.1. Ordering Policies

The model involves a fixed ordering cost *(K)*, a unit holding cost *(h)* and total procurement cost of purchasing z units is denoted by $(\tau(z))$ where

$$\tau(z) = \begin{cases} 0, & if \quad z = 0 \\ K + \tau z, & if \quad z > 0 \end{cases} \tag{4}$$

Furthermore, there is also the unitary shortage cost *(b)* which is a function of the excess of demand over supply at the end of the period. The initial inventory *(X=x)* plus delivered units (ξ) is denoted by *(y)* and the demands are discrete random variables with density (ϕ) and mean $(\eta > 0)$. The expected holding and shortage costs to be charged for each period are given by

$$L(y) = \begin{cases} h\sum_{\xi=0}^{y}(y-\xi)\phi(\xi), & if \quad y \geq 0 \\ \\ b\sum_{\xi=y}^{\infty}(\xi-y)\phi(\xi), & if \quad y < 0 \end{cases} \tag{5}$$

If the ordering and shortage costs are linear and ϕ has a monotone likelihood ratio, or more generally when $L(y)$ is convex, and the ordering costs is described as in Eq. (4), then there is always an optimal policy of the $(s,\ S)$ type without the need for additional conditions [Scarf, 1959; Aneja and H.Noori, 1987; Cetinkay and Parlar, 2002]. Next, denote by $C_n(x)$ the expected value of the discounted costs accumulated over the first n periods, under a discount factor α. Therefore, according to Cetinkay and Parlar [2002] $C_n(x)$ satisfies

$$C_n(x) = \underset{y \geq x}{Min} \left\{ \tau(z)(y-x) + L(y)\alpha \sum_{\xi=0}^{\infty} C_{n-1}(y-\xi)\phi(\xi)d\xi \right\} \tag{6}$$

However, in certain situations managers may find the $(s,\ S)$ policy inadequate. Indeed, many multi-item system are designed such to operate under an $(r,\ Q)$ policy [Federgruen and Zheng, 1992]. Keeping the same assumptions previously defined and according to Axsäter [2015] and considering the invariant distribution (γ_x) of the inventory level and the supply's demand (μ) as previously defined, the long-run average cost of such a policy is

$$C(r) = -b(r + \frac{Q+1}{2} - \mu) + (h+b)\sum_{x=1}^{r+Q} \gamma_x \tag{7}$$

Classical reorder point algorithms often make use of Eq. (6)-(7) to evaluate the performance by means of an algorithm described in section 3.2.3. A distinct approach that makes use of Markov decision processes is presented in Section 3.2.4.

3.2.2. Service Level

Service level is often a measure designed to quantity the proper functioning of an inventory policy. Among the most frequent measures one finds *fill rate (S_2)*, which is the portion of the demand immediately satisfied with the stock in hand, and the *ready rate (S_3)*, which can be defined as the probability that the stock is positive. At first, the fill rate may sound like the better choice. However, since the majority of supplies analyzed in this study have low demand ($\geqslant 90\%$ of zero demand during per period), the respective fill rate may present good results even if there is no stock in hands to satisfy the demand. Hence, the service level will be measured by ready rate given, i.e. the probability that the stock is positive.

Now, consider an inventory system with certain discrete demand, which is controlled by a continuous review $(r,\, Q)$ policy. Once steady state is reached, it follows that the inventory position is uniformly distribution in the interval $[r+1,\, r+Q] \in \mathbb{Z}$ [Axsäter, 2015]. Through this result, the *long-run Average cost* (η_π) can be calculated as follows:

$$\eta_\pi = \lim_{n\to\infty} \frac{1}{n} \sum_{t=1}^{n} \tau(X_t) = \sum_{j=1}^{M} \gamma_j \tau(j), \tag{8}$$

where $\gamma : S \to [0, 1]$ is the stationary distribution of the inventory process and π is the control policy, which prescribes a single order to be placed at any given state of the system.

