

Article

An Integrated Framework for Internal Replenishment Processes of Warehouses Using Approximate Dynamic Programming

İrem Kalafat ^{1,2,*} , Mustafa Hekimoğlu ¹, Ahmet Deniz Yücekaya ¹, Gökhan Kirkil ¹ , Volkan Ş. Ediger ¹
and Şenda Yıldırım ^{1,2}

¹ Department of Industrial Engineering, Faculty of Engineering and Natural Sciences, Kadir Has University, 34083 Istanbul, Türkiye; mustafa.hekimoglu@khas.edu.tr (M.H.); ahmety@khas.edu.tr (A.D.Y.); gokhan.kirkil@khas.edu.tr (G.K.); volkan.ediger@khas.edu.tr (V.Ş.E.); sendayildirim@outlook.com (Ş.Y.)

² Advanced Analytics Department, Dogus Technology, 34398 Istanbul, Türkiye

* Correspondence: irem.kalafat@stu.khas.edu.tr

Abstract

Warehouses are vital in linking production to consumption, often using a forward–reserve layout to balance picking efficiency and bulk storage. However, replenishing the forward area from reserve storage is prone to delays and congestion, especially during high-demand periods. This study investigates the strategic use of buffer areas—intermediate zones between forward and reserve locations—to enhance flexibility and reduce bottlenecks. Although buffer zones are common in practice, they often lack a structured decision-making framework. We address this gap by developing an optimization model that integrates demand forecasts to guide daily replenishment decisions. To handle the computational complexity arising from large state and action spaces, we implement an approximate dynamic programming (ADP) approach using certainty-equivalent control within a rolling-horizon framework. A real-world case study from an automotive spare parts warehouse demonstrates the model's effectiveness. Results show that strategically integrating buffer zones with an ADP model significantly improves replenishment timing, reduces direct picking by up to 90%, minimizes congestion, and enhances overall flow of intra-warehouse inventory management.

Keywords: forward–reserve; inventory management; ADP; replenishment; spare parts warehouse



Academic Editors: Ran Ji, Jose Machado and Zhengyang Fan

Received: 4 June 2025

Revised: 21 June 2025

Accepted: 26 June 2025

Published: 10 July 2025

Citation: Kalafat, İ.; Hekimoğlu, M.; Yücekaya, A.D.; Kirkil, G.; Ediger, V.Ş.; Yıldırım, Ş. An Integrated Framework for Internal Replenishment Processes of Warehouses Using Approximate Dynamic Programming. *Appl. Sci.* **2025**, *15*, 7767. <https://doi.org/10.3390/app15147767>

Copyright: © 2025 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) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Efficient warehouse management directly influences critical performance metrics such as operating costs, throughput, and customer satisfaction. Among all warehouse activities, order picking is one of the most resource-intensive and cost-sensitive processes, particularly in spare parts supply chains where timely and accurate order fulfillment is paramount [1]. In order to address the challenges, many warehouses employ a forward–reserve area configuration. The forward area serves to stock a limited quantity of spare parts for rapid order picking, while the reserve area holds bulk inventory to replenish the forward area as needed. This internal replenishment practice aims to balance accessibility and storage capacity, but it also introduces operational complexities.

The forward–reserve replenishment process often faces significant challenges, such as delays, congestion, and capacity constraints, particularly in high-demand environments. To alleviate these issues, some warehouses utilize buffer areas as interim holding spaces. The usage of a buffer area is a well-established practice in warehouse management, valued

for enhancing overall warehouse flexibility. Despite the benefits of reduced handling times and improved responsiveness, buffer areas can contribute to congestion if not properly managed. This study introduces a replenishment policy that strategically incorporates buffer areas into forward reserve operations. The policy prioritizes order picking from the buffer area. Additionally, it restricts returning stock to the reserve area to avoid operational disruptions. By focusing on reducing excess buffer inventory, the approach ensures a more efficient flow of goods, maximizing the benefits of buffer areas. Thus, buffer areas not only relieve pressure on reserve areas but also enable faster replenishment of forward areas, enhancing throughput and operational fluidity.

The significance of this study lies in its ability to address a critical gap in warehouse management practices: the absence of structured strategies for buffer area utilization. Although buffer areas are commonly used across industries, they are often implemented reactively as short-term solutions rather than as complementary components of a well-designed replenishment strategy. This study's novelty revolves around the development of a tractable, integrated, and adaptive framework for this complex, real-world problem. The methodological novelty lies in the synthesis of three components: (1) the explicit mathematical modeling of unstructured buffer areas within the replenishment problem, (2) the use of a rolling-horizon mixed-integer linear programming (MILP) model as a high-quality policy evaluation engine, and (3) its integration with a dual-loop, adaptive forecasting module. This combination provides a practical approximate dynamic programming (ADP) framework for optimizing a high-dimensional, stochastic system where traditional dynamic programming is computationally infeasible.

Motivational Case Study

The warehouse system of interest includes a reserve area, forward area, and buffer area. The internal replenishment process in this system is carried out using cages with constant batch sizes. On the other hand, the reserve area consists of narrow corridors that accommodate the bulk storage of all parts. This area is equipped with specialized machines capable of retrieving large cages that carry spare parts, which are not suitable for order picking. Therefore, it is only preferred that parts from this area be picked up if required. In certain instances of replenishment from reserve to forward, the cage is added to the existing inventory on the shelf, potentially exceeding its capacity. In such cases, the warehouse personnel use buffer areas instead of returning the excess inventory to the reserve area. Warehouse managers determine these buffer areas arbitrarily, and they are located next to the shelves without any physical designation as special zones. Their primary purpose is to prevent flow obstruction in the reserve area by avoiding the transportation of partially filled cages. However, if these areas are not effectively managed, they can accumulate a significant amount of stock, leading to congestion in the warehouse. To address this issue, the prioritization of this area is emphasized during the order picking process. Consequently, when inventory is present in both the buffer area and the shelves of the forward area, pickers are instructed to take the stock in the buffer area first. Figure 1 shows the overview of the system.

The findings reveal that incorporating a strategic buffer area policy within the forward–reserve replenishment framework significantly enhances warehouse operations. Key results include a dramatic reduction in direct picking activities up to 90% in some scenarios alongside optimized buffer area utilization and decreased handling costs. The use of demand forecasts, combined with the integration of ADP, enables the system to dynamically adapt to fluctuating demand, enhancing replenishment efficiency even under labor constraints. By showcasing these benefits, this study presents a scalable and adaptable model for optimizing warehouse operations in high-demand, multi-product environments.




Figure 1. Internal replenishment process.

This paper makes an essential contribution to the forward reserve area problem and reveals that internal replenishment processes can be carried out more smoothly with a well-managed buffer area.

The specific objectives of the work in this paper can be further described as follows:

- To overcome the computational challenges of in-warehouse replenishment processes and progressively improving policies by adopting an ADP approach with certainty equivalent.
- To evaluate the impact of the ADP method in managing intuitively placed buffer points for internal replenishment and develop strategies to alleviate warehouse congestion.
- To demonstrate that internal replenishment processes can be carried out more smoothly in a well-managed buffer area.

This paper is structured as follows: Section 2 reviews the relevant literature on the forward–reserve problem and identifies existing gaps. Section 3 details the methodology, including the ADP framework and the MILP model. Section 4 presents the experimental results from different scenarios and analyzes the model’s performance. Section 5 discusses managerial insights, limitations, and future research directions. Finally, Section 6 provides concluding remarks.

2. Literature Review

The forward–reserve area problem has been studied repeatedly in the literature. Ref. [2] solved the problem of selecting items to store in automated storage and retrieval systems and allocating space for inventory keeping. They developed a heuristic solution and applied it to a naval supply center’s small-sized items where it is assumed that they are continuously divisible. This paper is the first to solve the internal replenishment problem: restocking a primary area from a secondary location in the facility using a mathematical model. The model focuses on maximizing profit by allocating sufficient items in the volume base, considering the replenishment cost. Ref. [3] considered busy and idle periods in a distribution center. The busy period means the order picking operations will occur, and the idle period means replenishment operations of the forward area from the reserve area may happen. They also assumed that one trip from the reserve to the forward area is sufficient, replenishment is performed with one unit load, and order picking from the reserve area is possible at a higher cost. They formulated the model as a binary programming problem that minimizes the total labor time and aims to prevent accidents in busy periods. They solved it with greedy knapsack and dynamic programming. Also, a constraint for limiting replenishment capacity is added, and the computational outcomes are compared.

