A Markov Decision Process-based ICU Inventory Management System for a South African Hospital

A. Steenkamp[1*], M.M. van Zyl-Cillié[1] & C. du Plessis[1]

ARTICLE INFO

Article details
Presented at the 2025 Production and Operations Management Society Africa International Conference, held from 26 to 28 November 2025 in Pretoria, South Africa

Available online 8 Dec 2025

Contact details
* Corresponding author
anesteenkamp29@gmail.com

Author affiliations
1 School of Industrial Engineering, North-West University, South Africa

ORCID® identifiers
A. Steenkamp
https://orcid.org/0009-0002-4426-2574

M.M. van Zyl-Cillié
https://orcid.org/0000-0003-3320-706X

C. du Plessis
https://orcid.org/0000-0002-6968-0433

DOI
http://dx.doi.org//10.7166/36-3-3338

ABSTRACT

The South African healthcare system is divided into public and private healthcare. Although the private healthcare sector is funded by individuals or health insurance, private hospitals experience operational problems associated with the optimal allocation and usage of resources. The intensive care unit (ICU) of a private hospital in South Africa frequently experiences shortages of critical supplies owing to limited inventory visibility and inflexible replenishment policies. Demand variability and static inventory ordering policies are found to be the principal drivers of stock-outs in this specific ICU. This research aims to develop a probabilistic inventory management system for an ICU to enhance supply availability and to minimise risks to patient care. This framework is based on a Markov decision process (MDP), deployed as the methodological foundation for modelling the inventory state transition probabilities derived from historical usage data. A Gradio-based interface facilitates real-time interaction with the model outputs, ensuring accessibility for clinical and administrative staff. The proposed system improves visibility and supports proactive decision-making in response to stochastic demand. The findings demonstrate that an MDP, integrated with accessible data and interface tools, could reduce shortages and expired stock, thereby aligning the ICU inventory management policies more closely with clinical requirements.

OPSOMMING

Die Suid-Afrikaanse gesondheidsorgstelsel word verdeel in openbare en private gesondheidsorg. Alhoewel die private gesondheidsorgsektor deur individue of gesondheidsversekering befonds word, ervaar private hospitale operasionele probleme wat verband hou met die optimale toewysing en gebruik van hulpbronne. Die intensiewe sorgeenheid (ICU) van 'n private hospitaal in Suid-Afrika ervaar gereeld tekorte aan kritieke voorrade as gevolg van beperkte voorraadsigbaarheid en onbuigsame aanvullingsbeleide. Vraagvariasie en statiese voorraadbestellingsbeleide word as die hoofdryfvere van voorraadtekorte in hierdie spesifieke ICU beskou. Hierdie navorsing het ten doel om 'n probabilistiese voorraadbestuurstelsel vir 'n ICU te ontwikkel om voorraadbeskikbaarheid te verbeter en risiko's vir pasiëntsorg te minimaliseer. Hierdie raamwerk is gebaseer op 'n Markov-besluitnemingsproses (MDP), wat as die metodologiese fondament gebruik word vir die modellering van die voorraadtoestand-oorgangswaarskynlikhede afgelei van historiese gebruiksdata. 'n Gradio-gebaseerde koppelvlak fasiliteer intydse interaksie met die modeluitsette, wat toeganklikheid vir kliniese en administratiewe personeel verseker. Die voorgestelde stelsel verbeter sigbaarheid en ondersteun proaktiewe besluitneming in reaksie op stogastiese vraag. Die bevindinge toon dat 'n MDP, geïntegreer met toeganklike data en koppelvlak-instrumente, tekorte en vervalde voorraad kan verminder, waardeur die ICU-voorraadbestuursbeleide nouer met kliniese vereistes in lyn gebring kan word.

