

SPARE PARTS CONSUMPTION FORECASTING USING A HIERARCHICAL BAYESIAN MODEL

Marcello Dantas Gomes Júnior^{1*} and Pauli Adriano de Almada Garcia²

Received February 21, 2019 / Accepted October 21, 2019

ABSTRACT. One of industrial companies' challenges, especially for intensive-use plants and other assets, is the proper sizing of the stock of strategic spare parts, items that have a history of low consumption, but whose lack can cause delays in repair and maintenance services, at the extreme leading to operational shutdowns. Effects can be from small to large scales. While on one hand, having a large stock of strategic items can provide a greater guarantee of operational availability, on the other hand, it brings additional storage and preservation costs, in addition to fixed capital outlays. A compromise solution is needed. The use of traditional or simpler techniques to infer the ideal level of stock for each spare often suffers from lack of historical data, especially in installations in the initial phase of the operation and maintenance cycle. Another problem is the diversity of applications for some materials. The present work proposes a method based on reliability and Bayesian hierarchical models (HBMs) to overcome the problems of data scarcity, uncertainties and variability between applications of each spare part. The criticality of the equipment or assets in which the spare parts are applied is taken into account in the method. The hierarchical Bayesian model enables updating information as new consumption of strategic items is registered. The method is tested for a stationary offshore oil and gas unit.

Keywords: inventory management, strategic spares parts, reliability, hierarchical bayesian modeling.

1 INTRODUCTION

One of the challenges of industrial companies', especially for plants and physical assets subject to intensive use, is the adequate sizing of the stock of strategic spares, items that have a history of low consumption, but whose lack can cause delays in repair and maintenance services, leading to operational shutdown (Mor et al., 2021; Poppe et al., 2017; Sathish et al., 2019; van der Auwer-aer et al., 2019). The effects range from small to large scales. While on one hand, having a large stock of strategic items can bring greater guarantee of operational availability, on the other hand

*Corresponding author

¹Petróleo Brasileiro S.A., Avenida Henrique Valadares, 28. Centro, Rio de Janeiro, RJ, Brazil – E-mail: marcello.dantas@petrobras.com.br – <https://orcid.org/0000-0001-5875-8816>

²Universidade Federal Fluminense, Rua Desembargador Ellis Hermydio Figueira, 783, Aterrado, Volta Redonda, RJ, Brazil – E-mail: pauliadriano@id.uff.br – <https://orcid.org/0000-0002-3934-4332>

it causes additional storage and preservation costs, in addition to fixed capital outlays. Hence, a compromise solution is needed (Costantino et al., 2018; Tapia-Ubeda et al., 2019; Turrini & Meissner, 2019). The use of traditional or simpler techniques to infer the ideal level of stock for each spare part often faces lack of historical data, especially in installations in the initial phase of the operation and maintenance cycle (Farhat et al., 2018; Hu et al., 2018; Kazemi, Zanjani & Nourelfath, 2014; Rezaei et al., 2018). Another problem is the diversity of applications of some materials. Here we proposes a method based on reliability and Bayesian hierarchical models (HBMs) to overcome the problem of data scarcity, uncertainties and the variability of the applications of each spare. The problem regarding uncertain demand or consumption of spare parts is the limited amount of data available to build models and make predictions for future consumption in time windows. It is crucial that these spare parts form a strategic stock in a quantity that ensures a good level of equipment availability in industrial plants, while avoiding excessive inventory that would minimize immobilized capital. This trade-off is quite challenging (Hu et al., 2018). Determining appropriate levels of strategic spare parts inventory is possible through simulations of reliability, availability, and maintainability, as well as life cycle cost analysis. Although these methods are robust, industrial experience has shown that they are not scalable for many materials to be analyzed. Thus, this study aims to address two gaps: scarcity of historical data on strategic spare parts consumption and the difficulty of scaling solutions for a large quantity and variety of strategic spare parts. The Table 1 summarizes the current scenario and the desired outcome through the methodological procedure proposed in this study.

Table 1 – Comparison of MLE, Bayesian MLE, and HBM.

AS IS	TO BE
1. Forecasting consumption entirely dependent on the opinion or experience of experts.	1. Data-driven consumption forecasting with less reliance on expert opinions.
2. Impossibility of good predictions due to scarcity of data.	2 Feasibility of good predictions even with scarce data.
3. Little cohesion between stocks for the same equipment or plants.	3. Greater cohesion between stocks for the same equipment or plants.
4. Infeasibility of carrying out stock sizing of strategic spare parts in bulk.	4. Feasibility of carrying out stock sizing of strategic spare parts in bulk, through a scalable method.