Ref. [4] presented a mathematical model and a heuristic for assigning items and deciding the size of a warehouse’s five functional areas: receiving, shipping, forward, reserve, and cross-docking. They mention that determining the sizes of the areas and assigning products to those areas affect warehouse management on strategic, tactical, and

operational levels. The model is formulated as a mixed-integer linear programming model (MILP) and solved with the proposed heuristic and branch and bound algorithms. The computational results show that the heuristic algorithms take much less time. Ref. [5] applied three storage strategies to the forward area, analyzed the results, and determined which was better. The first one is the equal space strategy, which allocates equivalent space for each item; the other ones are the equal time strategies, which allocate space for each item that is sufficient for the same amount of time, and the optimal strategy. They applied strategies to small parts of the data by extending the model in [2] with new inventory level constraints based on the strategies. Ref. [6] considered a forward reserve area problem in a picker-to-belt system where order pickers collect orders from the forward area and put them on the conveyor. It is assumed that if the pickers face a stockout, the belt stops, and they must perform emergency replenishment. They have compared and evaluated the computational results of replenishment methods, and six heuristic policies have been proposed. Ref. [7] presented the results of early heuristic methodology studies and an alternative branch and bound algorithm that can solve the problem optimally. They applied the algorithms to two different data sets that belong to different warehouses, and the computational performance of the algorithms was presented. It is shown that when the number of stock-keeping units (SKUs) gets smaller, the optimality difference is higher for the aforementioned greedy heuristic in Hackman's study. It is also observed that even though the objective value is equal, the model's outcomes differ between the optimal and the heuristic solutions. Ref. [8] used simulation and what-if analysis to assign items to forward areas and allocate space to them. In order to achieve effective solutions in an acceptable amount of time, several storage policies have been implemented.

In earlier studies, the forward pick area was assumed to be continuously divisible; therefore, the models were fluid. Ref. [9] demonstrated a discrete space allocation for forward reserve problems and analyzed the difference between the discrete and the fluid models. They developed three models which serve different purposes. The first model solves the problem of dividing the forward area for a given set of items, which is defined in the paper as a discrete forward-reserve allocation problem (DFRAP). The second model unifies the allocation and assignment problem, and the model's outcome shows which items must be held in the forward area and the size that belongs to the particular items. The final model has an approach to the problem that considers the forward area size as a variable and how the products must be allocated. The results showed that the gap between the discrete and the fluid models is small and can be ignored.

The case study presented by [10] has been applied to a cosmetic firm's warehouse where they use a wave picking strategy. The objective is to avoid stockouts in an environment where picking and replenishment operations are performed simultaneously. They proposed three prioritizing policies for replenishments and simulated to compare the computational results and stockouts. Ref. [11] proposed an alternative storage allocation policy to reduce the picking time for the picking area of a warehouse. The proposed policy assumes several empty locations in the picking area that can be used for items whose requested quantities are higher than those in their stock. The study compared the proposed policy with the conventional warehouse layout by simulation and analyzed its performance regarding operator travel and order fulfillment time. The study also considered the impact of congestion on the proposed policy. The paper provides insights to improve the performance of the warehouse order-picking system.

Ref. [12] identified the best storage strategies for an automated storage and retrieval system. The research used analytical and numerical methods to compare the performance of different storage strategies. The paper also presents the factors affecting the performance of storage strategies and the experimental results used to determine the best strategies.

Ref. [13] aimed to design efficient work schedules for human pickers in mobile rack warehouses using human–robot co-coordinated order picking systems. The article discusses the challenges of designing good work schedules for human pickers in such systems. It proposes a model that allows mobile racks with different workloads to be assigned to pickers and schedules the racks assigned to each picker to minimize the expected total picking time. The paper uses a stochastic dynamic programming model and an approximate dynamic programming-based branch-and-price solution approach to solve the problem. The results show that the proposed approach can solve a moderate-sized problem of 50 racks in less than 2 min and produces high-quality solutions with picking times that are 10% shorter than those that do not account for schedule-induced fluctuations in the pickers' work states. The article concludes that the proposed approach can be applied to other warehouses to detect fluctuations in picker work states, incorporate these fluctuations into robot schedules, and increase the productivity of a picking system.

Ref. [14] considered the storage replenishment and routing problem (SRRP) in a warehouse environment where items are stored in separate forward and reserve storage areas and replenished using a common reserve storage area. The authors propose an MIP model and heuristic approaches to solve the problem. Based on the inventory routing problem (IRP) literature, the authors propose an MIP model to solve SRRP. They also propose a priori routing heuristics to solve the problem. These heuristics are based on the shortest path algorithm and graph theory-based heuristics. To further improve the replenishment routes, they propose an a posteriori routing step. Ref. [15] examined the ADP method for solving multidimensional bag problems. ADP provides fast and effective solutions to complex and large-scale optimization problems using approximate calculations of the value function. The results show that this approach provides high-quality and practical solutions, especially under multidimensional constraints. Ref. [16] developed an optimization tool intending to replace Zara's manual and experience-based pricing process with a more systematic approach based on demand forecasting and price optimization models. In the study, the certainty equivalent approach was used by substituting expected values for uncertain future sales, and an MIP model was formulated with this approach. The developed model was tested in controlled field experiments in Zara's Irish and Belgian stores, and this new process increased revenue from discount sales by 6%. As a result, this method began to be used in Zara's discount decisions worldwide.

A broader view of buffer management reveals its strategic importance across industries. In manufacturing, buffers act as decoupling points in Just-in-Time (JIT) systems, absorbing variability to maintain a smooth production flow [17]. In supply chain management, inventory buffers are strategically placed to hedge against demand and lead-time uncertainty, as analyzed in the seminal work by Simchi-Levi et al. [18]. While the function of these formal, planned buffers is well-understood, the strategic management of informal, dynamic buffer zones within a warehouse—which arise from operational constraints, as addressed in our study—remain underexplored in the literature.

Recent developments in warehouse automation, such as the deployment of Automated Storage and Retrieval Systems (AS/RS) and Autonomous Mobile Robots (AMRs), are changing the landscape of internal logistics. A comprehensive review in [19] shows that these technologies excel at executing physical movements efficiently. However, they require a high-level strategy to provide the decision logic for what to move, when to move it, and where to place it. The hierarchical nature of such systems, where strategic planning directs tactical execution, is a key theme in the automation literature. Our ADP framework is designed to serve as this strategic layer, providing optimized decisions that can then be passed to an automated system's execution layer, thereby addressing a critical need in the evolution towards smart warehouses.

In the context of evaluating stochastic optimization models like ours, it is a standard practice to benchmark performance against a theoretical optimum. This is often achieved by solving the problem with perfect future information, a method known in the literature as a “dynamic oracle” or perfect information benchmark [20]. This oracle provides a theoretical upper bound on performance, and the gap between the stochastic solution and the oracle’s performance quantifies the “cost of uncertainty” or the value of perfect information. This benchmark is crucial for contextualizing the performance of any forecast-driven model.

Despite this extensive body of work, a critical gap remains in the literature regarding the strategic and optimized use of buffer areas within the forward–reserve replenishment context. While buffers are acknowledged as operational tools, they are rarely integrated as a formal component of a dynamic, forecast-driven optimization model. Most studies focus on the direct link between the forward and reserve areas, overlooking the potential of a managed intermediate zone to absorb variability and reduce congestion.

This study fills this gap by explicitly modeling the buffer area and using an ADP framework to make intelligent, proactive decisions about its use, thereby justifying its contribution to the field. Alternative optimization methods, such as purely deterministic models, often struggle with the uncertainty of real-world demand, while simple heuristics may fail to find high-quality solutions. Our ADP approach provides a balance of computational tractability and responsiveness to dynamic conditions.

3. Methodology

3.1. Application of Approximate Dynamic Programming (ADP)

This study utilizes ADP to manage and optimize the decision-making process under uncertainty to address the complexities of internal replenishment in a spare parts warehouse. ADP is particularly suitable for environments where the state and decision spaces are too large for traditional dynamic programming (DP) techniques due to the curse of dimensionality. It simplifies the problem by approximating the value function, which assesses future costs associated with different states and decisions. While a formal mathematical proof of convergence for this type of applied framework is beyond the scope of this paper, we demonstrate its empirical stability in Section 4.3, where results are shown to be consistent across varying planning horizons, suggesting the policy does not diverge. The DP formulation for this problem is presented in Appendix A.

1. Demand Forecasting

The efficacy of the ADP framework is fundamentally dependent on the quality of its demand forecasts, which serve as the primary input for decision-making. To ensure predictive reliability, we conducted a comparative analysis of state-of-the-art machine learning models, including Extreme Gradient Boosting (XGBoost) and LightGBM, using time-series cross-validation. Evaluation metrics included Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). LightGBM was ultimately selected due to its superior performance across a holdout test set. To maintain robustness in a dynamic warehouse environment, the model operates within a dual-loop learning framework:

- **Daily Operational Loop:** At the end of each day, actual demand is captured and used to update the inventory state for the next day’s MILP optimization run.
- **Periodic Retraining Loop:** The same demand data is added to the historical training set, and the LightGBM model is retrained periodically to reflect evolving demand patterns.

