

Article

Optimizing Maintenance Contract Pricing Through Comprehensive Risk Assessment

Bruno Pereira ^{1,*}, Manuel Cruz ^{2,†} , Jorge Santos ^{2,*},† , Cristina Lopes ^{2,3,†} , Sandra Ramos ^{2,4,†}, Filipa Vieira ^{5,†} and Pedro Louro ^{5,†}

¹ Sonae MCretail, 4464-501 Matosinhos, Portugal

² LEMA-Laboratório de Engenharia Matemática, Instituto Superior de Engenharia do Porto, Instituto Politécnico do Porto, Rua Dr. António Bernardino de Almeida, 4249-015 Porto, Portugal; mbc@isep.ipp.pt (M.C.); cristinalopes@iscap.ipp.pt (C.L.); sfr@isep.ipp.pt (S.R.)

³ CEOS.PP-Centro de Estudos Organizacionais e Sociais do Politécnico do Porto, Instituto Superior de Contabilidade e Administração do Porto, Instituto Politécnico do Porto, Rua Jaime Lopes Amorim, 4465-004 São Mamede de Infesta, Portugal

⁴ CEAUL-Centro de Estatística e Aplicações da Universidade de Lisboa, Campo Grande, 1749-016 Lisboa, Portugal

⁵ Nors Group, S.A., 4149-010 Porto, Portugal; fivieira@nors.com (F.V.); plouro@nors.com (P.L.)

* Correspondence: brunopereira498@gmail.com (B.P.); jms@isep.ipp.pt (J.S.)

† These authors contributed equally to this work.

Abstract

This study develops a risk-informed pricing framework for maintenance contracts in the trucking industry. We apply a comprehensive methodology combining statistical segmentation, cost analysis, and Value at Risk (VaR) modeling to a dataset of nearly 2000 contracts. Contracts were grouped by duration and truck usage, and distributions were fitted to estimate costs and compute risk premiums. Two pricing models are proposed: a traditional VaR-based approach and an adaptive model that incorporates distribution tail heaviness. Results show that the adaptive model resolves the counterintuitive decline in prices at higher risk levels and yields more stable, flexible premiums. These findings underscore the importance of tail-risk metrics in contract pricing to better capture cost uncertainty. The approach supports more accurate risk management and sustainable pricing strategies for maintenance services.

Keywords: pricing strategies; maintenance contracts; value at risk; risk analysis; trucking industry; industrial mathematics

MSC: 91G70



Academic Editor: Anatoliy Swishchuk

Received: 13 March 2026

Revised: 21 April 2026

Accepted: 22 April 2026

Published: 26 April 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

It is common practice for individuals to enroll in warranty or service agreements with a view to safeguard their possessions. This trend is also evident in the automotive industry, where purchasing a vehicle mandates signing an insurance policy, with certain brands offering additional subscription options for maintenance contracts [1].

Ensuring adequate maintenance of trucks is pivotal to guaranteeing their safety, performance, and reliability. Given the strenuous working conditions they endure—encompassing long distances, heavy loads, and diverse weather and road conditions—maintaining trucks in optimal operational condition is imperative. Maintenance contracts have received significant attention due to increased profits for the service providers and reduction in risk for

owners through expert services provided by the service providers. To make maintenance contracts more effective, there is a need to develop mathematical models and understand future costs that could be built into the contract price [2,3].

A maintenance contract is a legal agreement between a service provider and a client that covers the costs of both preventive and corrective maintenance services over a specified period. The client pays a fixed upfront fee in exchange for the service provider's commitment to provide maintenance services during the contract period. Maintenance contracts play a crucial role in ensuring the reliability, safety, and long-term performance of industrial assets, particularly in sectors such as transportation and logistics. In the trucking industry, vehicles operate under demanding conditions, including long distances, heavy loads, and varying environmental factors, making maintenance both essential and costly. As a result, many companies rely on maintenance contracts to transfer operational risk and ensure predictable service costs. Given the unpredictable nature of maintenance costs, determining the appropriate break-even price for such a contract is a significant challenge [4,5]. From a pricing perspective, maintenance contracts share similarities with non-life insurance products, where premiums must reflect both expected costs and underlying uncertainty. Early studies in this area have emphasized the importance of stochastic modeling and cost variability, particularly in high-value assets such as aircraft engines [6]. More recent contributions have extended these approaches by incorporating uncertainty in repair processes, operational conditions, and contract structures [7], as well as learning effects, optimized maintenance policies, and information asymmetry [8]. Furthermore, data-driven approaches have gained increasing attention, enabling more accurate pricing by leveraging historical contract data and predictive analytics [3].

Recent literature [9] has emphasized the importance of data-driven and risk-based approaches for pricing maintenance contracts, highlighting their strong analogy with insurance products where premiums should reflect the underlying risk exposure to mitigate adverse selection.

A key challenge in this context is the incorporation of risk into pricing decisions. Among the various risk measures available, Value at Risk (VaR) has become one of the most widely used due to its simplicity and interpretability [10,11]. However, despite its widespread adoption, VaR has well-documented limitations, particularly its inability to adequately capture extreme tail events and the severity of losses beyond the selected confidence level [12]. These limitations are especially relevant in maintenance cost modeling, where rare but high-impact events may significantly affect profitability.