Source: Produced by the authors.

Offshore industrial facilities have an additional challenge to the formation of strategic stock items of their physical assets: the fact that they operate on the high seas (Gomes Junior, 2020; Gomes Junior & Rocha Filho, 2021). This operating environment imposes more severe consequences for poorly sized strategic inventory. The warehouses of stationary offshore production units have limited space, which brings limitations in the number of spare parts. However, the absence of a

strategic spares on board can cause huge safety and operational disruptions. One should not be overly conservative or optimistic involving stock sizing (van Horenbeek et al., 2013).

Another major problem, not unique to the offshore industry, is the scarcity of data or information (Cavalieri et al., 2008; Mirzahosseini & Piplani, 2011; Romeijnnders et al., 2012). This scarcity can be due to the fact that the plant is new or has just entered the operation and maintenance phase. Another possible factor is the culture of not accurately recording information related to maintenance, a problem that is aggravated in industries with intensively used assets. Asset-intensive industries operate 24 hours a day, seven days a week, with the exception of production downtime. So, the need to correct record-keeping problems can overwhelm teams at all hierarchical levels, making it difficult to structure processes that bring operational improvements and greater security. One process that can benefit from a culture of good records is the management of strategic spares (Hardwick & Lafraia, 2021). However, changing a culture at various hierarchical levels is a time-consuming process, and while fundamental, it takes a long time to work smoothly. It is therefore necessary to develop methods to work in this environment of scarcity of data, generating results that are neither excessively conservative nor optimistic. Here, we propose an initial approach using reliability and a hierarchical Bayesian model, using information from the consumption history of strategic spares, with the objective of dimensioning the stock of these spares for a future time window.

The text is structured as follows: this section is the introduction, where the problem of strategic spare parts inventory is contextualized. Section 2 briefly describes the theoretical foundation of hierarchical Bayesian modeling for reliability. Section 3 addresses our proposed method. Section 4 presents a case study, where the proposed method is applied to a specific class of physical assets. Finally, section 5 presents final discussions, conclusions and opportunities for future developments.

2 HIERARCHICAL BAYESIAN MODELING (HBM)

The larger the data sample, the closer the probability estimate asymptotically approaches the true value, becoming less biased. However, this is not true in the case of small datasets and/or rare events. The event tree approach requires knowing the probability of occurrence of the initial event and the failure of safety barriers. But the failure probabilities of safety barriers are difficult to obtain for scarce events, such as the need to replace strategic spares. So, to estimate the parameters of the statistical distributions to infer the size of the strategic stock, it is not recommended to use the traditional maximum likelihood technique (MLE). Here, the probability update is one-way, that is, from the initial event it is possible to calculate the probability of each consequence, but the opposite does not hold. The event tree is also a static model. One way out of this impasse is to convert the fault tree or event tree into a Bayesian network, where the quantitative relationships between nodes are represented in conditional probability tables (CPT). Thus, the updating of the probabilities can be done in both directions. To model the Bayesian network, it is necessary to know the a priori distribution of the root nodes and the conditional probability tables (CPT), which requires data or expert opinions. The problem is that in both cases (data and

experts), uncertainties are not taken into account, nor is source-to-source variability. To address these problems, hierarchical Bayesian modeling (HBM) is used.

A level is added to the parameters of a base distribution, which are inferred from a prior, which in turn is modeled by hyperparameters, given by a non-informative distribution. Then, from a non-informative distribution, the hyperparameters of the non-informative prior are generated. Uninformative priors are then used so that the uncertainty of the data is captured in the unbiased updated a posteriori distribution, called the informative posterior. But to obtain it, the data are first incorporated into a likelihood function that adequately describes the behavior of the data, which are combined with a non-informative prior through Bayes' theorem to obtain an informative posterior estimate of the model's parameters. Then the expected values of the parameters are calculated. These expected values are fed back to the base distributions to predict the number of spares that will be needed in the next time interval. The use of HBM has major advantages: the interdependence between FTA and ETA parameters is fully exploited in likelihood; and joint likelihood eliminates the need for CPTs and the resulting uncertainties of two-way analysis and source-to-source variability (Yu et al., 2017).