3.2.3. Reorder Point Algorithm

The $(r,\, Q)$ policy prescribes two parameters, namely the reorder point r and the lot size Q. An effective way to select reorder point is to make use of Eq. (7). Starting from an arbitrary $(r,\, Q)$ policy, one keeps Q constant and increases r by one unit. It follows that [Axsäter, 2015]:

$$C(r+1) - C(r) = -b + (h+b)S_3(r+1) \tag{9}$$

Note that Eq. (9) is convex because $S_3(r+1)$ increases with r. Hence, the optimal value for R^* is that which nullifies the right hand side of (9). Then, one can make use of the newly found r^* to improve the lot size to Q^* through algorithm 2.

Algorithm 2: Reorder Point Algorithm for (r,Q) policy

Result: R^* and Q^*
Starts with $r = -Q$;
Calculate $C(r+1) - C(r)$;
if $C(r+1) - C(r) \leq 0$ **then**
    $r \leftarrow r + 1$;
    Go to step 2 ;
else
    Break;
end
$Q \leftarrow Q - 1$;
Calculate $C(r+1) - C(r)$;
if $C(r+1) - C(r) \geq 0$ **then**
    $Q \leftarrow Q + 1$;
    Break ;
else
    while $C(r+1) - C(r) \leq 0$ **do**
        $Q \leftarrow Q - 1$;
    end
end

For $(s,\ S)$ policies Zheng and Federgruen [1991] has developed a similar two-phase scheme. Finally, note that a relation between shortage cost and service level is described in Eq. (10) [Axsäter, 2015]

$$S_3(r^*) \leq \frac{b}{h+b} \leq S_3(r^*+1) \tag{10}$$

The equation above enables us to prescribe effectively the shortage cost for a given prescribed service level, so long as we operate under a $(r,\ Q)$ policy.

3.2.4. Markov Decision Processes

Markov Decision Processes (MDP) allow us to model processes with probabilistic transitions among states, in which the decision maker is capable of interfering in the dynamics at specific decision epochs [Puterman, 1994]. However, the system keeps the Markov property, according to which the effect of an action on a state depends both on the action and on the current state. Every action has an immediate reward associated to it, which depends on the state that the process happens to be.

State Space. The system states represent the available inventory level of a given product. There is a defined limit ($\bar{b}$) for the number of units that can be back-ordered at any time. Any demand beyond this limit is lost and is counted as patient waiver. Furthermore, the inventory cannot exceed an arbitrary maximum level $\bar{S}$. Consequently, the state space is discrete and denoted by $S = \{\bar{b}, ..., -1, 0, 1, ..., \bar{S}\}$.

Actions Space. At the beginning of each period, an order can be placed by the manager, limited only by state space constraints and denoted by a. Let $A \in \{0, 1, ..., \bar{S}\}$ be the set of feasible actions. Note that if no other restriction is placed, you will end up with an $(s,\ S)$ policy. It happens because this parametric policy is optimal (see section 3.2.1). To reach an $(r,\ Q)$ policy, an additional constraint is implemented that only permits actions $a = Q$ or $a = 0$.

https://proceedings.science/p/106772?lang=pt-br

Transition Probabilities. Only the demand (defined in section 3.2) and orders (actions) can cause transitions between states. So, the transition denoted by T: S x A x S $\subset$ [0, 1] is the probability function of a system to switch to a state $s^{'} \in S$ given the process is in a state $s \in S$ and the agent performed and action $a \in A$.

Costs. Represent the cost incurred at a given state $s \in S$ while performing a given action $a \in A$. These costs are defined in Eq. (4) and (6) in Section 3.2.1.