This adaptive retraining strategy ensures that the forecasts remain aligned with changes in seasonality, customer behavior, or product life cycles. Inaccurate forecasts can lead to over-replenishment or stockouts, especially in a cost-sensitive MILP environ-

ment; hence, this mechanism is vital for sustaining high-quality replenishment decisions over time.

2. ADP Model with Certainty Equivalent Approach

The ADP model incorporates state variables for inventory levels, decision variables for replenishment and picking, and an objective to minimize total expected costs. Value function approximation is achieved by implementing a certainty-equivalent control policy, a practical approach for complex problems chosen here for its interpretability, data efficiency, and stability. In this approach, the stochastic future demand is replaced by deterministic point forecasts from the machine learning model. This simplification allows the problem to be resolved at each step by the finite-horizon MILP. The solution to this MILP provides the optimal actions, and its objective value serves as the practical approximation of the true value function for the current state. This method of using a deterministic optimization model within a larger stochastic framework is a practical approach for complex problems, and a similar strategy was successfully used by Caro et al. [16].

While this certainty-equivalent approach simplifies the problem at each step, uncertainty is managed over the long term through a rolling-horizon implementation. This mechanism is clarified by the daily operational cycle:

- a. Forecast and Optimize: At the start of the day, we generate a forecast for the planning horizon. The MILP is solved using the current inventory state, and this forecast produces an optimal plan.
- b. Implement: Only the decisions for the immediate first day are implemented.
- c. Observe and Update: At the end of the day, we observe the actual demand and update the inventory levels to reflect the true state of the system.
- d. Repeat: On the next day, the entire process repeats from the new, true state. The old plan is discarded, and a new optimization is performed.

This constant cycle of re-planning based on real-world feedback allows the system to continuously correct its course, effectively managing uncertainty without needing to model the full probability distribution of demand.

3. Rolling-Horizon Implementation

The implementation of ADP follows a rolling-horizon approach, which is particularly effective for managing operations in dynamic environments like warehouses. The optimization model is performed over discrete time periods such as daily. The decisions are made at the beginning of each period based on the system's current state using demand predictions as inputs. At the end of each period, the state information is updated based on the actual demand fulfilled and replenishment actions taken. This updated state information is then used to make decisions in the next period. A feedback loop is incorporated whereby the outcomes of previous decisions inform updates of the value function approximations, refining the decision-making process over time.

To effectively apply ADP, certain assumptions are made about the operational context, including but not limited to the predictability of demand variations, the reliability of supply chain logistics, and the stability of replenishment lead times. These assumptions help simplify the model while ensuring it remains robust enough to handle real-world complexities. This detailed application of the ADP model with certainty equivalence seeks to optimize internal replenishment processes by effectively handling uncertainties in demand and supply, thus enhancing operational efficiency and reducing costs in the spare parts warehouse setting.

3.2. ADP Pseudo Code

The core of our framework is the iterative solving of the MILP model within the ADP loop. The MILP serves as the policy function in our ADP framework. It is solved iteratively, once per day, in our rolling-horizon simulation. At the beginning of each day t , the current state (inventory levels) provides the initial boundary for the MILP. The MILP is then solved to find an optimal set of actions over the planning horizon.

In this approach, the MILP's finite-horizon objective function acts as the practical approximation for the theoretical ADP value function, $V(S)$. Since computing $V(S)$ for all possible states is intractable, we instead approximate the total cost-to-go by solving the detailed MILP over a finite lookahead period. This is a common and effective technique in ADP, allowing us to leverage the power of exact optimization for tactical planning within a larger, adaptive framework.

Step 0: Initialization

- 0.1 Load historical demand data from the database;
- 0.2 Initialize state variables: forward inventory, buffer inventory and reserve inventory;
- 0.3 Define the planning horizon T (e.g., the next seven days);
- 0.4 Set MILP parameters.

Step 1: Main Loop

For each spare part ($i = 1, 2, 3, \dots, N$):

For each planning day ($t = 1, 2, \dots, T$):

Step 1.1: Generate forecasts (certainty equivalent);

Step 1.2: Solve MILP model with recursive state transitions (1)–(17).

Step 2: Update Initial States (Inventory levels)

- 2.1 Record the actual demand at the end of the day;
- 2.2 Update the initial states and historical data with the recorded actual demand (iteratively refining value approximation).

The ADP process begins with Step 0: Initialization, where historical demand data is loaded, state variables are set, and MILP parameters are configured. This establishes the foundation for the iterative optimization. In Step 1: Main Loop, demand forecasts are treated as fixed values and serve as inputs to the MILP model, which determines optimized replenishment and buffer allocation decisions over the finite planning horizon. By solving the MILP iteratively each day, the approach captures the dynamic nature of inventory management while maintaining computational efficiency.

In Step 2: Update Initial States, when the actual demand at the end of the day is observed, we update the initial states of the model for the next day's run. These iterative updates continually improve the accuracy of the model's future decisions. It is important to note that optimal decisions are made not with a complex, intractable DP model but with the simple MILP model that we run continuously. This demonstrates that ADP offers a much more practical solution than traditional DP. While this methodology does not solve the DP problem exactly, it uses ADP principles to approximate the value function and optimize decisions iteratively. The rolling-horizon framework recalibrates the optimization at each step, balancing immediate and future costs dynamically. Although this method does not provide a truly optimal solution, it provides a much more feasible and practical solution, avoiding the difficulties of DP. The DP formulation of this problem and the DP pseudo code are explained in detail in Appendix A.

3.3. MILP Model for Forward Area Replenishment Problem with Buffer Area

A MILP model has been developed to address the specific challenges presented in this paper. This model is designed to optimize decision-making by capturing the complex dynamics and constraints of the problem. The model is structured around a key policy that prioritizes order picking from the buffer area whenever the required spare part is available there. By giving precedence to buffer stock, we aim to decrease the inventory in the buffer area, reducing congestion and improving operational flow. If the buffer stock is insufficient, the remaining quantity is picked from the forward area, ensuring that the forward area is replenished only when necessary. This policy-driven approach has demonstrated clear benefits in improving warehouse efficiency by reducing unnecessary movements in warehouse and streamlining the replenishment process.

In developing the mathematical model, certain assumptions related to warehouse operations were made to reflect realistic business practices. These assumptions include the operational characteristics and warehouse policies that are critical to this study, as outlined below.

1. Each spare part has a specific batch size for replenishment.
2. Due to the involvement of a limited number of staff members in the internal replenishment process, the number of replenishments that can be accomplished daily is limited.
3. The stock needed to cover the demand is available in the warehouse in the forward, reserve, or buffer areas.
4. A dedicated rack with a capacity of a specific batch size is allocated in the forward area for each part.
5. A complete cage is transferred from the reserve area to the forward area for replenishment purposes, i.e., all transferred cages are full.
6. If there is any surplus beyond the forward area's capacity, the remaining stock must be located in the buffer area, which is adjacent to the shelf.
7. If a cage is transferred from the reserve area while some parts exist in the associated shelf in the forward area, this partially filled cage in the forward area is transferred to the buffer area, and the new full cage is located in the forward area.
8. In cases where the order picker exclusively visits the forward area and encounters insufficient stock, they proceed to the reserve area to fulfill the order and perform direct picking.

The mathematical model explained in this paper will undergo comprehensive examination alongside the notations employed for delineating the internal replenishment problem. The methodologies for optimization utilized to address the issue will be explained. The notations that will be used in the mathematical model are given in Abbreviations.

The framework aims to efficiently ascertain the inventory levels of shelves and enhance product positioning. For this objective, the mathematical expression of the framework and the resolution methods are elaborated thoroughly. The MILP model details are given below.