Alternatives like Expected Shortfall enhance foundational risk measures by accounting for tail loss severity, providing coherence via subadditivity and positive homogeneity to aid risk aggregation and utility maximization in fat-tailed scenarios [13]. Developments such as modified Expected Shortfall for market robustness and the Wang transform [14] for tunable tail sensitivity offer nuanced tail risk tools, validated through backtesting [15]. Despite this, Value at Risk, with its simple quantile-based approach for benchmarking losses, remains useful in pricing maintenance contracts for trucking fleets, supporting efficient cost forecasting and compliance under usage limits.

The primary objective of this work is to develop a risk-informed pricing framework for maintenance contracts that addresses some of those limitations. Using a dataset of nearly 2000 contracts, covering either new and used trucks, we first construct statistically homogeneous cost groups based on contract characteristics. We then apply a VaR-based pricing model and identify a key limitation: the failure of premiums to increase monotonically with higher confidence levels. To overcome this issue, we propose an adaptive pricing model that explicitly incorporates the tail behavior of the cost distribution through a quantitative metric. This approach provides a flexible framework that adjusts pricing formulas according to the underlying distributional properties, leading to more consistent and robust premium structures.

This may be of the utmost importance to transport companies, fleet managers, leasing companies or insurers seeking to forecast the associated costs of automotive/heavy trucks maintenance.

In the following sections an overview of the dataset is presented, followed by the results from applying a pricing strategy based on the Value at Risk. In Section 4, a new mathematical model to overcome the problems arising from the VaR strategy is defined leading to the conclusions presented at the end of this paper.

2. Cost Analysis

In maintenance contracts, the prices set by the service-providing company strive to be competitive in the current market while also aligning with the company's vision regarding potential costs during the contract period. Therefore, it is essential to consider the entire history of previous contracts to establish a value that encompasses all known costs. The costs of a contract typically include preventive and corrective maintenance services, spare parts and corresponding labor hours. For confidentiality reasons, the costs are presented in monetary units (M.U.).

These contracts can be linked to automobile insurance, as the process is quite similar. In non-life insurance, particularly in automobile insurance, policyholders are typically categorized into risk groups, with all members of the same group paying the same premium designated as an "a priori" premium, which is later updated based on their costs during the first year of the contract [16]. This methodology was adapted for the work presented here, creating risk groups, later referred to as cost groups, based on the historical costs of contracts, with the difference being that the established premiums are not updated.

This section conducts an analysis of the costs of 1901 maintenance contracts, related to new and used long-distance trucks, insured by a Portuguese-based industrial group. These trucks are divided into two ranges, D-range and T-range, which are considered differently by the company that owns the contracts, regarding the different characteristics and prices of each range. This initial cost analysis aims to create groups of contracts whose costs are significantly different from each other. All statistical tests performed in this section were conducted using a significance level of 0.05.

2.1. Average Annual Cost

The duration of a maintenance contract plays a crucial role in determining its cost. Generally, longer contracts are expected to incur higher costs for the company compared to shorter ones. This is due to the wear and tear that occurs over time, resulting in more frequent and expensive repairs over the contract period.

In aiming to create heterogeneous cost groups, the initial step involves comparing costs across various contracts. However, this set of contracts exhibits a wide range of durations. Consequently, total cost may not be the most appropriate metric for comparing contract costs. This is because contracts with vastly different durations would naturally incur significantly different costs, making it challenging to establish meaningful cost comparisons.

To address this issue, the average annual cost variable was considered. This variable is obtained by dividing the total cost by the total duration of the contract. This approach allows for the comparison of costs between any two contracts, irrespective of their durations.

2.2. Grouping Similar Cost Contracts

Now, with the possibility of comparing costs significantly between the various contracts, it becomes necessary to divide them into groups. For this purpose, the duration variable was initially considered, and as a first approach, the contracts were divided into duration intervals with a range of 1 year, where, for example, the interval $[2, 3[$ includes all contracts with a duration of 2 years or more but less than 3 years. Figure 1 shows the

boxplots for the average annual cost for each duration interval. Note that the last interval, $[6, 8[$, has a range of 2 years, this is because the contracts of 6 and 7 years were grouped into the same group, since the sample size for these groups were very small.




Figure 1. Boxplots for the mean annual cost per duration interval.

Observing Figure 1, it is noted that contracts with short durations exhibit very high values of average annual costs that tend to decrease until contracts of 3 years, where they begin to increase again. This phenomenon partially contradicts what would be expected from the expected behavior of contract costs, as for longer-duration contracts, a greater number of repairs would normally be expected, especially in the later years, which would lead to higher costs. However, this does not happen in contracts with durations up to 3 years. It is also important to note that not all intervals are of the same size, and therefore, an extreme value in a small set influences that set more than in a larger set; however, it seems necessary to investigate the characteristics of these groups more thoroughly.

After a joint analysis with the maintenance contract specialist, the importance of truck usage was mentioned, i.e., if the truck is new or used at the time of contract signing, in contract costs. Comparing, for example, two 1-year contracts between a new and a used truck, it is natural that the cost behavior is not similar, especially due to the difference in wear between the trucks, which influences the number of times the vehicle goes to the workshop. Thus, it was concluded that a second division of contracts was necessary, this time also considering the usage of trucks.

Before making this division, it is important to understand if, within the intervals, there are indeed significant differences in costs between new and used trucks. The analysis began by examining the distribution of new and used trucks within each interval, as shown in Figure 2.

It is observed that as the duration of the contracts increases, the percentage of used trucks decreases, even becoming null from contracts of 5 years, demonstrating that the company has little interest in entering long-term contracts with used vehicles. In contrast, new trucks exhibit an opposite behavior, being less frequent in contracts of short duration and predominant in contracts of longer durations. This distribution is already an indication that this division is necessary because, as seen in Figure 1, there is a decreasing behavior of costs up to contracts of 3 years, where used trucks predominate, and an increasing behavior after these contracts, where new trucks predominate.

