DOI: 10.65322/j.srd.05192026.96

UDC: 629.3.083:658.787

OPTIMIZING SPARE PARTS EXCHANGE POOLS IN AUTOMOTIVE SERVICE ENTERPRISES SUBJECT TO A SERVICE LEVEL CONSTRAINT

Andrea D'Ariano,
E-mail: andrea.dariano@uniroma3.it
Roma Tre University,
Polvonov A., *Professor of the Department;*
E-mail: aspolvonov@mail.ru,
Abdusattorov N., *PhD, Dean of the Faculty;*
E-mail: nabdusattorov93@mail.ru,
Imomnazarov S., *PhD, Head of the Department,*
E-mail: sarvar1991082@gmail.com,
Namangan State Technical University

Abstract. *A spare-parts allocation that minimizes initial investment may still leave vehicles waiting longer for repairs. This trade-off motivates our examination of exchange pools in automotive service enterprises. We determine how demand rates, replenishment lead times and purchase prices influence the least-cost stock allocation under a specified average service requirement. A single-echelon base-stock model with independent Poisson demand is evaluated using immediate demand fulfilment and expected backorders; complete enumeration verifies the least-cost integer solution in an illustrative three-item example. At a 95% service requirement, the allocation (4, 3, 2) costs UZS 18 million and achieves a 96.38% fill rate, requiring 14.29% less initial capital than the uniform allocation (3, 3, 3). Yet its expected backorders and mean logistics waiting time are higher, showing that a favourable average fill rate can conceal delays associated with a less frequently demanded item. Increasing all replenishment lead times by 50% raises the reoptimized minimum cost to UZS 21 million. We therefore interpret the least-cost allocation as a solution to a defined inventory problem, whose operational suitability also depends on critical-part availability and acceptable waiting times. The findings support combining average service constraints with item-specific or waiting-time limits while assessing technical availability separately. All numerical results are illustrative calculations and require validation against enterprise records before operational or economic benefits can be established.*

Keywords: *automotive service; spare parts; exchange pool; pooled inventory; Poisson demand; service level; backorders; integer optimization.*

Introduction

Once an automotive service enterprise accepts a repair order, completion time depends on the availability of the required part as well as diagnostic work and operational capacity. If the required item is unavailable in stock, the vehicle must wait until the supply process is completed. However, increasing spare-parts stocks without limit is also undesirable: expensive components used infrequently tie up capital for extended periods. The relationship between spare-parts inventory, demand characteristics, repairability and supply networks is treated as a distinct problem in the literature [1, 2].

An exchange pool is a stock of serviceable components used to replace unserviceable components promptly and linked to a repair-and-return flow. Non-repairable parts constitute consumable inventory. Both categories may be managed within a common accounting system, but their replenishment sources differ. Part criticality, price and supply characteristics are important when selecting a strategy [3]. Planning must coordinate orders, repairs and returned components rather than focus solely on purchasing quantities [4].

DOI: 10.65322/j.srd.05192026.96

Intermittent demand makes it difficult to apply a single inventory norm to all items. The link between classification and forecasting and practical inventory control [5], together with the relationship between inventory policy and the stage of a product's service life [6], calls for a careful approach. Accordingly, this study examines a cost-differentiated allocation subject to a service level constraint rather than uniform stock quantities across items.

We ask how a specified average level of immediate demand fulfilment can be achieved at the lowest initial pool cost when the demand rate, replenishment lead time and purchase price of each item are known. Our contribution is to formulate this question for automotive service operations, distinguish inventory service levels from technical availability and verify a heuristic solution through complete enumeration for a small set of items. The Poisson model and base-stock policy are established approaches rather than new theory. We use them to examine an operationally relevant question: whether reducing initial inventory investment under an average service constraint also preserves acceptable waiting-time performance.

Materials and Methods

2.1 Study Design and Standards Framework

The study is theoretical and computational. No enterprise order logs, inventory movement records, failure statistics or repair protocols were provided. Therefore, the parameters below are illustrative values selected to evaluate the method; they must not be interpreted as actual performance data for a specific enterprise or vehicle fleet in the Namangan region.