There are two ways to estimate parameters from a data sample: the traditional MLE and the Bayesian MLE. Traditional MLE takes the parameters that maximize likelihood based on the data and the PDF. On the other hand, the Bayesian MLE has candidates for probability distribution whose parameters are sampled from a prior distribution, where the parameters come from maximizing the joint likelihood of the sampled data and the worst item. The following figure shows that Bayesian MLE is better because it considers the entire sample space of the parameter θ . By the traditional MLE, there is only one distribution of θ that can explain the data D . However, in these two cases, the data generated are considered to come from very specific scenarios, which is not in line with reality, especially in the case of source-to-source variability.




Figure 1 – Comparison of MLE, Bayesian MLE, and HBM.

Source: Adapted from (Yu et al., 2017).

Then the two-stage HBM is used. Like the Bayesian MLE, the HBM models the uncertainty of the sampled data in the parameter prior θ , which is sampled from hyperparameters of a non-informative distribution. In Figure 1.c above, $p(\theta|\alpha, \beta)$ is a prior of θ , while α and β are hyperparameters sampled by a non-informative distribution $P(A, B)$, which allows a wide range of values for α and β , so that $p(\theta|\alpha, \beta)$ is general enough to capture source-to-source variability. In the update phase, non-informative priors have no strong influence on the posteriors. So, the

Bayesian updating is completely dependent on the data. Thus, a posterior distribution captures the uncertainty of the data and more faithfully reflects the true value. Using a uniform uninformative prior brings the problem of having to work within an interval, which causes bias and invariability in case of reparameterization. To satisfy Jeffrey's rule, a fuzzy range can be used as an uninformative prior to model the uncertainty of the data. Uninformative priors can be updated using new data, to generate a posterior version of the data of interest, in three steps:

- a) The likelihood of the hyperparameters is obtained by calculating the expected likelihood value of the parameter of interest:

$$l(D|\alpha, \beta) = \int l(D|\theta) \cdot p(\theta|\alpha, \beta) \partial \theta \quad (1)$$

- b) The objective is to obtain a likelihood value that relates the field data regarding the spare parts with the hyperparameters. This makes use of the numerical knowledge of a physical asset to make inferences about another similar physical asset. So now we obtain the posterior distribution of the hyperparameters, through Bayes' theorem:

$$p(\alpha, \beta|D) = \frac{p(\alpha, \beta) \cdot l(D|\alpha, \beta)}{\iint p(\alpha, \beta) \cdot l(D|\alpha, \beta) \partial \alpha \partial \beta} \quad (2)$$

The denominator will need treatment via MCMC to sample α and β from the joint distribution $p(\alpha, \beta)$. Then, the integrals are approximated by the sample mean of $l(D|\alpha, \beta)$.

- c) In the third step, by marginalizing the parameters α and β , the posterior distribution of θ , given the data, can be obtained. It is also known as the informative posterior of θ . The double integral can be approximated with MCMC by the sample mean of $p(\theta|\alpha, \beta)$.

$$p(\theta, D) = \iint p(\theta|\alpha, \beta) \cdot p(\alpha, \beta|D) \partial \alpha \partial \beta \quad (3)$$

In this way, updating with the arrival of new data can be done by:

$$p(\theta|D', D) = \frac{p(\theta|D) \cdot p(D'|\theta)}{\int p(\theta|D) \cdot p(D'|\theta) \partial \theta} \quad (4)$$

Where $p(D'|\theta)$ captures new data in every time interval t .




Figure 2 – Graphic illustration of the HBM.

Source: Adapted from Yu et al. (2017).

3 METHODOLOGY

The research was fundamentally conducted following the following steps in Figure 3.




Figure 3 – Steps of the proposed methodological procedure.

Source: Produced by the authors.

The first three steps correspond to the collection of field data related to the use of strategic spare parts. The fourth stage is composed of Bayesian hierarchical modeling for the consumption behavior of each spare part in each installation. The last step is, from the lead time of each spare part and its consumption history, to use the hierarchical Bayesian model to infer the necessary strategic stock level. The future time window used corresponds to the time it takes for each spare part to be acquired and made available at the facility, to minimize the possibility of a shortage of spare parts and minimize the inventory level.

The Bayesian Hierarchical Model proposed here is structured in the following steps:

- a) Definition of the likelihood functions, first-stage and second-stage hyper-priors.
- b) Inference of the parameters of the predictive distribution.

- c) Inference of the required quantity of each spare in a future time window.