In this situation, the influence that the use of trucks has on contract costs can only be evaluated in the intervals referring to durations of less than 5 years because they are the only ones that have trucks from both usage categories.

Figure 3 presents the boxplots of the average annual costs within the duration intervals, between new and used trucks, and also represents the number of elements in each category. It is noted that in all intervals, the costs associated with used trucks are generally higher

when compared to new trucks, indicating once again that there are differences in costs between these two usage categories. Regarding the number of elements, it is observed that the intervals $[3,4[$ and $[4,5[$ are the ones that concentrate a higher number of contracts. On the other hand, the interval $[0,1[$ is the one with the lowest number of contracts, particularly containing only 5 contracts of new trucks.




Figure 2. Distribution of trucks by usage.




Figure 3. Comparison of average annual costs between new and used trucks, by interval.

To understand with a certain degree of certainty if there are indeed significant differences between the two categories within each interval, it was considered to initially apply the *t*-test for independent samples. This test requires the data to satisfy the normality condition, so the Shapiro–Wilk test [17] was applied to the data represented in Figure 4.




Figure 4. Histograms of the average annual cost for new and used trucks.

In all cases, the null hypothesis of the test was rejected, indicating that the data do not come from a normal distribution and therefore it is not advisable to apply the *t*-test for independent samples. As an alternative, the Mann–Whitney U test [18] was then used, which has the same purpose as the parametric test but does not assume normality assumptions. For each interval, the test indicated that there are significant differences in costs between new and used trucks, confirming that a second division taking into account truck usage is necessary.

2.3. New Trucks

As shown in Figure 2, new trucks are present in all duration intervals, albeit with different expressions. After this second division, it becomes important to look at the size of the resulting groups in order to compare possible disparities. It was concluded that the group $[0, 1[$, due to having only 5 elements, did not have sufficient representation to

continue in the analysis, and for that reason, it was excluded. Also, when dividing into the T and D ranges, it was found that the D range in new vehicles is not represented in all groups, and when it is, it is considerably low. Therefore, only the T range trucks were considered for the new trucks.

After all the changes, it remains to analyze the behavior of the new trucks in the T range. As depicted in Figure 5, these trucks appear to follow an increasing trend as the contract duration increases, which aligns with the initial expectation for cost behavior.




Figure 5. Boxplots for new trucks.

It becomes essential to understand if there are significant differences between the identified groups and, if so, in which groups these differences occur. For this purpose, the non-parametric Kruskal–Wallis test [19] was considered as a replacement for the ANOVA test, as its normality assumptions are not met. The test resulted in a p -value < 0.001 , allowing us to conclude with reasonable significance that there are differences in costs between the different groups. Subsequently, to identify in which groups those differences occur, two appropriate post hoc tests were applied, the Dunn test [20] and the Dwass–Steel–Critchlow–Fligner test [21]. The results of these tests are represented in Figure 6.




Figure 6. Post Hoc results for new trucks.

As observed, all tests returned the same result, indicating that there are no significant differences in costs between the groups $[1,2[$ and $[2,3[$ and between the groups $[5,6[$ and $[6,8[$. Therefore, these groups were merged, resulting in 4 final groups for the new trucks in the T range: $[1,3[$ (88 contracts); $[3,4[$ (482 contracts); $[4,5[$ (400 contracts); and $[5,8[$ (207 contracts).

2.4. Used Trucks

Similarly to what was done for the new trucks, it is essential to examine the distribution of contracts across different groups for used trucks. It was considered that all groups had acceptable sizes for further analysis, so none was excluded. Again, as with new trucks, the groups are also separated according to the D and T ranges. It was observed that most of the used trucks belong to the T range, while a residual quantity is associated with the D range. In fact, the latter range is not represented in three out of the five groups, and in the other two, its presence is significantly lower than that of the T range. Therefore, for the same reason presented for new trucks, it was decided to exclude trucks from the D range from the set of used trucks as well.

Following that, the analysis of the average annual costs was then carried out, only for the used trucks of the T range. As can be seen in Figure 7, there does not appear to be a clear trend in costs, suggesting that, in reality, these trucks have a more complex behavior than the new trucks.




Figure 7. Boxplots for used trucks.

To understand if there are significant differences in costs among the various groups, the Kruskal–Wallis test was applied again, replacing the ANOVA test, resulting in a p -value of 4.47×10^{-3} , indicating that there are relevant differences between the groups.

Next, the two post hoc tests, the Dunn test and the Dwass–Steel–Critchlow–Fligner test, were applied again to understand between which groups the previously detected differences exist. The results are shown in Figure 8.




Figure 8. Post Hoc results for used trucks.