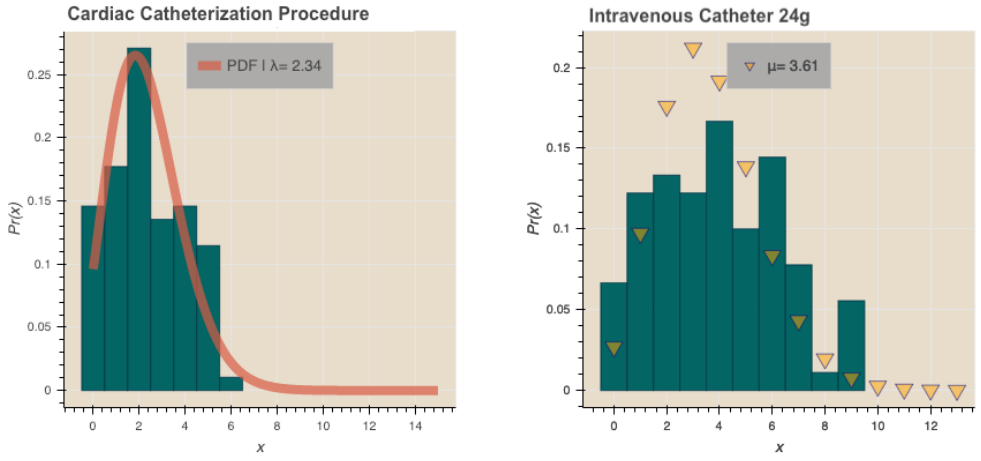
Objective Function. In order to find the optimal policy, one needs to assess the long-run average cost of the system, defined in Eq. (8). Recall that the objective is to find an optimal policy $\pi^* \in \Pi$ such that: $\eta_{\pi^*} \leq \eta_\pi, \forall \pi \in \Pi$, where Π is the set of all feasible control policies. In order to solve the problem with infinite horizon we used the Value Iteration Algorithm and for average reward criteria found in [Puterman, 1994]. For more details see Puterman [1994].

4. Numerical Results

This section features the results of a case study at a reference hospital in Cardiology in the city of Rio de Janeiro. The demand for each supply is estimated based on the procedure in Algorithm 1. It also illustrates the structure of (s, S) and (r, Q) policies and their respective inventory level distribution in steady state. And lastly, the performance of the reorder point algorithm was compared to that of the MDP model. The database is composed of 441 procedures, classified in four categories. Moreover there are 12307 entries that account for 235 individual supplies.

4.1. Stochastic demand estimation

Following the scheme proposed in Section 3.2, Figure 1a illustrate the approximation of a demand distribution for surgical procedures by a Poisson distribution with mean (λ).



(a) Histogram of arrivals of "Cardiac Catheterization"

(b) Demand histogram of "Intravenous Catheter"

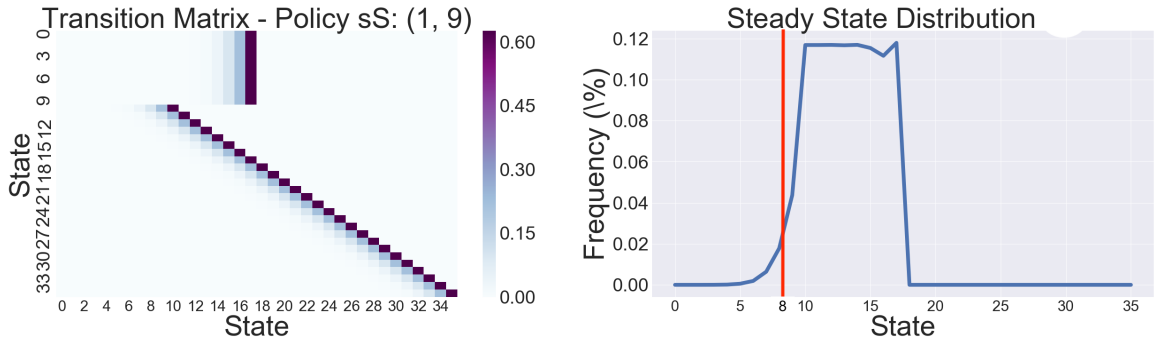
Figura 1: Patient demand divided by four categories for procedures

This result will be used in Eq. 3 to generate the supplier distribution depicted in Figure 1b, with mean (μ).