Model:

$$\min \sum_{i=1, t=1}^{|N|, T} Q_{it} c^r + C_{it}^D c^d + (WAR_{it} - DW_{it} + WBR_{it}) c^w + X_{it}^W c^{wx} \quad (1)$$

Subject to:

$$XAR_{it} = XBR_{it} + QN_{it} C_i - X_{it}^W \quad i \in N, 1 \leq t \leq T \quad (2)$$

$$XBR_{it+1} = XAR_{it} - D_{it}^S \quad i \in N, 1 \leq t \leq T \quad (3)$$

$$WAR_{it} = X_{it}^W + WBR_{it} \quad i \in N, 1 \leq t \leq T \quad (4)$$

$$WBR_{it+1} = WAR_{it} - D_{it}^W + C_{it}^D C_i - D_{it}^P \quad i \in N, 1 \leq t \leq T \quad (5)$$

$$X_{it}^W \leq Q_{it} C_i \quad i \in N, 1 \leq t \leq T \quad (6)$$

$$QN_{it} \leq Q_{it} M_i \quad i \in N, 1 \leq t \leq T \quad (7)$$

$$QN_{it} C_i + XBR_{it} - X_{it}^W \leq F_i \quad i \in N, 1 \leq t \leq T \quad (8)$$

$$D_{it}^W + D_{it}^S + D_{it}^P = D_{it} \quad i \in N, 1 \leq t \leq T \quad (9)$$

$$X_{it}^W \leq XBR_{it} \quad i \in N, 1 \leq t \leq T \quad (10)$$

$$D_{it}^P \leq C_{it}^D C_i \quad i \in N, 1 \leq t \leq T \quad (11)$$

$$D_{it}^W \leq WAR_{it} \quad i \in N, 1 \leq t \leq T \quad (12)$$

$$QN_{it} \leq C_i F_i \quad i \in N, 1 \leq t \leq T \quad (13)$$

$$XAR_{it} \leq F_i \quad i \in N, 1 \leq t \leq T \quad (14)$$

$$\sum_{i=1}^N Q_{it} \leq \tau \quad i \in N, 1 \leq t \leq T \quad (15)$$

$$XBR_{it}, XAR_{it}, WBR_{it}, WAR_{it}, QN_{it}, X_{it}^W, D_{it}^S, D_{it}^W, D_{it}^P, C_{it}^D \geq 0 \quad i \in N, 1 \leq t \leq T \quad (16)$$

$$Q_{it} \in \{0, 1\} \quad (17)$$

The objective function (1) minimizes overall costs, which include the costs associated with internal replenishment (c^r), direct picking (c^d), inventory holding in the buffer area (c^w), and transferring inventory from the forward area to the buffer area (c^{wx}). Based on Assumption 7, the cost of transferring inventory from the forward area to the buffer area is incorporated to prevent unnecessary handling activities.

In this model, the cost structure is defined with the assumption that direct picking incurs the highest cost, followed by forward-to-buffer area transfer, then buffer area usage, with replenishment having the lowest cost. This hierarchy reflects the labor intensity and resource requirements of each operation.

Constraints (2) and (3) determine the stock levels in the forward area based on the inventory from the preceding time period and the replenishment quantity within the current timeframe. Constraints (4) and (5) model the inventory movement between the buffer and forward areas. Constraints (6) and (7) are the either-or constraints that help the variables X_{it}^W , QN_{it} and the binary variable Q_{it} to be assigned correctly. In Equation (18), the replenishment decision is also shown. This binary decision is critical as it triggers the sequence of subsequent inventory adjustments and movements.

$$Q_{it} = \begin{cases} 1, & \text{if forward area is replenished} \\ 0, & \text{if forward area is not replenished} \end{cases} \quad (18)$$

Equation (19) shows the process of stock relocation from the forward area to the buffer area within a warehouse. If the transfer meets certain capacity conditions, the binary variable is set to 1, indicating replenishment. If no stock is moved, the variable is set to 0.

$$Q_{it} = \begin{cases} 1, & \text{if } C_i \geq X_{it}^W > 0 \\ 0, & \text{if } X_{it}^W = 0 \end{cases} \quad (19)$$

Function (20) shows the logic of the inclusion of QN_{it} is particularly important as it integrates the replenishment decision with the physical quantity of inventory moved, ensuring that our model reflects the practical constraints of the warehouse's operational capacity. Together, these equations fine-tune our inventory management and enhance

the overall efficiency of warehouse operations, leading to a reasonable balance between availability and the cost implications of inventory holding and movement.

$$Q_{it} = \begin{cases} 1, & \text{if } M_i \geq QN_{it} > 0 \\ 0, & \text{if } QN_{it} = 0 \end{cases} \quad (20)$$

Constraint (9) ensures that the total demand is met through a combination of stocks from the forward area, the buffer area, and direct picking. Constraint (10) ensures that the inventory transferred from the forward area to the buffer area does not exceed the existing stock in the forward area. Constraint (11) ensures that the cages required for direct picking are adequately supplied. Constraint (12) limits the amount of demand covered from the buffer area to the available stock within that area.

Constraints (8), (13), and (14) ensure that the quantity of parts brought from the reserve area for replenishment is limited to the capacity of the forward area. Finally, constraint (15) ensures compliance with the daily replenishment capacity, preventing overstocking and optimizing resource allocation. The objective function and the constraint set in (1)–(20) constitute our mathematical model for effectively managing warehouse inventory.

4. Results and Analysis

The results presented in this section are based on a one-month simulation conducted for spare parts groups of 50, 200, and 300 items. To provide robust validation for our findings, we performed 30 independent simulation runs for each scenario. In each run, the actual demand was generated stochastically to reflect real-world uncertainty, while the model only had access to its deterministic forecast. Different random seeds were used to ensure demand variability across the runs. All reported metrics for these scenarios show the mean and the 95% confidence interval. The proposed optimization model's performance is evaluated against three distinct benchmarks:

- **(s,S) Policy:** A standard inventory control policy from the academic literature, serving as a rigorous algorithmic benchmark.
- **Business as Usual:** The current, manual practice in the case study warehouse, serving as a practical baseline.
- **Simulation with Known Demand (Simulation KD):** A theoretical, perfect-information case, corresponding to the “dynamic oracle” benchmark discussed in the literature [20], used to establish an upper bound on performance.

Our proposed model, **Simulation with Predicted Demand (Simulation PD)**, is driven by the ADP framework with machine learning forecasts. The findings illustrate the benefits of integrating predictive demand models and buffer area strategies into warehouse operations.

4.1. Metrics and Scenarios

The performance of the optimization model was evaluated using four key metrics:

1. **Direct Picking:** This represents the percentage of items picked directly from the reserve area, bypassing forward and buffer areas. A lower percentage indicates a more streamlined replenishment process.
2. **Replenishment:** This tracks the frequency of stock restocking activities, ensuring forward areas remain adequately stocked to meet demand.
3. **Buffer Area Usage:** This measures the percentage of inventory stored temporarily in buffer areas. Lower usage reflects improved stock flow without reliance on intermediary storage.
4. **Forward-to-Buffer Transfers:** This reflects the movement of inventory from forward areas to buffer areas due to capacity constraints or imbalances.

4.2. Comparative Analysis of Handling Activities

The proposed solution framework was applied to a case with all part groups, and the results are presented in Table 1. The most striking achievement is the dramatic and statistically significant reduction in direct picking activities. For the 50-part group, our proposed model (Simulation PD) averages 13.8% direct picking, which is a substantial improvement over both the 19.5% from the standard (s,S) policy and the 28.1% from Business as Usual. This signifies a more streamlined picking process that translates into faster fulfillment and reduced labor costs—key factors for warehouse efficiency. Another notable outcome is the optimized use of buffer space, a critical resource in managing inventory flow. While buffer usage under Business as Usual is excessively high (56.8%), our model maintains a healthier balance (18.9%) compared to the more reactive (s,S) policy (25.4%), ensuring that space is used only when truly necessary to avoid unnecessary congestion and enhancing throughput. These performance trends hold and are often amplified as the part group size increases. For instance, in the 300-part group, our model reduces direct picking to just 5.5%, compared to 8.9% for the (s,S) policy and 19.2% for Business as Usual, demonstrating the scalability of the model's benefits. In essence, these results highlight the potential for real-world implementation, promising significant gains in operational performance and efficiency over both current practice and standard academic benchmarks.

Table 1. Handling activities for different part groups (50, 200, 300 parts).