The results of both tests were consistent, indicating that there are significant differences only between the groups [1, 2] and [3, 4]. However, these results complicate the formation

of distinct groups. For example, the tests did not point to significant differences between the groups $[1, 2[$ and $[2, 3[$, nor between the groups $[2, 3[$ and $[3, 4[$. This suggests that, according to transitivity, both pairs of groups could be combined, resulting in the interval $[1, 4[$. However, this fusion would place the intervals $[1, 2[$ and $[3, 4[$ in the same group, and as observed, these groups present significant differences. Thus, the obtained results reveal the intrinsic complexity in the behavior of costs of used trucks, highlighting the difficulty in forming heterogeneous groups based on the analyzed variables. Faced with this scenario, a deeper reflection on the feasibility of including used trucks in the analysis in question was necessary. Considering the lack of clarity in distinguishing the groups and the uncertainty regarding the representativeness of the results, the decision was made to set aside the used trucks from the scope of the current analysis. Instead, the focus was exclusively directed towards new trucks, where cost behavior patterns seem to be clearer and more distinct. This strategic choice was based on the search for more robust and conclusive results, thus avoiding premature conclusions or misconceptions. Moreover, postponing the analysis of used trucks may represent an opportunity for future more in-depth studies, in which more appropriate methodologies may be employed or other relevant variables could be considered.

3. VaR-Based Pricing Model

Establishing contract prices is a fundamental task for any company, and this is especially true in the context of maintenance contracts, where setting appropriate premiums plays a crucial role in financial and strategic management. The premiums were calculated based on the previously calculated costs of the contracts, reflecting not only the direct costs associated with maintenance but also the underlying risks involved.

In this section, the methodology was based on the working paper [16], particularly the formula for calculating the premium for each group, which is based on the VaR (Value at Risk) risk measure [22,23]. This approach allows for a more precise and comprehensive assessment of the financial risks associated with each contract, taking into account factors such as the probability of adverse events occurring and the magnitude of potential losses.

3.1. Value at Risk

A risk measure associated with a random variable X , which in this context represents a loss, can be defined as a functional that maps a real number to X . The commonly used risk measure in practice is the VaR at a certain confidence level, typically denoted by p , where $0 < p < 1$. VaR represents the maximum potential loss that will not be exceeded with probability p . In other words, it is the threshold value x where the cumulative distribution function $F(x)$ intersects the p level. VaR is also referred to as the quantile risk measure [23].

Formally, for a random variable X , the VaR at confidence level p is defined as

$$VaR_p(X) = \inf\{x, F_X(x) \geq p\} = F_X^{-1}(p)$$

where F_X^{-1} is the inverse function of the cumulative distribution function. Typically, this chosen confidence level is arbitrary. In practice, it can be high, like 99.95%, for the entire company, or much lower, like 95%, for a single unit or risk class within the company [24]. In this work, confidence levels between 95% and 99.9% were considered. To calculate VaR , one must first make an assumption regarding the behavior of the variable under consideration, primarily to estimate its distribution. In this study, the parametric approach was employed; however there exist other methods for this purpose [16]. The parametric approach assumes a specific form for the underlying distribution of the data and estimates the parameters of that distribution using available data. It offers a simpler methodology and requires fewer historical data points when compared to, for example, the non-parametric method.

3.2. Cost Curve Fitting

With the four cost groups, $[1,3]$, $[3,4]$, $[4,5]$, and $[5,8]$, defined, the next step is to estimate the distribution of the average annual costs for each group to understand the cost behavior and also to calculate the VaR. For this purpose, probability distributions available in the Scipy library for Python 3.12.7 were used [25]. Specifically, for each group, 106 probability distributions were fitted using the maximum likelihood method, and the distribution with the lowest sum of squared residuals was selected as the best fit. The p -values obtained from the Kolmogorov–Smirnov (KS) test for each group, are presented in the respective figure legend, supporting the selected distributions.

Figures 9–12 illustrate the fitting of a theoretical distribution curve to the data of average annual costs for each group. Once the distribution is chosen, the VaR for each group is calculated using the `scipy.stats.rv_continuous.ppf` function, which corresponds to the inverse of the cumulative distribution function. For this calculation of VaR, a confidence level of $p = 95\%$ was chosen as an example.




Figure 9. Fitting for group $[1, 3]$.




Figure 10. Fitting for group $[3, 4]$.




Figure 11. Fitting for group [4, 5].




Figure 12. Fitting for group [5, 8].

In this example, it can be observed that for the groups [1,3[, [3,4[, and [4,5[, the fitted distribution exhibits right skewness, while for the [5,8[group, this skewness is less pronounced. Concerning the VaR values, a growing trend is observed, indicating that the VaR increases as the group intervals progress.

Table 1 summarizes the descriptive statistics of the average annual costs for each of the four cost groups.

Table 1. Descriptive statistics of average annual costs (M.U.) per group.

Group	Mean	Median	Std	Skewness	Kurtosis
[1, 3[	1104	937	826	1.95	7.63
[3, 4[	1804	1623	1013	1.31	5.68
[4, 5[	2733	2533	1119	0.91	3.99
[5, 8[	3749	3653	1108	0.62	4.49

3.3. Final Price Formula

In [16], it is suggested that a risk measure should be considered in the premium calculation, specifically that the loading, i.e., the surcharge added to the average costs, be proportional to this risk measure. In particular, the formula considered is

$$P_i = \overline{G}_i + (1 - p) \text{VaR}_p(G_i) \quad (1)$$

Here, P_i represents the premium for group i , $\overline{G}_i$ denotes the mean of group i , p is the confidence level, and $\text{VaR}_p(G_i)$ represents the Value at Risk for group i at the confidence level p . The above formula is used in this work to establish the annual premium for each cost group.