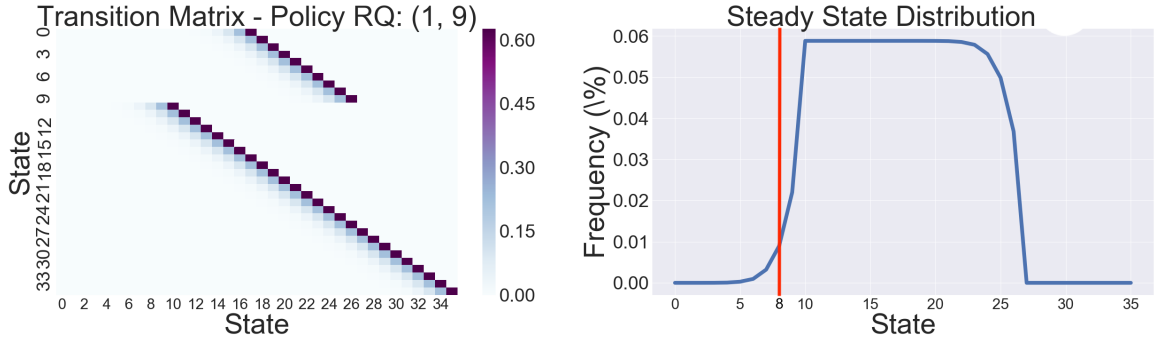
4.2. Inventory Level Distribution

Figure 2a(i) depicts the transition matrix of a (s, S) policy with $s = 1$ and $S = 9$. One can observe large probabilities of a steep increase when the inventory level is low. In contrast, the transition matrix approximates a stair for large inventories. On the right hand side of Figure 2a one can see that the steady state distribution of the inventory level is highly concentrated in a small

subset of relatively small inventories. Figure 2b(i) represents the transition matrix for an (r, Q) policy with $r = 1$ and $Q = 9$. Observe the similar structure for small and large inventories. On the right hand side, one can notice a subtle difference between the steady state distributions, with that for the (r, Q) policy tending to be more spread over the state space. Both representations are for the supply called "Intravenous Catheter", with maximum back-order of eight units (the maximum back-order allowed is equal the max demand in one period).



(a) (i)Transition Matrix for (s, S) Policy (ii) Inventory Level Distribution



(b) (i)Transition Matrix for (r, Q) Policy (ii) Inventory Level Distribution

Figura 2: Graphic Scheme of a given material and their respective Transition Matrix and IL Distribution

As previously stated, the right side of Figures 2a(ii) and 2b(ii)) represents the inventory level distributions, already described in section 3.2.2. One can see in both figures that the distribution resembles a classical uniform distribution. The second important thing to state is the vertical red line in the position $X = 0$ (Note that the back-order level is 8. Hence, the state 8 refers to an empty inventory). By recalling that $S_3 = P(X \geq 0)$, one can easily determine the service level for any parametric policy.

4.3. Safety Stock

Following the model described throughout section 3.2, there will be four models to be evaluated, namely the *Reorder Point Algorithm* and *Markov Decision Process*, for both ordering policies ((s, S) and (r, Q)). Table 2 summarizes the results for four different supplies. The output of the model is the respective optimal parametric policy, which represents the actions that must be executed at each decision epoch. The set of all optimal actions, one for each inventory level, comprises the optimal ordering policy for each material. The ordering policies presented should be adopted by the decision maker in order to minimize not only the average inventory costs, but also to keep a minimum Service Level (S_3) of at least 0.99, i.e., 99%. The Service Level was

arbitrarily chosen, however it was selected to be very large and therefore consistent with healthcare applications.