Part Group	Scenario	Direct Picking (%)	Replenishment (%)	Buffer Area Usage (%)	Forward-to-Buffer Area Transfer (%)
50 Parts	Simulation PD	13.83 ± 0.7	33.75 ± 1.1	18.91 ± 1.5	15.75 ± 1.2
	(s,S) Policy	19.5 ± 1.1	29.1 ± 1.8	25.4 ± 2.1	18.3 ± 1.9
	Business as Usual	28.06 ± 1.3	32.47 ± 1.4	56.78 ± 2.5	23.98 ± 2.0
	Simulation KD	2.66	36.16	15.41	22.5
200 Parts	Simulation PD	6.83 ± 0.4	17.85 ± 0.9	7.708 ± 0.8	6.5 ± 0.6
	(s,S) Policy	10.2 ± 0.8	15.5 ± 1.2	12.1 ± 1.4	9.8 ± 1.0
	Business as Usual	22.78 ± 1.2	17.41 ± 1.0	46.05 ± 2.8	12.01 ± 1.1
	Simulation KD	2	18.66	0.52	11.08
300 Parts	Simulation PD	5.541 ± 0.3	14.55 ± 0.8	5.70 ± 0.7	4.75 ± 0.5
	(s,S) Policy	8.9 ± 0.6	13.1 ± 1.0	9.5 ± 1.1	7.2 ± 0.8
	Business as Usual	19.15 ± 1.1	14.24 ± 0.9	43.88 ± 2.9	10.30 ± 1.0
	Simulation KD	3.11	14.33	1.02	7.83

4.3. Summary of Handling Activities: Total and Average Metrics

The findings presented in Table 2 offer valuable insights into the impact of optimization on warehouse handling activities. The key takeaway is the significant reduction in direct picking in our proposed Simulation PD compared to both the (s,S) policy and Business as Usual. This illustrates the efficiency gains achieved through predictive and demand-driven approaches, particularly Simulation KD, which leverages known demand to minimize unnecessary direct picking, one of the most labor-intensive and costly activities in warehouse operations. The results also highlight the role of replenishment in improving operational efficiency. In Simulation KD, higher replenishment rates indicate that the system is replenishing stock at more optimal times and quantities, which in turn reduces the reliance on direct picking. The minimal increase in replenishment activities demonstrates that this process is not overly burdensome and instead adds to the overall efficiency by ensuring the forward area is adequately stocked.

Table 2. Sum and mean of handling activities for different part groups (50, 200, 300 parts).

Part Group	Scenario	Direct Picking (Sum/Mean)	Replenishment (Sum/Mean)	Buffer Usage (Sum/Mean)	Forward-to-Buffer Transfer (Sum/Mean)
50 Parts	Simulation PD	18,822/15.69 ± 0.9	405/0.34 ± 0.03	3067/2.55 ± 0.21	3231/2.69 ± 0.24
	(s,S) Policy	26,520/22.1 ± 1.2	361/0.3 ± 0.04	57,720/48.1 ± 4.5	4944/4.12 ± 0.38
	Business as Usual	42,633/35.49 ± 1.5	390/0.32 ± 0.03	131,074/109.13 ± 5.1	23,479/19.54 ± 1.7
	Simulation KD	9711/8.09	434/0.36	15,658/13.04	3750/3.13
200 Parts	Simulation PD	32,381/6.74 ± 0.5	857/0.17 ± 0.02	4623/0.96 ± 0.1	5142/1.07 ± 0.11
	(s,S) Policy	47,040/9.8 ± 0.7	768/0.16 ± 0.02	7200/1.5 ± 0.18	6969/1.45 ± 0.15
	Business as Usual	66,492/13.84 ± 0.9	836/0.17 ± 0.02	478,690/99.70 ± 8.2	55,486/11.55 ± 1.0
	Simulation KD	26,842/5.59	896/0.19	234/0.05	9296/1.94
300 Parts	Simulation PD	33,124/4.60 ± 0.4	1044/0.15 ± 0.02	5592/0.78 ± 0.09	5879/0.81 ± 0.08
	(s,S) Policy	48,600/8.1 ± 0.6	936/0.13 ± 0.02	8280/1.15 ± 0.14	7560/1.05 ± 0.12
	Business as Usual	66,611/9.25 ± 0.7	1026/0.14 ± 0.02	302,231/41.97 ± 4.5	46,077/6.39 ± 0.6
	Simulation KD	27,302/3.79	1032/0.14	237/0.03	10,641/1.47

The buffer area plays a crucial role in operational efficiency, especially in smaller part groups. In the 50-part scenario, Simulation KD shows a higher reliance on buffer usage, as this model optimizes buffer space to balance demand fluctuations. As the part group size increases, buffer usage in Simulation KD decreases, suggesting that for larger part groups, the optimization model relies more on strategic replenishment and less on buffer storage. This reflects the flexibility of the optimization model, adjusting to the scale of the operation and using buffers only when necessary. If the buffer area were not used at all, we would expect a significant increase in direct picking activities and overall handling costs. Without the buffer, the system would rely more on direct fulfillment from the forward or reserve areas, leading to inefficiencies and increased labor costs. The results make it clear that buffer areas, when used strategically, help smooth out inventory management, reduce unnecessary work, and improve overall efficiency. In summary, the findings show that the use of optimization, particularly Simulation KD, enhances warehouse operations by reducing direct picking and strategically utilizing the buffer area.

The buffer's role diminishes in larger part groups due to more effective replenishment planning, but its presence still adds value to smaller-scale operations. This demonstrates that optimized replenishment and buffer management can significantly improve warehouse efficiency, minimizing costs and unnecessary handling activities across varying scales of operation.

4.4. Time Horizon Impact on Operations

Table 3 illustrates how handling activities vary across different planning horizons (1 week, 2 weeks, and 3 weeks) for the 50-part group. The results show that the optimization model maintains consistent performance regardless of the length of the planning period. For direct picking, the percentage remains stable across all planning horizons, with only a slight decrease from 13.83% in the 1-week and 2-week horizons to 13.73% in the 3-week horizon. This stability suggests that the optimization model is effective in consistently managing direct picking activities, regardless of the planning period. Replenishment activities also show minimal variation, staying nearly identical across the different horizons, with a minor increase to 33.97% in the 3-week horizon. This indicates that the model can reliably manage replenishment needs over extended periods without significant changes in the handling activity percentage. Buffer area usage and forward-to-buffer area transfers exhibit

a similarly stable pattern, with only slight variations as the planning horizon increases. For instance, buffer area usage increases marginally from 18.91% in the 1-week and 2-week horizons to 19.067% in the 3-week horizon. This consistency across metrics underscores the robustness of the optimization model in handling inventory activities effectively, even as the planning horizon extends.

Table 3. Handling activities for different planning horizons with predicted demands (50 parts).

Planning Horizon	Direct Picking (%)	Replenishment (%)	Buffer Area Usage (%)	Forward-to-Buffer Area Transfer (%)
1 Week	13.83 $\pm$ 0.7	33.75 $\pm$ 1.1	18.91 $\pm$ 1.5	18.91 $\pm$ 1.2
2 Weeks	13.83 $\pm$ 0.7	33.75 $\pm$ 1.1	18.91 $\pm$ 1.5	18.91 $\pm$ 1.2
3 Weeks	13.73 $\pm$ 0.7	33.97 $\pm$ 1.1	19.067 $\pm$ 1.5	19.067 $\pm$ 1.2

Overall, Table 3 demonstrates that the optimization model's performance remains steady over different planning periods, ensuring reliable and efficient handling activities regardless of the planning horizon.

Table 4 presents the sum and mean values of handling activities across different planning horizons (1 week, 2 weeks, and 3 weeks) for the 50-part group. The key takeaway is that both the total handling activities and average efficiency remain consistent, regardless of the planning period. Direct picking activities exhibit minimal variation, with a slight decrease in the 3-week horizon, indicating that the model effectively manages operations over longer periods without significant changes in the workload. Similarly, replenishment activities maintain stable sum and mean values across all planning horizons, suggesting that the optimization model ensures consistent efficiency. While buffer usage shows a minor increase over the 3-week horizon, its impact on overall performance is minimal. In summary, the results in Table 4 highlight the model's effectiveness in maintaining consistent handling activities and operational efficiency across varying planning horizons.

Table 4. Sum and mean of handling activities for different planning horizons (50 parts).

Planning Horizon	Direct Picking (Sum/Mean)	Replenishment (Sum/Mean)	Buffer Usage (Sum/Mean)	Forward-to-Buffer Transfer (Sum/Mean)
1 Week	18822/15.685 $\pm$ 0.9	405/0.3375 $\pm$ 0.03	3067/2.55 $\pm$ 0.21	3231/2.692 $\pm$ 0.24
2 Weeks	18822/15.685 $\pm$ 0.9	405/0.3375 $\pm$ 0.03	3067/2.55 $\pm$ 0.21	3231/2.692 $\pm$ 0.24
3 Weeks	18710/15.592 $\pm$ 0.9	409/0.341 $\pm$ 0.03	3095/2.575 $\pm$ 0.21	3241/2.701 $\pm$ 0.24

5. Discussion

5.1. Cost and Efficiency Trends

Total cost trends were observed over time in four scenarios—Simulation PD, (s,S) Policy, Simulation KD, and Business as Usual—for three different parts groups (50, 200, and 300 parts). This analysis provides valuable information about the performance of each scenario and the effectiveness of the optimization models. It highlights how each approach impacts costs, particularly in the context of an automotive parts warehouse.