3.4. Risk Assessment and Final Prices

To account for the risk associated with the prices established for each group, 200 simulations were conducted using a cross-validation procedure where the group data were randomly divided into training sets comprising 80% and test sets comprising 20%. The contract data in the training set were fitted to theoretical distributions, from which the corresponding Value at Risk was calculated. Additionally, the mean and standard deviation of the costs of each group were also computed. Using the mean and VaR, the formula in (1) was then applied to each group to determine the annual premiums. Subsequently, for each contract in the test set, the premium corresponding to its group was applied, and it was verified whether the total contract premiums given by the model covered the actual costs of the test set. The percentage of simulations where these costs were not covered defines the risk of this model; for example, if in 5 of the 200 simulations performed, the costs were not covered, the risk would be 2.5%. The calculation of VaR employed confidence levels of 95%, 99%, 99.5%, and 99.9%, with simulations conducted once for each confidence level.

In Table 2, the annual premiums are presented, derived from the average premiums across the 200 simulations, along with the associated risk at varying confidence levels.

Table 2. Mean (std) annual premiums for each group and associated risk.

Group	Annual Premium (M.U.)			
[1, 3[	1207 (188)	1141 (114)	1123 (100)	1100 (83)
[3, 4[	1991 (378)	1853 (360)	1831 (348)	1811 (320)
[4, 5[	2975 (70)	2794 (62)	2766 (58)	2740 (49)
[5, 8[	4022 (405)	3814 (380)	3783 (372)	3754 (350)
Risk	0.5%	28%	39%	50.5%
p	95%	99%	99.5%	99.9%

In this table, it is evident that cost groups with longer durations tend to exhibit higher premiums, which aligns with expectations, as previously observed, these are associated with higher costs. Regarding risk, a considerable difference is observed across various confidence levels. Notably, at the 95% confidence level, the risk is quite low at 0.5%, while at 99%, the risk increases to 28%. One possible explanation for this phenomenon is related to the premiums, which are higher for a 95% confidence level and gradually decrease up to 99.9% confidence level.

However, this decrease in prices with an increase in the confidence level does not align with expectations, as with an increase in the confidence level of the VaR, a higher percentage of losses is considered, with their value being higher. Therefore, it would be expected that with the increase in the VaR confidence level, the annual pre-

mium would also be higher. For the same cost group, considering two confidence levels $0 < p_1 < p_2 < 1$, by Formula (1), the difference between the premiums would be given by $(1 - p_2)VaR_{p_2} - (1 - p_1)VaR_{p_1}$. As mentioned earlier, it is expected that this difference be positive, resulting in the following expression:

$$VaR_{p_2} > \alpha VaR_{p_1} \quad (2)$$

where $\alpha = \frac{1-p_1}{1-p_2} > 1$. This expression essentially indicates that to satisfy the equation, the value of the VaR associated with the higher confidence level must be greater than a proportion of the VaR at the lower confidence level. In this context, the increase in VaR is intrinsically linked to the behavior of the theoretical distribution considered, particularly its tail. Therefore, the “weight” of the distribution’s tail defines, in this situation, whether Equation (2) is satisfied. In this way, it is understood that the fitted theoretical distributions are not “heavy” enough to apply the approach presented in (1).

4. Adapted Model for the Pricing Challenge

As seen so far, an increase in the confidence level of VaR does not result in an increase in the contract price. Therefore, there is a need to consider a new approach to the contract price formula. By definition, if $p_2 > p_1$, one always expects that $VaR_{p_2} \geq VaR_{p_1}$, which means that the satisfaction of the previous condition depends solely on the factor $\frac{(1-p_1)}{(1-p_2)}$, which in this case can be seen as a percentage of VaR_{p_1} . The problem can be formulated as follows: When increasing the confidence level from p_1 to p_2 , how much does VaR_{p_2} need to increase relative to VaR_{p_1} ? In fact, this problem can be answered by substituting the values of p_1 and p_2 . Substituting the values of confidence levels 95%, 99%, 99.5%, and 99.9%, one obtains the following three transitions:

$$VaR_{99} > 5VaR_{95}$$

$$VaR_{99.5} > 2VaR_{99}$$

$$VaR_{99.9} > 5VaR_{99.5}$$

These values indicate that, for example, when transitioning from 99.5% to 99.9%, the VaR needs to increase by more than five times.

Since VaR is a measure of risk associated with the tail of the distribution, a natural question arises: how heavy does the tail of the distribution need to be to satisfy these conditions?

Hence, arises the need to quantify the tail of the distribution.

4.1. Quantify the “Fat-Tailedness” of the Distribution

In [26], a metric is introduced for comparing the “fat-tailedness” of distributions. The metric is defined as follows:

Let $X_1, \dots, X_n$ be independent and identically distributed (i.i.d.) random variables with finite mean, i.e., $E(X) < +\infty$. Let $S_n = X_1 + X_2 + \dots + X_n$ be a partial sum. Define $M(n) = E(|S_n - E(S_n)|)$ as the expected mean absolute deviation from the mean for n summands. The “rate” of convergence for n additional summands starting with n_0 is defined as follows:

$$\kappa_{n_0, n} = 2 - \frac{\log(n) - \log(n_0)}{\log\left(\frac{M_n}{M_{n_0}}\right)} \quad (3)$$

$\kappa_{n_0, n} \in [0, 1]$, where 0 represents maximally thin-tailed (Gaussian) distributions and 1 represents maximally fat-tailed distributions.

Using this metric allows for the quantification of heavy-tailed distributions that meet the aforementioned condition, thereby establishing a significant lower threshold for this pricing formula.

4.2. Finding a Distribution That Satisfy the Condition

The primary objective in this section is to define a lower bound on the κ metric that satisfies Equation (2). Consequently, any distribution with a metric value exceeding this bound will also satisfy the equation.