Tabela 2: Ordering Policies obtained through Reorder Point Algorithm

	Ini. Policy	Reorder Point Algorithm			MDP Model		
Code/Policy	π	π	S_3	η_π	π	S_3	η_π
03360 - (s, S)	(1, 5)	(8, 17)	0.9983	19.54	(1, 3)	0.9957	3.23
03360 - (r, Q)	(1, 5)	(7, 17)	0.9938	24.65	(0, 2)	0.9931	3.82
52293 - (s, S)	(1, 9)	(4, 8)	0.9946	13.43	(3, 11)	0.9934	10.53
52293 - (r, Q)	(1, 9)	(3, 8)	0.9944	12.91	(2, 10)	0.9997	13.08
51848 - (s, S)	(1, 3)	(0, 3)	0.9908	2.67	(1, 2)	0.9998	2.50
51848 - (r, Q)	(1, 3)	(0, 3)	0.9955	4.23	(0, 1)	0.9931	2.16
50842 - (s, S)	(1, 9)	(3, 7)	0.9952	8.56	(3, 7)	0.9952	8.56
50842 - (r, Q)	(1, 9)	(2, 7)	0.9949	10.60	(2, 7)	0.9949	10.60

Is important to note that this algorithm requires an initial policy to start, and this guess can dramatically change the final ordering policy. Once the service level (S_3) is guaranteed by the model, the long-run average cost (η_π) may produce very poor results if compared to MDP Model, as it can be notice for code "03360". MDP Models often outperform the Reorder Point Algorithm (RPA) as expected, however for code "52293 - (r, Q)", the Reorder Point Algorithm (RPA) produces a slightly better result (η_π). That happened because the MDP algorithm was trapped in a sub optimal region with very high service level (S_3 = 0.9997). Perhaps a heuristic search could be useful here to avoid that.

Next, is possible to note that (s, S) policies are optimal and outperform equivalent (r, Q) policies, as stated in section 3.2.1. That can be verified for code "51848", in which case the RPA produces similar bounds for both (s, S) and (r, Q) policies. However, the (s, S) policy yields a lower cost. Lastly, is possible to get exactly the same results for both models as seen for code "50842"when solved by either the RPA or an MDP algorithm.

5. Concluding remarks and further research

This work presented an inventory management model for medical supplies with stochastic demand. It is up to the manager to choose the inventory model. He/she can select either a Markov Decision Process, which produces better results in terms of cost or a Reorder Point Algorithm for the sake of simplicity. He/she can also choose an ordering policy (s, S), which is optimal or opt for a (r, Q) policy for the sake of convenience. Whichever model you choose the goal is to minimize the average operational costs, once a minimum service level is already defined. Moreover, through the results presented in section 4.3, it is possible to conclude that modeling the problem as an MDP is worth the effort. Furthermore, and (s, S) policy should be the default policy in any case, except when the amount to be purchased must be fixed, in which case a (r, Q) policy is more convenient.

The paper contributes with the inventory control literature of hospital supplies by introducing a decision making model for inventory management that uses Markov Decision Process and a minimal Service Level. The MDP allows for a simple treatment of long-term operations while also enabling a simple, easy to implement solution strategy based solely on the storage level. Moreover, the service level plays a critical role to ensure the proper functioning of the hospital.

The model presented in this paper can be improved. A heuristic search to find a near-optimal initial policy can ensure better results for Reorder Point Algorithm. Other heuristic can be

https://proceedings.science/p/106772?lang=pt-br

developed improve the search of a shortage cost that allows the MDP model to reach the desired service level while also optimizing performance. It can also be useful to decrease the computation time. A sensitive analysis of the ordering cost is also recommended. Finally, no lead-time was considered in this study. However, in order for us to properly account for the reality of Brazilian's public hospitals, a stochastic lead-time should be considered.

Referências

S. Ahiska, S. Appaji, R. E. King, and D. P. Warsing. A Markov decision process-based policy characterization approach for a stochastic inventory control problem with unreliable sourcing. *International Journal of Production Economics*, 144:485–496, 2013.

Y. Aneja and A. H.Noori. The optimality of (s, S) policies for a stochastic inventory problem with proportional and lump-sum penalty cost. *Management Science*, 33:750–755, 1987.