GOST R 27.102-2021 [7] was used as a methodological source for distinguishing reliability and availability terminology, GOST 18322-2016 [8] for maintenance and repair concepts, and GOST R 53392-2017 [9] for the general framework of logistics analysis. The latter replaces GOST R 53392-2009. GOST R documents are Russian national standards; their use in this article does not imply mandatory applicability in Uzbekistan. Their local adoption requires separate verification. These standards do not prescribe the calculation algorithm used below; the mathematical model is developed from the scientific literature and explicit assumptions.

2.2 Input Data and Assumptions

The model comprises a single-echelon shared warehouse and M non-interchangeable item types. For item i, r_i is the aggregate demand rate in units/day; L_i is the replenishment lead time in days; S_i is the base-stock level in units; and c_i is the unit purchase price in UZS/unit. If the individual failure rate λ_i is known for m_i installed components operating under identical conditions, $r_i \approx m_i\lambda_i$ may be used. Where operating conditions differ, the aggregate rate is estimated from the sum of individual rates or from a time series of service orders.

Demand is assumed to follow independent, stationary Poisson processes, with each request requiring one part. Under the base-stock policy, a one-unit replenishment order is placed after each issue, and unmet demand is backordered. In the numerical example, L_i is constant. Repair capacity constraints, rejected components, non-return to stock, supplier disruptions and common-cause failures are excluded from the baseline calculation. Where these factors are present, the probability distribution or state representation of the model must be extended.

The Poisson model may provide an initial approximation for low, independent demand, but its suitability must be tested empirically. Croston's method [10], which forecasts demand size and inter-arrival intervals separately, and its corrected estimators [11] offer alternative approaches to intermittent demand. A recent review also supports selecting forecasting methods jointly with demand characteristics and inventory decisions [12]. These methods are not calculated in this study; they are identified as alternatives for comparison once operational data become available.

Table 1

Input parameters for the illustrative numerical example

Item	r_i, units/day	L_i, days	$a_i = r_iL_i$	c_i, million UZS/unit	w_i

DOI: 10.65322/j.srd.05192026.96

P_1	0.08	10	0.80	1	0.5714
P_2	0.04	20	0.80	2	0.2857
P_3	0.02	30	0.60	4	0.1429

P_1–P_3 do not denote specific automotive component models. Demand shares are calculated as $w_i = r_i / \Sigma r_j$. Rounded shares are shown only in the table; exact values are used in the calculations. All times are expressed in days.

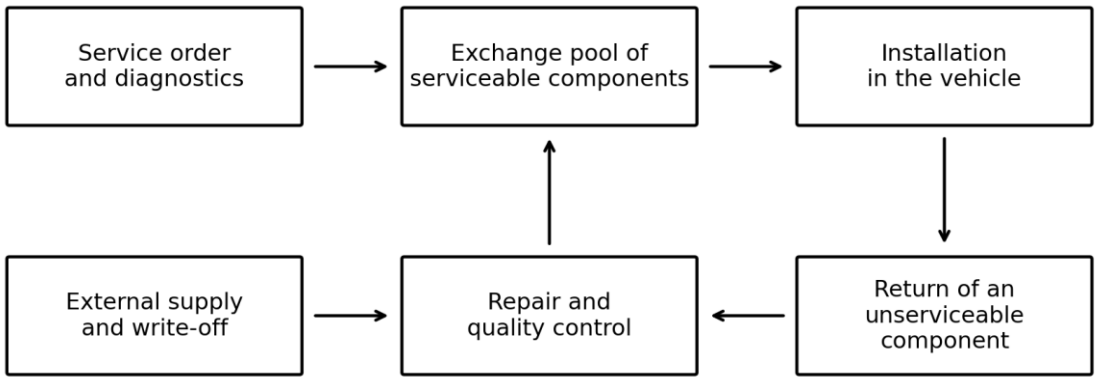

Figure 1. Conceptual flow of a repairable-component exchange pool. The diagram illustrates the scope of the computational model and does not represent a field observation