To achieve this, Equation (2) was rewritten as follows:

$$\frac{VaR_{p_2}}{\frac{1-p_1}{1-p_2} VaR_{p_1}} > 1 \quad (4)$$

One way to establish a lower bound is to transform (4) into an equality. Thus, the aim is to identify a theoretical distribution that satisfies this equality across the three aforementioned transitions.

For this purpose, two distributions known to have heavy tails were considered: the Pareto distribution and the t-student distribution.

The Pareto distribution, $Par(x_m, \alpha)$, has two parameters: a scale parameter x_m , and a shape parameter α , also known as the tail index. x_m is fixed at 1, and several levels of α are considered. α cannot be less than or equal to 1 because in such cases, the mean is not finite.

The t-student distribution, $t(\nu)$, has only one parameter, ν , which controls the amount of probability mass in the tails. For $\nu = 1$, the Cauchy distribution is obtained, which has very “fat” tails and does not have a finite mean. As $\nu \rightarrow \infty$, it approaches the standard normal distribution, which has very “thin” tails. Thus, ν greater than 1 is considered.

The values of the quotient of the left-hand side of Equation (4) for each transition are depicted in Figure 13 for the two distributions with parameters 1.01 and 1.001.




Figure 13. Quotient per transition and distribution.

It is observed that the Pareto distribution with parameter 1.001 is the one that comes closest to 1, and therefore, most closely approaches the desired equality. Thus, one may calculate the metric for this distribution and propose a lower bound on the κ scale. To calculate the metric $\kappa_{n_0, n}$ for the distribution, the values $n_0 = 1$ and $n = 2$ are utilized following [26] original idea. κ_1 will be used as an abbreviation. The result for the Pareto distribution with parameters $Par(1, 1.001)$ was $\kappa_1 \approx 0.773$.

Thus, the following scheme can be considered:

The spectrum in Figure 14 indicates that if the adjusted theoretical distribution has a κ greater than 0.773, Formula (1) may be applied, as condition (4) will be satisfied.




Figure 14. Spectrum of κ for pricing formulas (first iteration).

4.3. A Proposal for Small Values of κ

The problem that arises now is what to do when the value of $\kappa_1 < 0.773$. To address this, the situation where the variation in VaR is insignificant was first considered. In this scenario, it is proposed that the loading factor remains constant. If the increase in VaR is very small, indicating low risk, the increase in price should also be minimal, achieved by maintaining a constant factor.

The variation in VaR across different confidence levels can be constrained between the lowest level, 95%, and the highest, 99.9%, by the following equation:

$$VaR_{99.9} = c VaR_{95} \quad c > 1 \quad (5)$$

To define a significance threshold, a value for the constant c can be established. After choosing the value of c , the graph of this constant c as a function of the usual confidence level can be considered.

Thus, the slope of the line connecting the points (95,1) and (99.9,c) can be calculated, and this slope is considered as the factor multiplying the VaR.

The selection of the value of c is subjective and depends on individual risk perception. In this study, the value of 1.33, representing a 33% increase, was chosen, but this may be easily adapted to other situations. Therefore, the slope is approximately 0.0673 (Figure 15), resulting in the corresponding price formula:

$$P_p = \bar{G} + 0.0673 VaR_p \quad (6)$$

A 33% increase between 95% and 99.9% translates to an average increase of 10% across the four confidence levels. To find a metric representing this type of variation, the Pareto distribution is reconsidered.




Figure 15. Slope lines between the 95 and 99.9 confidence levels.

The goal is to find the parameter α of this distribution from which there is an average increase, between the three transitions, equal to or less than 10%. The Pareto distribution is chosen due to its easily computable quantile function:

$$VaR_p = F^{-1}(p) = \frac{1}{(1-p)^{1/\alpha}} \quad (7)$$

This problem can be translated into the following inequality:

$$\frac{F^{-1}(0.99)}{F^{-1}(0.95)} + \frac{F^{-1}(0.995)}{F^{-1}(0.99)} + \frac{F^{-1}(0.999)}{F^{-1}(0.995)} - 3 \leq 0.3 \quad (8)$$

Simplifying this equation yields:

$$2 \times 5^{1/\alpha} + 2^{1/\alpha} \leq 3.3 \quad (9)$$

Given the difficulty in calculating α analytically, the problem can be transformed into finding the zero of a function, after which the Newton–Raphson method can be applied. The solution obtained is $\alpha \approx 13.752$.

This result indicates that $\alpha \approx 13.752$ is the smallest value of the Pareto distribution parameter for which the average transition is insignificant. Any value of α between 0 and 13.752 will result in significant VaR variations. Therefore, the metric for the distribution $Par(1, 13.752)$ can be considered as an upper limit for having insignificant increases in the tail.

The metric for this distribution is $\kappa_1 \approx 0.265$, thus the scheme presented in Figure 14 becomes the one presented in Figure 16.




Figure 16. Spectrum of κ for pricing formulas (second iteration).

4.4. A Heuristic for Intermediate Values of κ

Finally, the problem arises of understanding what to apply when the value of κ_1 falls between 0.265 and 0.773.

Based on the previous observations, distributions with a metric less than or equal to 0.265 are expected to have average increases less than or equal to 10% of the VaR. In other words, for $p_1 > p_2$, it follows that $VaR_{p_1} \leq 1.1 VaR_{p_2}$. Therefore, distributions with a metric greater than 0.265 must satisfy, on average, $VaR_{p_1} > 1.1 VaR_{p_2}$.