Figure 2 presents the total daily costs for the scenarios of the 50-part group in a line graph. The Simulation KD scenario, which has the perfect knowledge of demand, consistently outperforms all other approaches, establishing a theoretical lower bound on costs. This outcome underscores the significant advantage of accurate demand forecasting in inventory management. The Simulation KD scenario's lower total costs suggest that

when demand can be known with certainty, inventory systems can be optimized to a level of efficiency that minimizes the total cost, likely through more effective stock positioning and reduced necessity for last-minute adjustments that typically incur higher costs. In contrast, the Simulation PD scenario, which relies on predicted demand, yields a total cost profile that exceeds that of the Business as Usual and (s,S) Policy scenarios at various points. This observation suggests that while predictive models aim to improve standard practices, the inherent uncertainty and potential inaccuracies in demand forecasting may lead to suboptimal decision-making that can inadvertently increase costs. It is important to note, however, that the higher costs associated with the Simulation PD model do not necessarily denote a flawed model; instead, they may reflect the complex trade-offs between different operational objectives, such as service level targets, risk of stockouts, and the cost of holding inventory. Moreover, the peaks in the cost graph for Simulation PD could indicate moments where the predictive model may have over- or under-estimated demand, leading to increased costs due to excess stock or emergency restocking. These variations in the Simulation PD scenario emphasize the critical role of forecasting accuracy and the potential cost implications of deviations from actual demand.




Figure 2. Total cost comparison over time for 50 parts.

Notably, the (s,S) Policy scenario falls consistently between Simulation PD and Business as Usual, validating its role as a reasonable academic benchmark but highlighting its limitations in adapting to dynamic demand. Meanwhile, the Business as Usual approach remains the most costly throughout, with significant variability and frequent cost spikes, illustrating the inefficiencies of manual decision-making and lack of systematic forecasting.

This graph does not dismiss predictive models' value but highlights the complexity and challenges associated with demand forecasting. It also illustrates the substantial benefits realized when demand is known precisely, presenting a best-case scenario for optimization. Thus, these findings serve as a catalyst for investing in the accuracy and robustness of demand forecasting methods within predictive models to bridge the gap between the idealized outcomes of known demand and the fluctuating results of predicted demand scenarios.

For the 200-part group, it is shown in Figure 3 that the Simulation PD scenario fluctuates significantly and, like in the 50-part set, sometimes exceeds the Business as Usual scenario in terms of total cost. This could indicate that the predictive model may struggle with scale, where the complexity of managing a more extensive inventory set might amplify the cost of inaccuracies in demand forecasting. The Simulation KD scenario maintains a significant advantage over other approaches, reinforcing that accurate demand knowledge allows for superior inventory optimization. The costs are consistently lower across the observed time frame, suggesting scalable efficiency gains when demand is precisely known. Business as Usual exhibits the highest costs among the three scenarios, which may signal

inefficiencies that are magnified as the inventory set increases. This scenario appears to lack the dynamic adjustments possible in the simulated scenarios, leading to less optimized cost outcomes. The (s,S) Policy provides a consistently moderate performance, superior to Business as Usual but generally inferior to Simulation PD and far from the efficiency of Simulation KD. Its performance reflects a rule-based approach that, while systematic, lacks responsiveness to real-time demand dynamics.




Figure 3. Total cost comparison over time for 200 parts.

As is shown in Figure 4, in the 300-part group, the Simulation PD scenario again shows variability in total costs, with peaks potentially reflecting moments of over- or under-estimation of demand by the predictive model. This variability seems less pronounced compared to the 200-part inventory, suggesting some degree of scalability in the model's predictive accuracy or a possible indication of the model's adaptive mechanisms that become more effective with larger data sets. Simulation KD continues to outperform the other scenarios, suggesting that the benefits of known demand in cost optimization are evident and consistent across different inventory sizes. The model's cost efficiency is apparent, with the lowest costs throughout, indicating that the ability to forecast demand directly correlates with cost savings precisely. Business as Usual remains relatively stable but at a higher cost level compared to Simulation KD, yet it occasionally dips below the costs observed in Simulation PD. This might reflect the inherent inefficiencies in a system that does not utilize optimization models to respond to changing demand patterns.




Figure 4. Total cost comparison over time for 300 parts.

Forward reserve area replenishment optimization analysis highlights consistent trends and critical insights across all part groups, namely 50, 200, and 300. Simulation PD consistently demonstrates lower direct picking frequencies than Simulation KD and Business as Usual scenarios. This pattern suggests the efficacy of predictive demand optimization in reducing direct picking occurrences, thereby potentially enhancing operational efficiency and cost-effectiveness.

Furthermore, Simulation KD consistently exhibits lower overall costs allocated to direct picking and higher costs allocated to pulling inventory from forward to buffer areas compared to Simulation PD and Business as Usual. This indicates a strategic allocation of resources towards optimizing inventory movement scenarios with known demand, potentially resulting in more streamlined operations and reduced overhead costs.

In contrast, the Business as Usual, consistently relies on traditional inventory management practices, with higher direct picking frequencies and costs allocated to buffer area usage. This suggests a potential opportunity for optimization through the adoption of predictive analytics and more efficient inventory management strategies.

The cost profile for the Simulation PD scenario highlights the practical limits of forecast-driven optimization. The observed cost spikes, particularly when compared to the stable, low costs of the Simulation KD scenario, can be attributed to the model's reaction to forecast inaccuracies. When the predicted demand deviates significantly from actual demand, the certainty-equivalent model makes decisions that are suboptimal in hindsight. For instance, over-forecasting can lead to excessive replenishment and high inventory holding costs, while under-forecasting can trigger costly, last-minute direct picking to avoid stockouts. This analysis underscores that while the ADP framework is robust, its performance is intrinsically linked to the quality of its forecast inputs.

5.2. Managerial Insights

The results from this study provide actionable insights for warehouse managers looking to optimize operations in dynamic environments. Key takeaways include the following:

- Buffer areas should be viewed as a flexible resource for managing demand fluctuations, but they should be used thoughtfully. Overusing them can cause inefficiencies, while underusing them can lead to congestion in forward areas. Managers should focus on clearing out excess inventory in buffer areas first, as shown in the model, to keep operations smooth and efficient.
- The variability in Simulation PD outcomes highlights the critical need for robust and accurate demand forecasting systems. Investments in machine learning or advanced analytics can close the gap between predictive and perfect knowledge scenarios.
- The significant cost savings and efficiency gains observed for larger part groups indicate that the proposed model is particularly beneficial for high-volume warehouses.
- Managers overseeing complex operations should prioritize strategic replenishment planning, which reduces direct picking and minimizes labor costs.
- While Simulation KD demonstrates the lowest costs, achieving perfect demand knowledge is often unrealistic. Managers should aim to combine predictive tools with adaptive strategies to mitigate forecasting errors and maintain flexibility.

These insights serve as a roadmap for implementing advanced replenishment strategies that align with real-world constraints and opportunities.

Furthermore, managers should be aware of the model's sensitivity to cost parameters. The optimal policies for replenishment and buffer use will shift based on the relative costs of different operations. For example, if the cost of direct picking were to decrease significantly, the model would naturally favor more direct fulfillment from the reserve area and rely less on maintaining high stock levels in the forward and buffer areas. Conversely, an increase in inventory holding costs would incentivize the model to maintain leaner inventories and use the buffer more sparingly. A sensitivity analysis of these cost parameters could provide valuable insights for tailoring the strategy to a specific operational environment.

5.3. Limitations and Future Research

Our study, while providing a robust framework, has certain limitations that open avenues for future research. First, the model's performance is highly dependent on the accuracy of demand forecasts. The certainty-equivalent approach does not explicitly model the stochastic nature of demand, which can lead to suboptimal decisions as discussed. Future work could explore more advanced stochastic optimization techniques or two-stage programming to better account for uncertainty.

Second, practical implementation requires a sophisticated IT infrastructure. This includes systems for real-time inventory tracking, automated data pipelines for the demand forecasting model, and sufficient computational power to solve the MILP on a daily basis.