1. INTRODUCTION

South Africa's healthcare system is divided into the public and private sectors. All citizens in South Africa can access public healthcare services, which are funded and operated by the state [1]. Private healthcare, on the other hand, is funded by individuals or health insurance. Hospitals in the private healthcare sector are autonomous, profit-driven organisations [2] and generally have better access to resources. Similar to hospitals around the world, these private hospitals experience operational problems associated with the optimal allocation and usage of resources [3]. In addition, inventory as a key resource is often difficult to manage in a dynamic setting such as a hospital [4].

"Inventory management" in a hospital setting describes methodical planning, monitoring, and regulation of inventories to satisfy demand while lowering expenses. In more traditional organisations, inventory management [5] guarantees product availability without requiring excessive capital to be tied up in storage. In hospitals, medical supplies have specified uses, expiration dates, and storage requirements, such as refrigerated vaccines or sterile surgical kits [6]. Applying efficient inventory management principles, therefore, underpins operational efficiency in hospitals.

Building on the concept of inventory management, its role and importance in hospitals is critical for both clinical and financial results. Medical supplies and equipment, such as bandages, syringes, life-saving pharmaceuticals, and surgical instruments, are necessary for patient care [6]. According to Lee *et al.* [6], efficient inventory management guarantees that medical personnel have instant access to these supplies, facilitating prompt and efficient treatment. Effective inventory management also lowers financial risks by reducing stockouts, which can disrupt patient care, and overstocking, which raises expenses because of storage or expired medications [5]. In South African hospitals, effective inventory management is crucial for operating efficiency and compliance with regulatory criteria, such as the National Health Act (No. 61 of 2003), which demands appropriate medical supply availability [7]. Moreover, inventory management influences numerous hospital systems, including procurement, storage, distribution, and billing, which in turn depend on accurate inventory data for precise patient invoicing.

In this research, we investigate the inventory management of the ICU in a 197-bed private hospital in South Africa. The hospital provides both general and specialised care through the medical wards, the ICU, paediatrics, and the neonatal intensive care unit (NICU), in addition to emergency care, renal care, and operating theatres. The pharmacy serves as the main consumable and medication supplier for all wards, including the ICU. The pharmacy operates under a first-in, first-out (FIFO) inventory policy, and uses a centralised AS-400 inventory management system alongside paper records for daily stock checks and weekly reorder quantities.

However, the hospital's decentralised billing and inventory management system requires each ward to manage its billing and significant portions of stock control independently, resulting in irregular inventory tracking. This creates problems for the ICU, which requires an uninterrupted supply of high-value stock such as ventilatory medications and IV fluids to enable critical interventions. The current reordering system depends on monthly averages of product usage to determine order quantities, which results in frequent stock unavailability because of unpredictable demand fluctuations.

The ICU has emerged as a pivotal area of focus owing to its elevated stock turnover, frequent patient admissions, and dependence on a continuous supply of high-value inventory, including ventilatory medications and intravenous fluids, which are essential for life-saving interventions. Variability in demand compounds these problems, as the current replenishment approach establishes stock levels based on the previous month's average usage, and fails to accommodate seasonal fluctuations or the complexity of cases. The objective of this research, therefore, is to formulate and design an inventory management system that enhances inventory availability in the ICU.

In Section 2, we explore inventory management principles and concepts that could assist in developing the solution. We provide an overview of the phased-approach methodology that is followed to design the solution in Section 3. In Section 4, we provide a brief demonstration of the system's functionality before concluding with the system's validation in Section 5.

2. LITERATURE REVIEW

This literature review explores potential solutions to the difficulties of inventory management in ICUs, with a particular emphasis on methodologies that apply to settings characterised by high demand variability and constrained inventory visibility. The review integrates research on demand forecasting models, probabilistic techniques, classification-based inventory control, integrated management systems, and real-time visibility solutions to lay the groundwork for the development of an effective ICU inventory management system.