The goal is to find a function $f(p)$ such that the average transition is at least 1.1 and also “close” to $1 - p$, to avoid exorbitant prices. It should be noted that these requirements do not uniquely determine the function $f(p)$, but rather define a class of admissible functions. In this context, the approach followed here is heuristic: the objective is not to identify an optimal specification, but to select a function that simultaneously satisfies the imposed constraints in a stable and tractable manner. Several candidate functions were explored during this process. However, it was observed that many of them either failed to consistently satisfy the required growth condition across all transitions or deviated excessively from the reference function $1 - p$, leading to undesirable pricing behavior. The function adopted below was selected because it provides a satisfactory balance between these two competing objectives:

$$f(p) = \frac{p^{10} \sqrt[10]{1-p}}{25} \quad (10)$$

For the transition $p_1 > p_2$ it holds

$$VaR_{p_1} > \frac{p_2}{p_1} \sqrt[10]{\frac{1-p_2}{1-p_1}} VaR_{p_2} \quad (11)$$

For the three transitions, it holds that

$$VaR_{99} > 1.1 VaR_{95}$$

$$VaR_{99.5} > 1.1 VaR_{99}$$

$$VaR_{99.9} > 1.2 VaR_{99.5}$$

which, on average, leads to $VaR_{p_1} > 1.1 VaR_{p_2}$.

The closeness between the proposed function and $1 - p$ was measured in the space of bounded functions $\mathcal{B}([0.95, 1])$ using the supremum norm. The resulting maximum deviation is approximately 0.022, which is small in the context of the interval under analysis and ensures that the proposed function remains close to $1 - p$ without leading to inflated pricing. For sake of brevity, the details of the calculation are omitted. We emphasize that the proposed specification is not unique, and alternative functions satisfying the same criteria could also be considered. The function in Equation (10) should therefore be interpreted as one representative element within a broader admissible family. Its selection is justified by its ability to satisfy the required growth condition while remaining sufficiently close to the baseline behavior, without introducing excessive pricing distortions.

Observing Figure 17, it is noted that Equation (10) is relatively close to the function $1 - p$ in the interval from 0.95 to 0.999 as intended.




Figure 17. Plot of $f(p)$ and $1 - p$.

This way, the final obtained scheme is presented in Figure 18.




Figure 18. Spectrum of κ for pricing formulas (final iteration).

4.5. Application to the Dataset

With the complete spectrum, there are premium calculation formulas for all possible scenarios, allowing this approach to be applied to the dataset provided by the company. Following a similar procedure to that in Section 3.4, a new pricing scheme was achieved, and its results are summarized in Table 3.

Table 3. Mean (std) annual premiums for each group and associated risk.

Group	Annual Premium (M.U.)			
[1, 3[	1166 (72)	1229 (87)	1268 (100)	1403 (161)
[3, 4[	1908 (278)	1975 (297)	2002 (305)	2059 (325)
[4, 5[	3024 (35)	3117 (45)	3156 (49)	3246 (61)
[5, 8[	4001 (300)	4088 (315)	4130 (322)	4244 (346)
Risk	14.5%	4.5%	3.5%	0.5%
p	95%	99%	99.5%	99.9%

It is noted that, similar to the previous model, premiums are higher for cost groups with longer durations; however, they now increase with the confidence level, thereby contributing to risk reduction. These data indicate that the proposed model performed well for these dataset, allowing for more flexibility regarding the weight of the tail of the adjusted theoretical distribution.

5. Conclusions

In conclusion, this study delved into the risk analysis of maintenance contracts, initially focusing on cost analysis. Through statistical tools, significant segmentation of contracts into heterogeneous cost groups was achieved. Subsequently, a pricing model based on the Value at Risk (VaR) measure was applied, with theoretical distributions fitted to the defined cost groups. However, it was observed that this model yielded unsatisfactory results, particularly as prices decreased with increasing VaR confidence levels. To address this, a second model was proposed, introducing a metric to gauge the “weight” of distribution tails, offering greater flexibility in pricing formula selection. This second model, essentially a spectrum based on this metric, defined three pricing formulas according to the tail weight of the cost data distributions. Remarkably, this model effectively resolved the issues of the previous model and demonstrated good adaptability to the given problem. The data from the used trucks subset were discarded due to the lack of clear statistical separation of cost groups and missing features. As a result, the generalizability of the proposed model remains limited to new vehicles. Future work should address this limitation by incorporating additional explanatory variables or alternative modeling approaches. From a practical perspective, the proposed adaptive pricing framework provides decision-makers with a flexible tool to align premiums with risk exposure, ensuring consistent pricing behavior across confidence levels and data distributions, thereby supporting more informed and robust maintenance contract management strategies.

Author Contributions: Conceptualization, B.P., M.C., F.V. and P.L.; methodology, B.P. and M.C.; software, B.P., F.V. and P.L.; validation, F.V. and P.L.; formal analysis, M.C., J.S., C.L. and S.R.; investigation, B.P., M.C., J.S., C.L. and S.R.; resources, B.P., F.V. and P.L.; data curation, B.P., F.V. and P.L.; writing—original draft preparation, B.P.; writing—review and editing, M.C., J.S., C.L. and S.R.; visualization, B.P.; supervision, M.C., F.V. and P.L.; project administration, M.C. and P.L.; funding acquisition, M.C. and P.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The datasets presented in this article are not readily available because they belong to a private company and therefore are confidential. Requests to access the datasets should be directed to Pedro Louro.