2.3 Service Level and Backorder Model

Under the stated assumptions, the stationary distribution of X_i, the number of units in the replenishment pipeline in a single-echelon base-stock system, is Poisson. The theoretical framework for modelling repairable spare parts is discussed in [13], while the classical approach to multi-echelon pools is addressed in [14]. The full classical multi-echelon METRIC model is not used here.

$$a_i = r_i L_i, \quad P(X_i = k) = e^{-a_i} \frac{a_i^k}{k!}, \quad k = 0,1,2,\ldots \qquad (1)$$

The dimensionless parameter a_i represents the expected number of units in the replenishment pipeline. A serviceable unit is available in stock immediately before a new demand arrival only if $X_i \leq S_i - 1$. Since Poisson arrivals observe the time-average state of the system, the fraction of demand fulfilled immediately is expressed as follows:

$$\beta_i(S_i) = \sum_{k=0}^{S_i-1} e^{-a_i} \frac{a_i^k}{k!}, \quad \beta_i(0) = 0 \qquad (2)$$

The threshold $S_i - 1$ is essential: when $X_i = S_i$, no unit remains available for the arriving request. β_i is an inventory service level, not the proportion of time for which a vehicle is operationally available. The aggregate measure across items is weighted by the shares of incoming demand:

$$\bar{\beta}(\mathbf{S}) = \sum_{i=1}^{M} w_i \beta_i(S_i), \quad w_i = \frac{r_i}{\sum_{j=1}^{M} r_j}, \quad \sum_{i=1}^{M} w_i = 1 \qquad (3)$$

An unweighted sum of probabilities may exceed 1 and cannot serve as an aggregate availability coefficient. A product of probabilities for different parts is meaningful only for a separately defined joint event under an independence assumption. No such product is used in this study. Expected backorders are calculated as an additional measure:

DOI: 10.65322/j.srd.05192026.96

$$EBO_i(S_i) = E[\max(X_i - S_i, 0)] = \sum_{k=S_i+1}^{\infty} (k - S_i)\, P(X_i = k) \qquad (4)$$

$$EBO_i = a_i P(X_i \geq S_i) - S_i P(X_i \geq S_i + 1), \quad EBO = \sum_{i=1}^{M} E\,BO_i \qquad (5)$$

EBO describes the number of requests awaiting fulfilment at a randomly selected point in time. In a stationary, stable system, the mean logistics waiting time is estimated as $\bar{W} = EBO / \Sigma r_i$. This average includes requests with zero waiting time. Immediate fulfilment from stock does not imply immediate completion of assembly and installation.

2.4 Optimization and Solution Verification

The objective function is the initial acquisition cost of the pool. It is not equated with total annual cost or enterprise profit. The principal optimization problem is:

$$\min C(\mathbf{S}) = \sum_{i=1}^{M} c_i\, S_i, \quad \bar{\beta}(\mathbf{S}) \geq \beta^*, \quad S_i \in \{0,1,2, \dots\} \qquad (6)$$

Here, β* is the required average fraction of requests served immediately. For items affecting safety or whose absence causes a complete operational shutdown, individual constraints $\beta_i(S_i) \geq \beta_i^*$ are added. These are not applied in the numerical example because empirical information on item criticality is unavailable.

An initial solution is obtained by comparing the service level improvement from adding one stock unit with its purchase price. At each step, the item with the largest g_i is selected:

$$g_i(\mathbf{S}) = \frac{w_i[\beta_i(S_i + 1) - \beta_i(S_i)]}{c_i} = \frac{w_i P(X_i = S_i)}{c_i} \qquad (7)$$

This procedure is a stock-increment heuristic. It should not be described as coordinate descent or as an algorithm guaranteed to produce the least-cost solution in every case. Discrete increments and different unit prices may cause the final service level overshoot to cost more than an alternative allocation. In the numerical example, the heuristic provides an upper cost bound for complete enumeration.