The objective is to predict the amount of a specific spare that is expected to be consumed in a time window at each installation. Then the data collection results are used to estimate the consumption of each spare for each platform in a past time window (Kelly & Smith, 2011). The next step is to define the likelihood distribution to be used. What is known from field data is the amount consumed or used of each spare per platform in a certain time period. We want to estimate the amount of each spare to be consumed by that platform in a future time interval. So the likelihood function that best suits this type of problem is the Poisson function. The parameter to be inferred is the consumption rate λ of each spare per platform. Equation 5 shows the Poisson distribution:

$$P[N(t) = X] = \frac{e^{-\lambda t} \cdot (\lambda t)^x}{x!} \quad (5)$$

Where x is the amount of spares to be consumed in a time window t . The random variable x follows the Poisson distribution.

Each spare is consumed on multiple platforms. To correlate the consumption profile of each spare among the platforms, a first-stage hyper prior is used. Without loss of generality, in the present work we use the gamma distribution, which is combined with the Poisson distribution. This takes advantage of the knowledge of all other platforms to infer the consumption of each spare part on a specific platform. This mitigates the negative effects of the scarcity of data from just one platform and captures the uncertainties and variability between the consumption of each spare on each platform.

$$p(\lambda|\alpha, \beta) = \frac{\beta^\lambda \cdot \lambda^{\alpha-1} \cdot e^{-\beta\lambda}}{\Gamma(\alpha)} \quad (6)$$

The hyper prior distribution is also described by parameters. To correlate these parameters to all consumption data for each spare, it is necessary to use a second-stage hyper prior distribution. This is responsible for inferring the parameters of the first-stage hyper prior. In order not to bias the first-stage a priori distribution, the second stage hyper prior is a fuzzy distribution. The proposed method uses a diffuse gamma, also known as Jeffrey's prior, according to Equation 7 and Equation 8. Both the parameters α and β of the first-stage priors follow Jeffrey's priors.

$$\alpha \sim \Gamma(0.00001, 0.00001) \quad (7)$$

$$\beta \sim \Gamma(0.00001, 0.00001) \quad (8)$$

The third step is to obtain the posterior distribution of the parameter λ of the Poisson distribution. It is also known as the informative posterior of λ . The double integral can be approximated with MCMC by the sample mean of $p(\lambda|\alpha, \beta)$, where D is the observed consumption data of the spare parts.

$$p(\lambda|D) = \iint p(D|\alpha, \beta) \cdot p(\alpha, \beta|D) \partial\alpha \partial\beta \quad (9)$$

Finally, an informative posterior is obtained to estimate the amount of each spare to be consumed in a future time window.

$$f(x|\lambda, t) = f(\lambda|D) \cdot f(\alpha, \beta) \quad (10)$$

In other words, the value x of the Poisson distribution is sought, where the consumption rate λ of each spare has already been inferred in the hierarchical Bayesian model (HBM). A credibility level must be chosen to estimate the value of x .

$$P[x] = \frac{e^{-\lambda t} \cdot (\lambda t)^x}{x!} \quad (11)$$

4 CASE STUDY

The definition of the physical asset targeted by the present work considers its criticality for operational safety or operational continuity, or both. Physical assets have a bill of materials (BOM), which is nothing more than a list of parts and items that make up the equipment. Some of them will not be replaced during the lifetime of the equipment, so they have been excluded from the study. Other items are considered consumables and their consumption forecast is already well defined, so they have also been excluded from the study.

This case study applies the method presented in the previous section to the centrifugal fire-fighting pump. These pumps are installed in nine stationary oil and gas production units. The spares list has 58 items. Among them, we chose a specific pump gasket. A survey was performed of how many of these gaskets were consumed in a time window of each platform. The data are reported in Table 2 and are the model's input data. We performed 200,000 Monte Carlo simulations with two Markov chains. The future time window for forecasting consumption of gaskets was 5 years. An alternative would be to use the lead time of each spare as the future time window. Two values were defined for the hyperparameter α and two for the hyperparameter β just to initialize the two Markov chains.

After the simulations, the results of the hierarchical Bayesian model proposed were obtained. Figure 4 shows the probability density functions of the first-stage prior hyperparameters:




Figure 4 – First-stage prior hyperparameters

Source: Produced by the authors.

Table 2 – Gasket consumption history of each pump.

Pump	Number of gaskets consumed	Pump age (years)
A	1	5.8
B	1	3.2
C	1	2.5
D	6	7.7
E	0	1
F	8	14.5
G	3	9.1
H	3	4
I	1	8

Source: Produced by the authors.