S. Axsäter. *Inventory Control*. Springer, New York NY, 3 edition, 2015.

F. J. Beier. The management of the supply chain for hospital pharmacies: A focus on inventory management practices. *Journal of Business Logistics*, 16:153–177, 1995.

L. Benkherouf and P. Suresh. Optimality of (s,S) policies for a stochastic inventory model with proportional and lump-sum shortage costs. *Operations Research Letters*, 38:252–255, 2010.

S. Cetinkay and M. Parlar. Optimality conditions for an (s,S) policy with proportional and lump-sum penalty cost. *Management Science*, 48:1635–1639, 2002.

I. Civelek, I. Karaesmen, and A. Scheller-Wolf. Blood platelet inventory management with protection levels. *European Journal of Operational Research*, 243:826–838, 2015.

M. E. Dallora and A. Forster. Gerenciamento custos de material de consumo em um hospital de ensino. *RAS*, 15:46–52, 2013.

A. Federgruen and Y. Zheng. An efficient algorithm for computing an optimal (R,Q) policy in continuous review stochastic inventory systems. *Oper. Res.*, 40:808–813, 1992.

G. O. Ferreira, E. F. Arruda, and L. G. Marujo. Inventory management of perishable items in long-term humanitarian operations using Markov Decision Processes. *International Journal of Disaster Risk Reduction*, 31:460–469, 2018.

S. G. Johansen and A. Thorstenson. Emergency orders in the periodic-review inventory system with fixed ordering costs and compound Poisson demand. *International Journal of Production Economics*, 157:147–157, 2014.

C. Kouki, Z. Jemaï, and S. Minner. A lost sales (R, Q) inventory control model for perishables with fixed lifetime and lead time. *International Journal of Production Economics*, 168:143–157, 2015.

S. Minner and E. Silver. Replenishment policies for multiple products with compound-poisson demand that share a common warehouse. *International Journal of Production Economics*, 108: 388–398, 2007.

K. P. Mubiru. Joint replenishment problem in drug inventory management of pharmacies under stochastic demand. *Brazilian Journal of Operations & Production Management*, 15:302–310, 2018.

C. A. Minou O. Keizer, H. T. Ruud, and J. Veldman. Joint condition-based maintenance and inventory optimization for systems with multiple components. *European Journal of Operational Research*, 257:209–222, 2017.

E. Porras and R. Dekker. An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods. *European Journal of Operational Research*, 184: 101–132, 2008.

Portal Comprasnet, 2018. URL `http://comprasnet.gov.br/ConsultaLicitacoes/download/download_editais_detalhe.asp?coduasg=250059\&modprp=5\&numprp=502018`. Accessed: 2019-01-24.

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

M. L. Puterman. *Markov Decision Processes*. Wiley, Hoboken NJ, 1994.

H. Scarf. The optimality of (s, S) policies in the dynamic inventory problem. Applied Mathematics and Statistics Laboratory 11, Stanford University, 1959.

H. L. F. Soares, L. Bahiense, and E. F. Arruda. Otimização e controle da oferta de bolsas de sangue em centros hemoterápicos por meio de ônibus itinerantes: o caso do hemorio. In *Anais do XLIX Simpósio Brasileiro de Pesquisa Operacional*, pages 4215–4224, Rio de Janeiro, 2017. SOBRAPO.

A. Syntetos and J. Boylan. The accuracy of intermittent demand estimates. *International Journal of Forecasting*, 21:303–314, 2005.

L. Turrini and J. Meissner. Spare parts inventory management: New evidence from distribution fitting. *European Journal of Operational Research*, 273:118–130, 2018.

J. Volland, A. Fügener, J. Schoenfelder, and J. O. Brunner. Material logistics in hospitals: A literature review. *Omega (United Kingdom)*, 69:82–101, 2017.

Y. Zheng and A. Federgruen. Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Operations Research*, 39:654–665, 1991.