2.1. Demand-forecasting models

Accurate demand forecasting is a critical component of effective inventory management in critical care settings, where consumption variability and urgent clinical needs make complex forecasting difficult. Traditional statistical methods such as the autoregressive integrated moving average and exponential smoothing perform adequately in environments with relatively stable demand patterns, but they are less suitable in the non-stationary demand contexts that typify ICU operations [8],[9]. A significant limitation of these conventional models is their inability to account for censored demand during stockouts, which results in systematic underestimation in subsequent periods [10].

To address these limitations, Barrow and Kourentzes [11] advocate forecast combination approaches that integrate multiple models to reduce error variations and improve inventory efficiency. This methodology proves particularly beneficial in unstable, data-limited contexts by mitigating the risks associated with model-specific biases. Tiacci and Saetta [12] emphasise the need to align forecasting methodologies with existing stock control procedures, demonstrating that even highly accurate forecasts can result in poor inventory performance without appropriate replenishment policies.

Segmentation-based forecasting approaches offer another suitable avenue for improved inventory management in an ICU setting. Kalchschmidt *et al.* [13] propose demand segmentation by item category, patient type, or clinical use to increase accuracy, recognising that different products in ICU environments have varying consumption drivers and predictability patterns. Advanced modelling techniques, including neural networks and entropy-based models, have been developed to manage complex inter-product interactions and high variability, with Lei *et al.* [14] proposing entropy-driven forecasting models that consider product associations to be particularly relevant in ICUs, where certain items are used interdependently. Karamshetty *et al.* [15] emphasise that operational viability is critical to implementing successful forecasting, noting that inadequate training, lack of digital infrastructure, or staffing shortages frequently cause sophisticated tools to malfunction in private hospital settings.

2.2. Markov chains and probabilistic models

Healthcare inventory systems benefit substantially from Markovian models for managing variability and uncertainty. Markov chains, characterised by their memoryless property when future outcomes depend solely on current circumstances rather than on historical events, prove well-suited to dynamic healthcare environments [16]. Vila-Parrish *et al.* [17] demonstrate the application of two-stage Markov chain models in managing consumable medical supplies under erratic demand, showing how these models adapt to regular fluctuations in consumption patterns and improve inventory control for time-sensitive medical supplies.

Building on this foundation, He and Jiang [18] formulate belief-based Markov models, which are a variation of Markov models in which the true state is not observable. Belief-based Markov models are adept at functioning with ambiguous or incomplete inventory data, rendering them particularly advantageous in contexts characterised by intermittent data collection or diminished inventory visibility. MDPs, which augment Markov chains with decision-making capabilities, can be used to optimise reorder decisions continuously in real-time, while simultaneously balancing cost-effectiveness with supply availability [19]. In the context of humanitarian healthcare logistics, Ferreira *et al.* [20] applied MDPs to sustain service levels amid supply chain disruptions and significant variability; nevertheless, the applicability of MDPs in resource-constrained healthcare systems could be limited by their computational demands and the requisite level of technical expertise.

2.3. Classification-based inventory control

Priority inventory segmentation using classification systems provides a proven method for improving stock reliability and efficient resource allocation [21]. The combination of ABC analysis, which groups products by value or annual consumption, with XYZ analysis, which categorises items by demand variability, enables differentiated inventory policies in which high-value variable items receive stricter controls while stable low-value items require less oversight [22]. Vila-Parrish and Ivy [23] advocate the vital, essential, non-essential classification system, which, when combined with ABC analysis, helps medical professionals to prioritise life-saving or essential medications for more stringent monitoring and restocking.

Al-Qatawneh and Hafeez [24] propose multi-criteria inventory classification models that incorporate supply risk, lead time, and unit cost to enhance decision-making capabilities. This approach ensures the continuous availability of critical supplies while minimising the excess inventory of low-priority items. Classification-based inventory control operates on the principle that not every item requires the same level of attention, allowing managers to concentrate resources where they have the greatest impact on costs and service levels [25]. However, classification techniques alone cannot resolve problems with demand forecasting and real-time inventory visibility; they require integration into comprehensive inventory control frameworks that incorporate data integration, supply chain coordination, and demand planning.