For future research, several promising directions can extend the contributions of this study. First, to provide a more comprehensive benchmark, our ADP model could be evaluated against established warehouse heuristics and purely deterministic models to better quantify the value added by the adaptive framework. Furthermore, the model could be enhanced by explicitly incorporating sustainability objectives. Inspired by the work of Stanković et al. [21] on optimizing dock door allocation to save energy, our objective function could be augmented to include energy consumption costs associated with internal vehicle movements, creating a trade-off between traditional operational costs and environmental impact. Building on this, the advanced stochastic optimization methods discussed by Zhong et al. [22] could be explored to manage the dual uncertainties of demand and energy pricing, moving beyond our certainty-equivalent approach to a more distributionally robust framework. Finally, the integration of our framework with warehouse automation presents a significant opportunity. The current ADP model could be adapted to generate optimal task sequences for an Automated Storage and Retrieval System (AS/RS), where the model's decisions would directly command robotic agents. This extension would involve reformulating cost functions to reflect robotic operational parameters and could leverage specialized optimization techniques, such as the cross-entropy method proposed by Foumani et al. [23], to solve the resulting complex scheduling problems.

6. Conclusions

This study comprehensively examined the internal replenishment process in an automotive spare parts warehouse, a critical component of the supply chain that directly impacts operational efficiency, costs, and customer satisfaction. One of the most significant challenges faced with this problem is that the system expands into an infinite state space over time due to its nonstationary distribution and multi-product structure. This creates computational difficulties, especially the curse of dimensionality, making it nearly impossible to solve with traditional methods like DP. The DP approach for this problem is highly complex and computationally costly, as it must consider each situation and decision.

By combining advanced methodologies such as ADP with machine learning techniques for demand forecasting, the research has developed a robust model that addresses the complexities of inventory management while improving decision-making under uncertainty. The findings indicate that maintaining optimal stock levels in the forward area significantly reduces the inefficiencies of direct picking from the reserve area, leading to a more effective replenishment process. The ADP application, designed to operate as a daily system, prioritizes parts that require more urgent replenishment and adapts to limited staff capacity, ensuring timely and efficient replenishment according to fluctuating demand patterns. Being a system that can be operated daily allows the warehouse to continuously adjust replenishment and buffer area management according to real-time demands. In this way, demand fluctuations are responded to immediately, both stock levels are kept balanced, and labor and costs are reduced. Daily operations enable warehouse processes to

be more flexible and dynamic, thus creating a structure that is more resistant to sudden changes and increasing customer satisfaction.

The study also emphasizes the crucial role of buffer areas in managing excess stock and preventing congestion. In an automotive warehouse environment, where parts vary significantly in size, demand frequency, and urgency, buffers are particularly vital for managing demand changes, facilitating efficient picking, and reducing congestion in smaller part groups. As part groups increase, the optimization model shifts focus to precise replenishment timing, reducing buffer dependence while keeping the forward area adequately stocked. Without buffer area utilization, handling costs would rise dramatically, especially in the dynamic environment of an automotive spare parts warehouse, where rapid access to parts is critical. The findings clearly demonstrate that strategic replenishment and buffer integration are essential for reducing labor costs, optimizing space usage, and ensuring smooth operations in a demand-sensitive warehouse.

The research suggests strategies that enable more efficient use of space and resources by prioritizing critical areas during order picking, thus supporting smoother operations even with limited staff.

In conclusion, this study's successful application of integrated ADP and machine learning models highlights the potential to improve warehouse operations and provides a basis for further innovation in inventory management applications. As businesses continue to meet the challenges of the evolving supply chain environment, the methodologies developed here offer strategic advantages, increased operational efficiency, and greater resilience to market fluctuations. This research contributes to a deeper understanding of internal replenishment processes, offering a compelling example of how advanced analytical tools can effectively manage complex logistics operations.

Author Contributions: İ.K. (Corresponding author): conceptualization, formal analysis, investigation, methodology, validation, visualization, writing—original draft, writing—review & editing. M.H.: supervision, methodology, writing—review & editing. A.D.Y.: supervision, methodology, writing—review & editing. G.K.: supervision, methodology, writing—review & editing. V.Ş.E.: supervision, methodology, writing—review & editing. Ş.Y.: conceptualization, formal analysis, investigation. All authors have read and agreed to the published version of the manuscript.

Funding: This study is funded by The Scientific Technological Research Council of Turkey (TUBITAK) under the 2244 Industrial PhD Fellowship Program, grant number 119C085.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting the findings of this study are proprietary and were provided under a research agreement. Due to confidentiality agreements and competitive considerations, the raw data cannot be made openly available. However, simulated datasets and the core model code can be made available to academic researchers upon reasonable request.

Conflicts of Interest: Author İrem Kalafat and Şenda Yıldırım are employed by Doğuş Technology as salaried staff members. However, this research was conducted independently and was not funded by the company. The authors declare no other conflicts of interest. Authors Mustafa Hekimoğlu, Ahmet Deniz Yücekaya, Gökhan Kirkil and Volkan Ş. Ediger declare no conflicts of interest.

Abbreviations

This are the abbreviations that are used in this manuscript.

Parameters:

N	Set of spare parts in the warehouse;
T	The planning horizon length in days;
D_{it}	Demand for product i at time t ;
C_i	Average batch size for the replenishment of product i ;
F_i	Maximum inventory capacity of the forward area for product i ;
τ	The number for daily replenishment capacity;
M_i	The maximum amount of batch that can be needed for replenishment of product i .

Costs:

c^d	Cost of direct picking from the reserve area;
c^w	Cost of keeping inventory in the buffer area;
c^{wx}	Cost of pulling inventory from the forward area to the buffer area;
c^r	Cost of internal replenishment.

Decision Variables:

XBR_{it}	Inventory level of product i in the forward area at time t before replenishment;
XAR_{it}	Inventory level of product i in the forward area at time t after replenishment;
X_{it}^W	Amount of product i that is pulled into buffer area for replenishment at time t ;
Q_{it}	Replenishment decision for product i at time t ;
QN_{it}	Number of cages needed for replenishing product i at time t from the reserve area;
WBR_{it}	Inventory level of product i in the buffer area at time t before replenishment;
WAR_{it}	Inventory level of product i in the buffer at time t after replenishment;
D_{it}^W	Amount of demand covered from the buffer area for product i at time t ;
D_{it}^S	Amount of demand covered from the forward area for product i at time t ;
D_{it}^P	Amount of demand covered from the reserve area (direct picking) for product i at time t ;
C_{it}^D	Number of batches needed for direct picking for product i at time t .

Appendix A

A. Dynamic Programming Formulation

1. State Variables:

$S_t = (XBR_{it}, XAR_{it}, WBR_{it}, WAR_{it})$: State at time t representing the inventory levels before and after replenishment in both the forward and buffer areas for each product i .

2. Decision Variables:

$a_t = (Q_{it}, QN_{it}, X_{it}^W, D_{it}^S, D_{it}^W, D_{it}^P, C_{it}^D)$: Action at time t representing the decisions for each product i .

3. Value Function:

$V_t(S_t)$: The minimum cost-to-go from the state (S_t) at time t to the end of the planning horizon.

4. Bellman Equation:

$$cost_t = \min_{a_t} \left\{ \sum_{i=1}^N \left(Q_{it}c^r + C_{it}^Dc^d + (WAR_{it} - DW_{it} + WBR_{it})c^w + X_{it}^Wc^{wx} \right) \right\}$$

$$V_t(S_t) = \min_{a_t} \{ cost_t + E[V_{t+1}(S_{t+1})] \}$$

$$V_t(S_t) = \min_{a_t} \left\{ \sum_{i=1}^N \left(Q_{it} c^r + C_{it}^D c^d + (WAR_{it} - DW_{it} + WBR_{it}) c^w + X_{it}^W c^{wx} \right) + E[V_{t+1}(XBR_{it+1}, XAR_{it+1}, WBR_{it+1}, WAR_{it+1})] \right\}$$

5. Constraints

State Transition Equations:

$$XAR_{it} = XBR_{it} + QN_{it}C_i - X_{it}^W \quad i \in N, 1 \leq t \leq T$$

$$XBR_{it+1} = XAR_{it} - D_{it}^S \quad i \in N, 1 \leq t \leq T$$

$$WAR_{it} = X_{it}^W + WBR_{it} \quad i \in N, 1 \leq t \leq T$$

$$WBR_{it+1} = WAR_{it} - D_{it}^W + C_{it}^D C_i - D_{it}^P \quad i \in N, 1 \leq t \leq T$$

Capacity Constraints:

$$X_{it}^W \leq Q_{it}C_i \quad i \in N, 1 \leq t \leq T$$

$$QN_{it} \leq Q_{it}M_i \quad i \in N, 1 \leq t \leq T$$

$$QN_{it}C_i + XBR_{it} - X_{it}^W \leq F_i \quad i \in N, 1 \leq t \leq T$$

$$QN_{it} \leq C_i F_i \quad i \in N, 1 \leq t \leq T$$

$$XAR_{it} \leq F_i \quad i \in N, 1 \leq t \leq T$$

Demand Fulfillment:

$$D_{it}^W + D_{it}^S + D_{it}^P = D_{it} \quad i \in N, 1 \leq t \leq T$$

$$X_{it}^W \leq XBR_{it} \quad i \in N, 1 \leq t \leq T$$

$$D_{it}^P \leq C_{it}^D C_i \quad i \in N, 1 \leq t \leq T$$

$$D_{it}^W \leq WAR_{it} \quad i \in N, 1 \leq t \leq T$$

Total Replenishment Capacity:

$$\sum_{i=1}^N Q_{it} \leq \tau \quad i \in N, 1 \leq t \leq T$$

Non-negativity and binary:

$$XBR_{it}, XAR_{it}, WBR_{it}, WAR_{it}, QN_{it}, X_{it}^W, D_{it}^S, D_{it}^W, D_{it}^P, C_{it}^D \geq 0 \quad i \in N, 1 \leq t \leq T$$

$$Q_{it} \in \{0, 1\}$$

B. Dynamic Programming Pseudo Code

This algorithm outlines the standard, computationally intractable backward induction method for solving the full DP formulation of the problem, based on foundational methods described in the literature [24].