For verification, B denotes the cost of the feasible solution obtained by the heuristic. Since $c_i > 0$, any better solution must lie in the region $C(S) \leq B$, with $0 \leq S_i \leq \text{floor}(B/c_i)$ for each item. All integer vectors in this finite region are enumerated, and the least-cost solution satisfying the service constraint is selected. Among solutions with equal cost, the one with the higher $\bar{\beta}$ is chosen. This procedure verifies the global minimum cost in the small example, although the search space grows rapidly for larger item sets.

The calculations were checked in two forms: using the Poisson sum and its cumulative distribution function. The infinite EBO sum was evaluated using the tail probabilities in Equation (5). In the sensitivity analysis, all L_i values were multiplied by the same factor while demand rates and prices remained unchanged; stock levels were reoptimized for each scenario.

Results

3.1 Baseline Calculation

For β* = 0.95, the globally least-cost solution is S = (4, 3, 2). Its acquisition cost is 4 × 1 + 3 × 2 + 2 × 4 = UZS 18 million. The item-level service levels are $\beta_1 = 0.990920$, $\beta_2 = 0.952577$ and $\beta_3 = 0.878099$, respectively. The weighted average is $\bar{\beta} = 0.963848$, or 96.38%.

Although the service level for the third item is below 95%, the average constraint is satisfied. This demonstrates that an aggregate measure can conceal shortages of a low-demand item. If P3 is critical, the solution (4, 3, 2) cannot be adopted without an individual service constraint.

Table 2

DOI: 10.65322/j.srd.05192026.96

Stock allocations evaluated against the same service requirement

Allocation	S	C, million UZS	β̄, %	EBO, units
Uniform stock	(3, 3, 3)	21	95.60	0.025192
Optimized	(4, 3, 2)	18	96.38	0.039227
High stock	(4, 4, 4)	28	99.17	0.003674

The uniform allocation (3, 3, 3) also satisfies the 95% requirement. Relative to this allocation, optimization reduces initial capital by (21 − 18)/21 × 100 = 14.29%. This calculated difference must not be interpreted as an annual economic benefit: holding, repair, logistics, obsolescence and downtime costs are excluded. The high-stock allocation (4, 4, 4) provides a different service level, so presenting savings relative to it as an outcome of equivalent service quality would be unjustified.

The optimized allocation has a higher β̄ but also a higher EBO. The fraction of requests not served immediately and the duration of the backlog capture different aspects of performance. Reducing stock for P3, which has a long replenishment lead time, increases the consequences of waiting. For the baseline solution, $\bar{W}$ = 0.039227/0.14 = 0.2802 days; for (3, 3, 3), it is 0.1799 days. Thus, constraining β̄ alone does not ensure the shortest mean waiting time.

3.2 Effect of Changes in the Service Requirement

Table 3

Required service level and minimum pool cost

β*, %	Optimal S	C, million UZS	Achieved β̄, %	EBO, units
90	(4, 3, 1)	14	91.68	0.161129
95	(4, 3, 2)	18	96.38	0.039227
98	(5, 3, 3)	23	98.23	0.014701
99	(5, 4, 3)	25	99.33	0.005621

Increasing the service requirement from 95% to 99% raises the pool cost from UZS 18 million to UZS 25 million, an increase of 38.89%. Because stock quantities are integers, the achieved service level exceeds the specified threshold. In this example, the cost of improving service is not uniform; it is better assessed through discrete solutions than represented by a constant percentage relationship.