Conflicts of Interest: Author Bruno Pereira was employed by the company Sonae MCretail and author Filipa Vieira, Pedro Louro were employed by the company Nors Group. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

- Barroso, S.; Cruz, M.B.; Ramos, S.F.; Pina, M. Price Modelling on Heavy-Duty Assistance Contracts. In *Progress in Industrial Mathematics: Success Stories*; Springer International Publishing: Cham, Switzerland, 2021; pp. 193–208.
- Rahman, A.; Chattopadhyay, G. Cost estimation for Maintenance contracts for complex asset/equipment. In *Proceedings of the 2008 IEEE International Conference on Industrial Engineering and Engineering Management*; IEEE: Piscataway, NJ, USA, 2009; pp. 1355–1358. [\[CrossRef\]](#)
- Deprez, L.; Antonio, K.; Boute, R. Pricing service maintenance contracts using predictive analytics. *Eur. J. Oper. Res.* **2021**, *290*, 530–545. [\[CrossRef\]](#)
- Deprez, L.; Antonio, K.; Boute, R. Pricing full-service maintenance contracts: A data analytics approach. *PHM Soc. Eur. Conf.* **2018**, *4*. [\[CrossRef\]](#)
- Jing, H.; Tang, L.C. A Risk-based Approach to Managing Performance-based Maintenance Contracts. *Qual. Reliab. Eng. Int.* **2017**, *33*, 853–865. [\[CrossRef\]](#)
- Bowman, R.A.; Schmee, J. Pricing and Managing a Maintenance Contract for a Fleet of Aircraft Engines. *Simulation* **2001**, *76*, 69–77. [\[CrossRef\]](#)
- Huber, C.; Spinler, S. Pricing of Full-Service Repair Contracts. *Eur. J. Oper. Res.* **2012**, *222*, 113–121. [\[CrossRef\]](#)
- Huber, C.; Spinler, S. Pricing of Full-Service Repair Contracts with Learning, Optimized Maintenance, and Information Asymmetry. *Decis. Sci.* **2014**, *45*, 791–819. [\[CrossRef\]](#)
- Deprez, L.; Antonio, K.; Boute, R. Empirical risk assessment of maintenance costs under full-service contracts. *Eur. J. Oper. Res.* **2023**, *304*, 476–493. [\[CrossRef\]](#)
- Dempster, M.A.H. *Risk Management: Value at Risk and Beyond*; Cambridge University Press: Cambridge, UK, 2002.
- Linsmeier, T.J.; Pearson, N.D. Value at Risk. *Financ. Anal. J.* **2000**, *56*, 47–67. [\[CrossRef\]](#)
- Sarykalin, S.; Serraino, G.; Uryasev, S. Value-at-Risk vs. Conditional Value-at-Risk in Risk Management and Optimization. In *INFORMS Tutorials in Operations Research*; INFORMS: Hanover, MD, USA, 2014; pp. 270–294. [\[CrossRef\]](#)
- Du, Z.; Escanciano, J.C. Backtesting Expected Shortfall: Accounting for Tail Risk. *Manag. Sci.* **2016**, *4*, 940–958.
- Wang, S.S. A Class of Distortion Operators for Pricing Financial and Insurance Risks. *J. Risk Insur.* **2000**, *67*, 15–36. [\[CrossRef\]](#) [\[PubMed\]](#)
- Armstrong, J.; Brigo, D. Risk managing tail-risk seekers: VaR and expected shortfall vs S-shaped utility. *J. Bank. Financ.* **2019**, *101*, 122–135. [\[CrossRef\]](#)
- Alemany, R.; Bolance, C.; Guillen, M. *Accounting for Severity of Risk When Pricing Insurance Products*; Working Papers 2014-05; Universitat de Barcelona, UB Riskcenter: Barcelona, Spain, 2014.
- Shapiro, S.S.; Wilk, M.B. An analysis of variance test for normality (complete samples). *Biometrika* **1965**, *52*, 591–611. [\[CrossRef\]](#)
- Dodge, Y. Mann–Whitney Test. In *The Concise Encyclopedia of Statistics*; Springer: New York, NY, USA, 2008; pp. 327–329. [\[CrossRef\]](#)
- Kruskal, W.H.; Wallis, W.A. Use of Ranks in One-Criterion Variance Analysis. *J. Am. Stat. Assoc.* **1952**, *47*, 583–621. [\[CrossRef\]](#)
- Dinno, A. Nonparametric Pairwise Multiple Comparisons in Independent Groups using Dunn’s Test. *Stata J.* **2015**, *15*, 292–300. [\[CrossRef\]](#)
- Hollander, M.; Wolfe, D.A.; Chicken, E. *Nonparametric Statistical Methods*; Wiley: Hoboken, NJ, USA, 2015.
- Stambaugh, R.F. Risk and Value at Risk. *Eur. Financ. Manag.* **1996**, *2*, 281–300. [\[CrossRef\]](#)
- Kaas, R.; Goovaerts, M.; Dhaene, J.; Denuit, M. *Modern Actuarial Risk Theory*; Number 978-3-540-70998-5 in Springer Books; Springer: Berlin/Heidelberg, Germany, 2008. [\[CrossRef\]](#)
- S. Klugman, H.P.; Willmot, G. *Loss Models: From Data to Decisions*; John Wiley & Sons: Hoboken, NJ, USA, 2012; Volume 715.

25. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [[CrossRef](#)] [[PubMed](#)]
26. Taleb, N.N. Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications. *arXiv* **2022**, arXiv:2001.10488. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.