Step 0: Initialization

0.1 Load historical demand data from the database

0.2 Define state variables, decision variables, and planning horizon T

0.3 Initialize value function $V_t(S_t) = 0$ for final stage T

Step 1: Backward Induction

For $t = T - 1, T - 2, \dots, 1$:For each state S_t : $V_t(S_t) = \infty$ For each action $a_t = (Q_{it}, QN_{it}, X_{it}^W, D_{it}^S, D_{it}^W, D_{it}^P, C_{it}^D)$:

Compute immediate cost:

$$cost_t = \min_{a_t} \left\{ \sum_{i=1}^N \left(Q_{it}c^r + C_{it}^D c^d + (WAR_{it} - DW_{it} + WBR_{it})c^w + X_{it}^W c^{wx} \right) \right\}$$

Compute the next state S_{t+1} according to the given constraints:

$$XAR_{it} = XBR_{it} + QN_{it}C_i - X_{it}^W$$

$$XBR_{it+1} = XAR_{it} - D_{it}^S$$

$$WAR_{it} = X_{it}^W + WBR_{it}$$

$$WBR_{it+1} = WAR_{it} - D_{it}^W + C_{it}^D C_i - D_{it}^P$$

Evaluate the next value:

$$\text{Next Value} = V_{t+1}(S_{t+1})$$

Update value function:

$$V_t(S_t) = \min(V_t(S_t), cost_t + \text{Next Value})$$

Step 2: Forward Simulation

Initialize starting state S_0 For $t = 1, 2, \dots, T$:Determine optimal action a_t^* based on $V_t(S_t)$ Apply optimal action a_t^* :Execute optimal replenishment decisions $Q_{it}^*, QN_{it}^*, X_{it}^{W*}, D_{it}^{S*}, D_{it}^{W*}, D_{it}^{P*}, C_{it}^{D*}$ Update state S_{t+1}

Record actual demand and update historical data

Update initial states for the next iteration

End

References

- De Koster, M.B.M. Warehouse Assessment in a Single Tour. In *Warehousing in the Global Supply Chain*; Springer: London, UK, 2012; pp. 457–473.
- Hackman, S.T.; Rosenblatt, M.J.; Olin, J.M. Allocating Items to an Automated Storage and Retrieval System. *IIE Trans.* **1990**, *22*, 7–14. [\[CrossRef\]](#)
- van den Berg, J.P.; Sharp, G.P.; Gademann, A.J.R.M.; Pochet, Y. Forward-Reserve Allocation in a Warehouse with Unit-Load Replenishments. *Eur. J. Oper. Res.* **1998**, *111*, 98–113. [\[CrossRef\]](#)
- Heragu, S.S.; Du, L.; Mantel, R.J.; Schuur, P.C. Mathematical Model for Warehouse Design and Product Allocation. *Int. J. Prod. Res.* **2005**, *43*, 327–338. [\[CrossRef\]](#)
- Bartholdi, J.J.; Hackman, S.T. Allocating Space in a Forward Pick Area of a Distribution Center for Small Parts. *IIE Trans.* **2008**, *40*, 1046–1053. [\[CrossRef\]](#)
- Gagliardi, J.-P.; Ruiz, A.; Renaud, J. Space Allocation and Stock Replenishment Synchronization in a Distribution Center. *Int. J. Prod. Econ.* **2008**, *115*, 19–27. [\[CrossRef\]](#)
- Gu, J.; Goetschalckx, M.; McGinnis, L.F. Solving the Forward-Reserve Allocation Problem in Warehouse Order Picking Systems. *J. Oper. Res. Soc.* **2010**, *61*, 1013–1021. [\[CrossRef\]](#)
- Accorsi, R.; Manzini, R.; Bortolini, M. A Hierarchical Procedure for Storage Allocation and Assignment within an Order-Picking System. *A Case Study. Int. J. Logist. Res. Appl.* **2012**, *15*, 351–364. [\[CrossRef\]](#)
- Walter, R.; Boysen, N.; Scholl, A. The Discrete Forward-Reserve Problem—Allocating Space, Selecting Products, and Area Sizing in Forward Order Picking. *Eur. J. Oper. Res.* **2013**, *229*, 585–594. [\[CrossRef\]](#)

10. de Vries, H.; Carrasco-Gallego, R.; Farenhorst-Yuan, T.; Dekker, R. Prioritizing Replenishments of the Piece Picking Area. *Eur. J. Oper. Res.* **2014**, *236*, 126–134. [[CrossRef](#)]
11. Bahrami, B.; Aghezzaf, E.-H.; Limère, V. Enhancing the Order Picking Process through a New Storage Assignment Strategy in Forward-Reserve Area. *Int. J. Prod. Res.* **2019**, *57*, 6593–6614. [[CrossRef](#)]
12. Wu, W.; de Koster, R.B.M.; Yu, Y. Forward-Reserve Storage Strategies with Order Picking: When Do They Pay Off? *IIE Trans.* **2020**, *52*, 961–976. [[CrossRef](#)]
13. Jiang, M.; Huang, G.Q. Intralogistics Synchronization in Robotic Forward-Reserve Warehouses for e-Commerce Last-Mile Delivery. *Transp. Res. Part E Logist. Transp. Rev.* **2022**, *158*, 102619. [[CrossRef](#)]
14. Çelik, M.; Archetti, C.; Süral, H. Inventory Routing in a Warehouse: The Storage Replenishment Routing Problem. *Eur. J. Oper. Res.* **2022**, *301*, 1117–1132. [[CrossRef](#)]
15. Bertsimas, D.; Demir, R. An Approximate Dynamic Programming Approach to Multidimensional Knapsack Problems. *Manag. Sci.* **2002**, *48*, 550–565. [[CrossRef](#)]
16. Caro, F.; Gallien, J. Clearance Pricing Optimization for a Fast-Fashion Retailer. *Oper. Res.* **2012**, *60*, 1404–1422. [[CrossRef](#)]
17. Hopp, W.J.; Spearman, M.L. *Factory Physics*; Waveland Press: Long Grove, IL, USA, 2011.
18. Simchi-Levi, D.; Kaminsky, P.; Simchi-Levi, E. *Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies*; McGraw-Hill: New York, NY, USA, 2008.
19. Azadeh, K.; de Koster, R.; Roy, D. Robotized and Automated Warehouse Systems: A Review and Recent Developments. *Transp. Sci.* **2019**, *53*, 917–945. [[CrossRef](#)]
20. Besbes, O.; Gur, Y.; Zeevi, A. Non-Stationary Stochastic Optimization. *Oper. Res.* **2015**, *63*, 1227–1244. [[CrossRef](#)]
21. Stanković, R.; Rogić, K.; Šafran, M. Saving Energy by Optimizing Warehouse Dock Door Allocation. *Energies* **2022**, *15*, 5862. [[CrossRef](#)]
22. Zhong, J.; Xie, S.; Li, Y.; Liu, J.; Zhou, B. Synergistic Operation Framework for the Energy Hub Merging Stochastic Distributionally Robust Chance-Constrained Optimization and Stackelberg Game. *IEEE Trans. Smart Grid* **2025**, *16*, 1037–1050. [[CrossRef](#)]
23. Foumani, M.; Moeini, A.; Haythorpe, M.; Smith-Miles, K. A cross-entropy method for optimising robotic automated storage and retrieval systems. *Int. J. Prod. Res.* **2018**, *56*, 6450–6472. [[CrossRef](#)]
24. Powell, W.B. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd ed.; John Wiley & Sons: Hoboken, NJ, USA, 2011.

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.