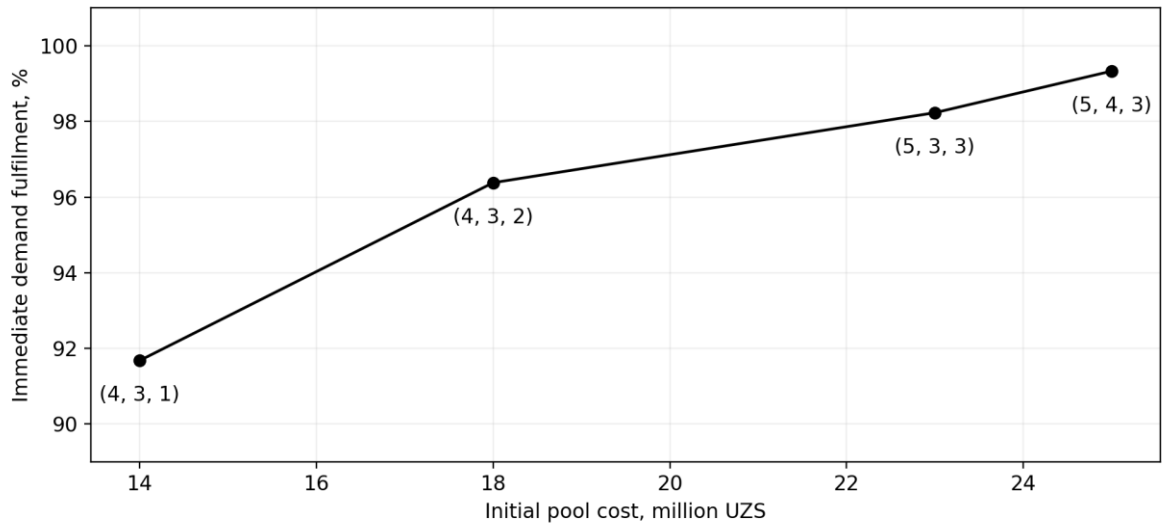

Figure 2. Optimized pool cost and achieved service level in the illustrative example. Vectors adjacent to the markers indicate stock quantities; the line only connects the calculated points

DOI: 10.65322/j.srd.05192026.96

3.3 Sensitivity to Replenishment Lead Time

Table 4

Reoptimization under lead time changes at β* = 95%

Lead time factor	Optimal S	C, million UZS	β̄, %	EBO, units
0.75	(4, 2, 2)	16	95.25	0.039536
1.00	(4, 3, 2)	18	96.38	0.039227
1.25	(5, 3, 2)	19	95.02	0.073034
1.50	(5, 4, 2)	21	95.34	0.090385

Reducing all lead times by 25% lowers the pool cost to UZS 16 million; increasing them by 50% raises it to UZS 21 million. The latter represents a 16.67% increase relative to the baseline. The stock levels in the table are reselected for each scenario; they do not indicate the resilience of a fixed pool to longer lead times. With a 25% increase, the achieved $\bar{\beta} = 95.02\%$ is very close to the threshold. A small parameter estimation error could violate the constraint, so practical stock selection requires evaluation against a credible high-demand or long-lead-time scenario.

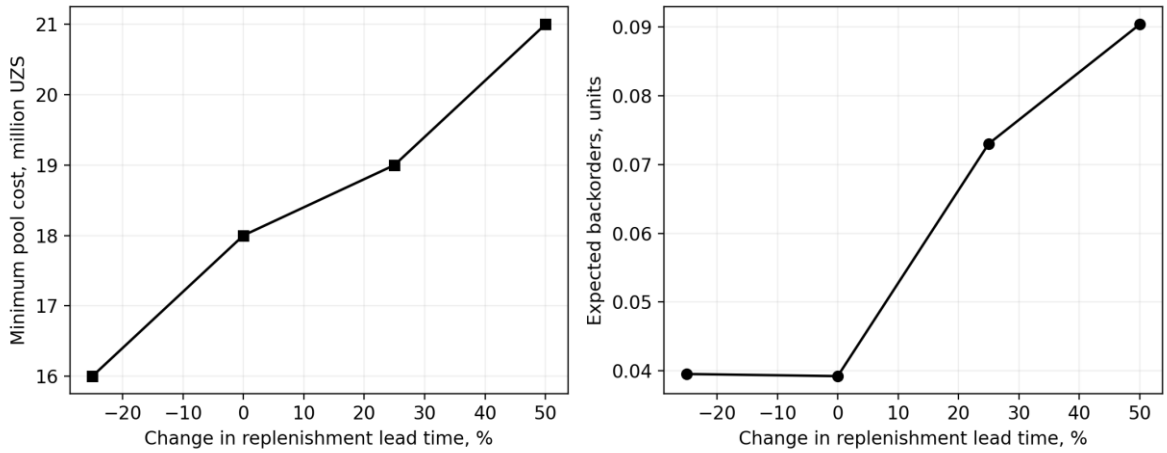

Figure 3. Minimum initial pool cost and EBO under changes in replenishment lead time. The plots are based on results from the illustrative model