Figure 5 shows the probability density function (pdf) of the inferred consumption rates for this specific pump gasket in each stationary production unit. A similar consumption rate profile can be seen between the units. It is necessary to choose a credibility interval to define a consumption rate value. For example, when choosing a 95% credibility interval, the information is extracted that the consumption rate that defines the 95% limit of the pdf graph area. So, it can be said that, with 95% credibility, the consumption rate will be the value λ_{95} obtained from the graph. The parameter λ of each unit was inferred considering the consumption data of the unit in question, also considering the influence of the consumption profile of the other units. In other words, to infer the consumption rate of each unit, we took advantage of our knowledge about the consumption of a specific pump gasket in the other units.

With the information of the consumption rate λ in hand, a predictive posterior distribution is used to estimate the number of gaskets that will be consumed on each platform in the next 5 years. This estimate is also presented in the form of a probability function, as it is a discrete random variable. Such estimates for each stationary production unit are shown in Figure 6. The reasoning regarding the level of credibility of the estimate is the same for the definition of the consumption rate λ .

Table 3 lists the consumption estimates of gaskets on each platform over the next 5 years, considering various levels of credibility. We point out that some stationary production units do not have a history of consumption of this spare, but using the knowledge of the consumption profile in the other units, it is possible to establish an estimate for such units as well.