Discussion

The calculations show that allocating stock solely in proportion to the number of items overlooks differences in demand rates, prices and lead times. Minimum acquisition cost alone, however, is not a sufficient criterion for a technical system. We interpret the higher EBO and mean waiting time for (4, 3, 2) as evidence that satisfying an average service requirement does not resolve every operational consequence of a stock allocation. In this example, the initial capital saving must be considered alongside the longer logistics wait associated with the allocation. This supports supplementing the service constraint with a waiting-time constraint and is conceptually consistent with models emphasizing system-level service criteria [2].

Technical availability is the ratio of time in an operable state to the specified observation time or required time. For a common, explicitly defined accounting period, an assessment may take the following form:

$$K_t = \frac{T_y}{T_y + T_x + T_r + T_l} \quad (8)$$

T_γ denotes time in an operable state; T_x, unavailability due to scheduled maintenance; T_r, repair time; and T_l, additional logistics waiting time. These time categories must be mutually

DOI: 10.65322/j.srd.05192026.96

exclusive; otherwise downtime is counted twice. The accounting period and the types of downtime included must be specified separately. $\bar{\beta}$ and K_t must not be equated or multiplied without a basis for independence. No numerical K_t is calculated in this study because actual time records are unavailable.

A pooled warehouse allows separate service enterprises to share inventory, but the benefits of centralization are constrained by last-mile delivery time. In a network where enterprises transfer spare parts to one another, local stocks, lateral transshipments and emergency deliveries must be treated as joint decisions [15]. The single-echelon calculations in this article cannot be transferred automatically to two- or multi-echelon networks.

For consumable parts under periodic review, safety stock generally needs to cover demand over the protection period $L_i + R_i$, where R_i is the review interval. However, the service level of a periodic-review system cannot be calculated exactly merely by replacing L_i with $L_i + R_i$ in a continuous-review formula: the within-cycle inventory state and order quantity also matter. Accordingly, the numerical results in this article apply to a continuous-review, one-for-one replenishment policy.

For a repairable pool, L_i includes the full turnaround time for transport, queueing, diagnostics, repair and quality control. When repair capacity is heavily utilized, L_i may depend on the demand rate, potentially violating the assumption of independent replenishment [13]. If components are rejected or not returned, the repair flow is supplemented by external procurement. Counting an unserviceable component as equivalent to serviceable stock overstates availability.

Practical validation requires recording, for each SKU, demand and stock-issue timestamps, backordered requests, actual delivery durations, component returns, rejections, purchase prices and the causes of vehicle downtime. Parameters should be estimated for an earlier period, with service levels, waiting times and pool costs compared in a subsequent independent period. Requests cancelled because of stockouts must also be recorded; otherwise demand is underestimated. Without these observations, the model's effectiveness in actual operations cannot be regarded as demonstrated.

The main limitations are the use of three illustrative items, stationary demand, constant lead times and an objective function restricted to initial acquisition cost. Further work requires a model incorporating seasonal demand, common-cause failures, obsolescence, finite repair capacity and lost demand. Introducing individual service constraints for critical items or an upper bound on EBO would align the optimization objective more closely with the enterprise's actual operational goals.

Conclusions

A single-echelon base-stock model was formulated to establish a spare-parts exchange pool subject to a service level constraint. The fraction of demand fulfilled immediately, expected backorders and technical availability were defined as separate measures. Using an unweighted sum of probabilities as aggregate availability and interpreting a heuristic search as unconditionally global optimization were rejected.

In the illustrative three-item example, we identify (4, 3, 2) as the globally least-cost stock allocation for a 95% average service requirement. The pool costs UZS 18 million and achieves a 96.38% service level, with initial capital 14.29% lower than for (3, 3, 3). Its poorer EBO and mean waiting-time performance limits the interpretation of this saving: the allocation is optimal within the chosen objective and constraint, while its suitability for an enterprise depends on the importance of the affected items and the waiting time that operations can tolerate.