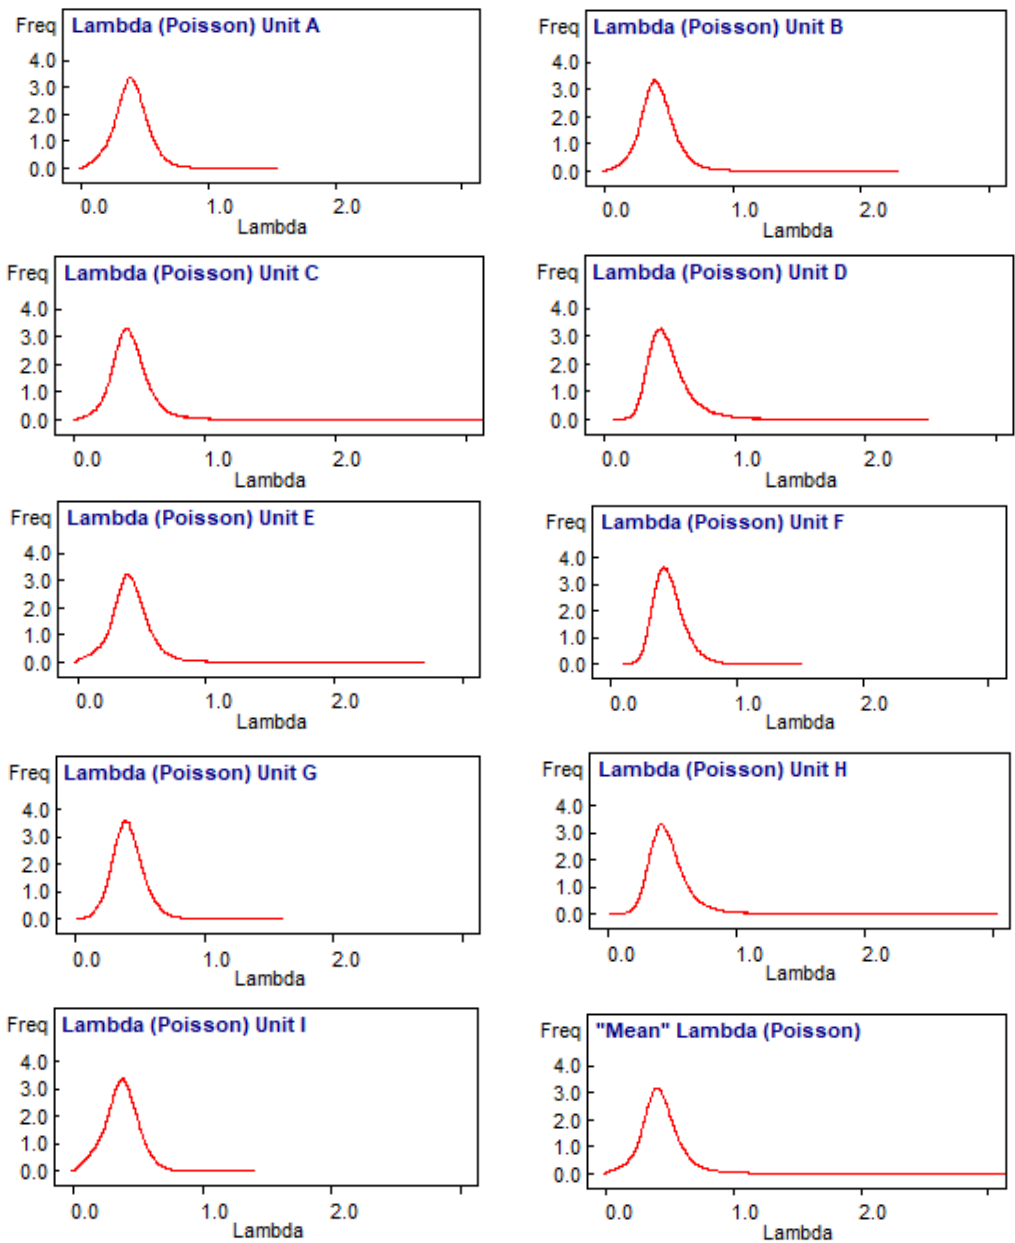


Figure 5 – Lambda parameter inferred from each unit.

Source: Produced by the authors.

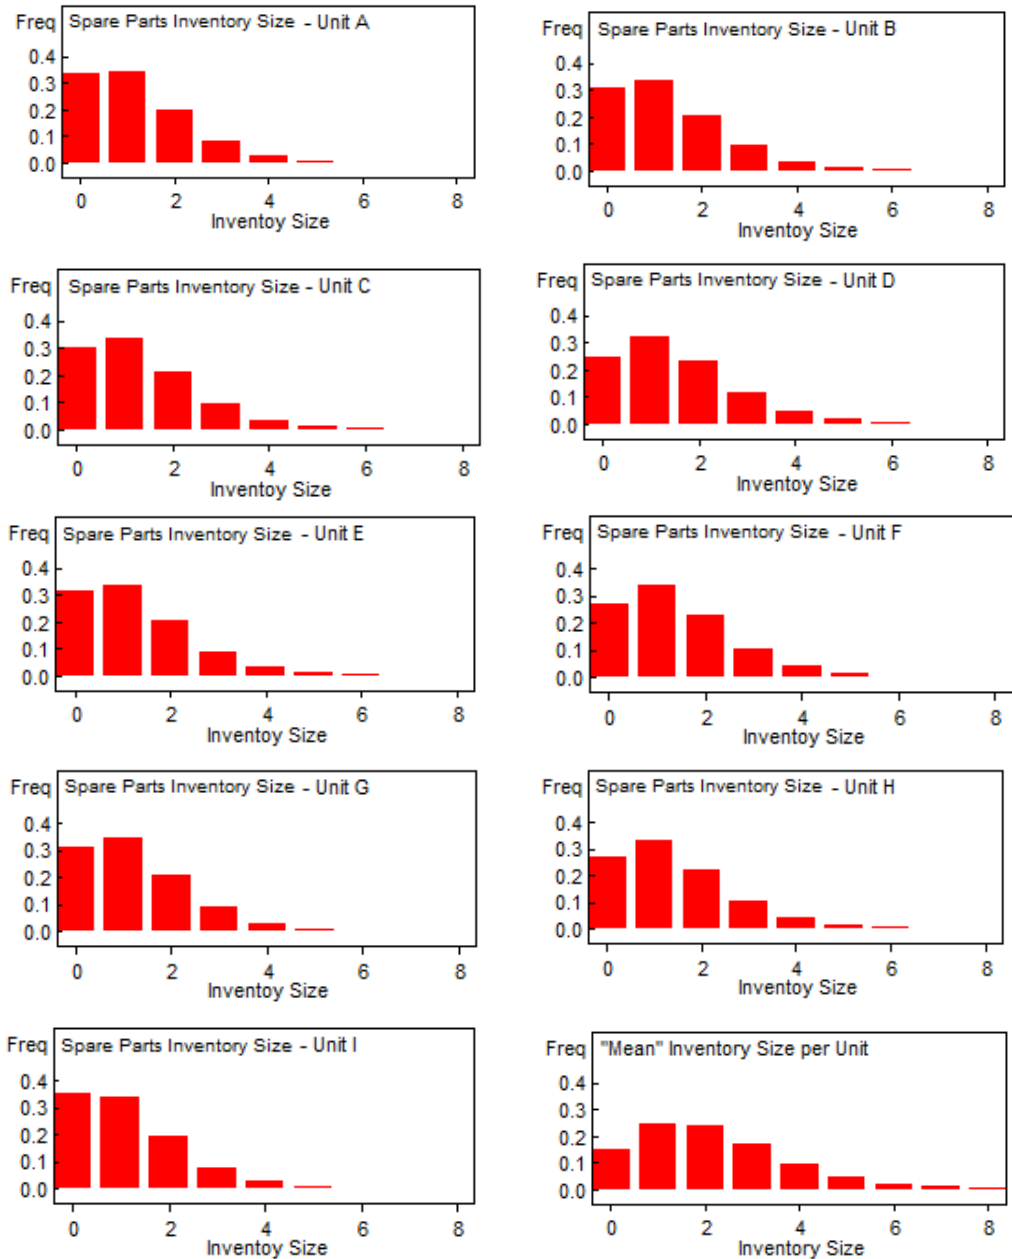


Figure 6 – Inventory size of each unit.

Source: Produced by the authors.

The last row of Table 3 shows an average size, by platform, of the fire pump's gasket stock. If the total inventory is measured based on the average value per platform instead of the individual value per platform, a more conservative estimate will be obtained, that is, a greater inventory of this spare. This is because the average value per platform encompasses all the uncertainties of the nine stationary production units.

Table 3 – Level of inventory size.

UNIT	Credibility level for inventory size				
	Median	75%	90%	95%	97.5%
Unit A	1	2	3	3	4
Unit B	1	2	3	4	4
Unit C	1	2	3	4	4
Unit D	1	2	3	4	5
Unit E	1	2	3	4	4
Unit F	1	2	3	4	4
Unit G	1	2	3	3	4
Unit H	1	2	3	4	4
Unit I	1	2	3	3	4
“Average” unit	2	3	4	5	6

Source: Produced by the authors.

5 CONCLUSIONS AND OPPORTUNITIES FOR FUTURE RESEARCH

The results show the capacity of the hierarchical Bayesian model to support analysis and decision making in data scarcity scenarios. Platforms without consumption history can have their inventory of strategic items scaled using information from similar facilities. From our experience, there is a tendency to oversize the strategic stock through the opinion of specialists such as maintenance engineers or operators. On the other hand, there is a tendency to undersize on the part of those who manage inventories. Another interesting point is the possibility of sizing the stock of each spare part based on the level of credibility, as seen in Table 3. The higher the level of credibility, the more conservative is the dimensioning of the stock, but with a greater degree of certainty that this level will meet the needs of maintenance and operation. For example, if the spare part in question belongs to a device with maximum criticality within the classification adopted by the company, its stock level can be dimensioned with a credibility of 97.5%. On the other hand, if this spare part only serves equipment with the least possible criticality, its stock level can be scaled with a credibility of 50% or median. According to the phase of the life cycle of the unit, the level of stock of the same spare can be different depending on where it is being

used. The hierarchical Bayesian model can provide this information. The present work demonstrates a very simple Bayesian approach to strategic inventory sizing. As a possibility for future developments, dynamic Bayesian networks can be adopted in conjunction with this method. The equipment failure rate information will update the consumption rate of each spare part in an interdependent manner, making the estimates closer to operational reality.

Another application would be to use zero inflated models in the likelihood distribution of the Bayesian model, to deal with even rarer events of consumption of spare parts.