A 50% increase in replenishment lead times raises the reoptimized pool cost to UZS 21 million. Practical decisions should monitor the availability of critical parts and logistics waiting time alongside the average service level. Until validation against enterprise records has been

DOI: 10.65322/j.srd.05192026.96

conducted, these calculations cannot establish demonstrated economic benefits or improvements in technical availability in actual operations.

References

1. Kennedy W. J., Patterson J. W., Fredendall L. D. An overview of recent literature on spare parts inventories. International Journal of Production Economics. 2002. Vol. 76, No. 2. P. 201–215. DOI: 10.1016/S0925-5273(01)00174-8

2. Basten R. J. I., van Houtum G. J. System-oriented inventory models for spare parts. Surveys in Operations Research and Management Science. 2014. Vol. 19, No. 1. P. 34–55. DOI: 10.1016/j.sorms.2014.05.002

3. Huiskonen J. Maintenance spare parts logistics: Special characteristics and strategic choices. International Journal of Production Economics. 2001. Vol. 71, No. 1–3. P. 125–133. DOI: 10.1016/S0925-5273(00)00112-2

4. Driessen M., Arts J., van Houtum G. J., Rustenburg J. W., Huisman B. Maintenance spare parts planning and control: a framework for control and agenda for future research. Production Planning & Control. 2015. Vol. 26, No. 5. P. 407–426. DOI:10.1080/09537287.2014.907586

5. Bacchetti A., Saccani N. Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. Omega. 2012. Vol. 40, No. 6. P. 722–737. DOI: 10.1016/j.omega.2011.06.008

6. Zhang S., Huang K., Yuan Y. Spare Parts Inventory Management: A Literature Review. Sustainability. 2021. Vol. 13, No. 5. Article 2460. DOI: 10.3390/su13052460

7. GOST R 27.102-2021. Dependability in technics. Dependability of an item. Terms and definitions. Effective date: 1 January 2022. Official source (accessed 30 September 2026). No DOI assigned.

8. GOST 18322-2016. Maintenance and repair system of technics. Terms and definitions. Effective date in the Russian Federation: 1 September 2017. The official catalogue record includes the correction published in IUS 11-2025. Official source (accessed 30 September 2026). No DOI assigned.

9. GOST R 53392-2017. Integrated logistic support. Logistic support analysis. General provisions. Replaces GOST R 53392-2009. Effective date: 1 July 2017. Official source (accessed 30 September 2026). No DOI assigned.

10. Croston J. D. Forecasting and Stock Control for Intermittent Demands. Operational Research Quarterly. 1972. Vol. 23, No. 3. P. 289–303. DOI: 10.1057/jors.1972.50

11. Syntetos A. A., Boylan J. E. The accuracy of intermittent demand estimates. International Journal of Forecasting. 2005. Vol. 21, No. 2. P. 303–314. DOI: 10.1016/j.ijforecast.2004.10.001

12. Pinçe Ç., Turrini L., Meissner J. Intermittent demand forecasting for spare parts: A Critical review. Omega. 2021. Vol. 105. Article 102513. DOI: 10.1016/j.omega.2021.102513

13. Guide V. D. R. Jr., Srivastava R. Repairable inventory theory: Models and applications. European Journal of Operational Research. 1997. Vol. 102, No. 1. P. 1–20. DOI: 10.1016/S0377-2217(97)00155-0

14. Sherbrooke C. C. Metric: A Multi-Echelon Technique for Recoverable Item Control. Operations Research. 1968. Vol. 16, No. 1. P. 122–141. DOI: 10.1287/opre.16.1.122

15. Paterson C., Kiesmüller G., Teunter R., Glazebrook K. Inventory models with lateral transshipments: A review. European Journal of Operational Research. 2011. Vol. 210, No. 2. P. 125–136. DOI: 10.1016/j.ejor.2010.05.048