References

- CAVALIERI S, GARETTI M, MACCHI M & PINTO R. 2008. A decision-making framework for managing maintenance spare parts. *Production Planning & Control*, **19**(4): 379–396. Available at: <https://doi.org/10.1080/09537280802034471>
- COSTANTINO F, DI GRAVIO G, PATRIARCA R & PETRELLA L. 2018. Spare parts management for irregular demand items. *Omega*, **81**: 57–66. Available at: <https://doi.org/10.1016/j.omega.2017.09.009>
- FARHAT M, AKBALIK A, SAUER N & HADJ-ALOUANE AB. 2018. *Inventory replenishment planning with a full truckload cost and a full return policy*. **511**(1): 259–263. Available at: <https://doi.org/10.1016/j.ifacol.2018.08.286>
- GOMES JUNIOR MD. 2020. *Deactivating a motor generator set in a stationary unit for the production of offshore oil in decommissioning using reliability, availability and maintainability analysis and life cycle cost*. **20**. Available at: <https://doi.org/10.48072/2525-7579.rog.2020.248>
- GOMES JUNIOR MD & ROCHA FILHO LB. 2021. ESTUDO DE CONFIABILIDADE DA PROPOSTA DE REORGANIZAÇÃO DAS VOTAÇÕES DE DETECTORES DE GÁS DOS FPSOS DA CESSÃO ONEROSA. *Congresso Brasileiro de Manutenção e Gestão de Ativos*. Rio de Janeiro Brazil.
- HARDWICK J & LAFRAIA JRB. 2021. *Living Asset Management Maturity*.
- HU Q, BOYLAN JE, CHEN H & LABIB A. 2018. OR in spare parts management: A review. *European Journal of Operational Research*, **266**(2). Available at: <https://doi.org/10.1016/j.ejor.2017.07.058>
- KAZEMI ZANJANI M & NOURELFATH M. 2014. Integrated spare parts logistics and operations planning for maintenance service provider. *International Journal of Production Economics*, **158**. Available at: <https://doi.org/10.1016/j.ijpe.2014.07.012>
- KELLY D & SMITH C. 2011. *Bayesian Inference for Probabilistic Risk Assessment*. Springer. Available at: <https://doi.org/10.1007/978-1-84996-187-5>
- MIRZAHOSSEINIAN H & PIPLANI R. 2011. A study of repairable parts inventory system operating under performancebased contract. *European Journal of Operational Research*, **214**(2): 256–261. Available at: <https://doi.org/10.1016/j.ejor.2011.04.035>
- MOR RS, BHARDWAJ A, KHARKA V & KHARUB M. 2021. Spare Parts Inventory Management In the Warehouse: A Lean Approach. *International Journal of Industrial Engineering & Production Research*, **32**(2): 1–11. Available at: <https://doi.org/10.22068/ijiepr.32.2.1>

POPPE J, BASTEN RJI, BOUTE RN & LAMBRECHT MR. 2017. Numerical study of inventory management under various maintenance policies. *Reliability Engineering and System Safety*, **168**: 262–273. Available at: <https://doi.org/10.1016/j.ress.2017.06.012>

REZAEI H, BABOLI A, SHAHZAD MK & TONADRE R. 2018. A new methodology to optimize target stock level for unpredictable demand of spare parts: A Case Study. *Business Aircrafts' Industry*, **5111**: 538–543. Available at: <https://doi.org/10.1016/j.ifacol.2018.08.374>

ROMEIJNDERS W, TEUNTER R & VAN JAARSVELD W. 2012. A two-step method for forecasting spare parts demand using information on component repairs. *European Journal of Operational Research*, **2202**: 386–393. Available at: <https://doi.org/10.1016/j.ejor.2012.01.019>

SATHISH N, ANBUUDAYASANKAR SP, DEEPAN M & NARASSIMA MS. 2019. Procurement methodologies to optimize the inventory levels of spare parts. *International Journal of Recent Technology and Engineering*, **83**: 1662–1669. Available at: <https://doi.org/10.35940/ijrte.C4437.098319>

TAPIA-UBEDA FJ, MIRANDA PA, GUTIÉRREZ-JARPA G & DURÁN O. 2019. Supplier selection for spare parts supply chain networks. *IFAC-PapersOnLine*, **5213**: 2237–2242. Available at: <https://doi.org/10.1016/j.ifacol.2019.11.538>

TURRINI L & MEISSNER J. 2019. Spare parts inventory management: New evidence from distribution fitting. *European Journal of Operational Research*, **2731**: 118–130. Available at: <https://doi.org/10.1016/j.ejor.2017.09.039>

VAN DER AUWERAER S, BOUTE RN & SYNTETOS AA. 2019. Forecasting spare part demand with installed base information: A review. *International Journal of Forecasting*, **351**: 181–196. Available at: <https://doi.org/10.1016/j.ijforecast.2018.09.002>

VAN HORENBEEK A, BURÉ J, CATTRYSSE D, PINTELON L & VANSTEENWEGEN P. 2013. Joint maintenance and inventory optimization systems: A review. *International Journal of Production Economics*, **1432**: 499–508. Available at: <https://doi.org/10.1016/j.ijpe.2012.04.001>

YU H, KHAN F & VEITCH B. 2017. A Flexible Hierarchical Bayesian Modeling Technique for Risk Analysis of Major.

How to cite

GOMES JÚNIOR MD & GARCIA PAA. 2023. Spare parts consumption forecasting using a hierarchical Bayesian model. *Pesquisa Operacional*, **43**: 269646. doi: 10.1590/0101-7438.2023.043.00269646.