2.4. Integrated inventory management systems

Integrated inventory management systems involve the coordinated use of technology, procedures, and data to oversee supply chain operations in multiple departments in healthcare institutions. Huijsman *et al.* [26] suggest that integration in healthcare supply chains requires coordinating information exchanges, material flows, and decision-making processes to improve resource efficiency, responsiveness, and patient outcomes. These systems typically connect inventory data with the clinical, administrative, and procurement processes through enterprise resource planning (ERP), pharmacy management software, and mobile technology to provide visibility and control [27].

While ERP systems have the potential to solve inventory management problems in ICU settings, Mabizela *et al.* [28] demonstrate in a South African case study how inadequate inventory visibility and unprioritised stock control caused frequent stockouts in critical care units, despite their having ERP systems in place. Mabizela *et al.* [28] suggest instead that request procedures and reordering stock be automated according to clinical risk. This strategy resulted in significantly improved item availability and reduced manual intervention. Saha and Ray [29] suggest augmenting ERP systems with dashboards and mobile-based tracking tools to increase real-time visibility and to reduce delays caused by data fragmentation. Shenoy [30] supports this perspective, arguing that coordinating clinical workflows with inventory data ensures prompt replacements, reduced duplication, and proactive planning. Vila-Parrish and Ivy [23] warn that inventory problems frequently force clinical professionals to perform supply chain tasks such as locating replacements manually or modifying requisition orders.

2.5. Real-time inventory visibility solutions

Real-time inventory visibility solutions address the fundamental problem of tracking and displaying current stock levels using various technological approaches. Barcode systems provide cost-effective, user-friendly methods for tracking medical supplies at the point of use. However, their effectiveness depends on consistent manual scanning, which may prove unreliable in high-stress ICU environments [31],[32]. Radio frequency identification (RFID) systems offer automatic identification and tracking capabilities without requiring manual handling, which would be ideal for real-time tracking of valuable or critical inventory through mass, non-line-of-sight scanning. However, implementation costs and infrastructure requirements may limit their adoption [32].

Internet of Things (IoT) approaches enable comprehensive tracking of expiration dates, product usage, and storage conditions through interconnected sensor networks [33]. Tadayonrad and Ndiaye [34] found that IoT-based systems significantly improve inventory key performance indicators (KPIs), including stock accuracy and lead time, by facilitating real-time warning and predictive replenishment. Digital dashboards serve as crucial tools for converting raw inventory data into actionable insights, particularly when paired with forecasting tools to enable proactive stock management and reduce stockout risks [22].

Each technological solution contributes to bridging the gap between visibility and inventory availability, but successful implementation requires its adoption by personnel, infrastructure capability, and hospital digital maturity-factors that vary significantly among South African healthcare facilities. The literature reveals that effective inventory management in ICU environments requires integrated approaches that combine accurate demand forecasting, probabilistic modelling for uncertainty management, strategic classification for resource prioritisation, and a real-time visibility tool, with successful implementation relying heavily on contextual factors that are specific to each healthcare institution.

3. METHODOLOGY

To develop an inventory management solution that would ensure stock availability in the ICU of the hospital where this research was conducted, we followed a phased-approach methodology. In the first phase (analysis), a thorough investigation into the causes leading to inventory unavailability was conducted. Through this analysis, several of the problem causes were identified, and this led to the second phase (system requirements), in which system requirements were identified to ensure that the final solution addressed the current problem's causes. In the third phase (system design) of the solution development, we provided the details of the inventory management system solution's design.

3.1. Phase 1: Analysis

Our initial investigation showed that the inventory management difficulties specific to the ICU under investigation consisted of : (i) stock disappearance because of inadequate tracking processes; (ii) inter-ward borrowing that complicates inventory forecasting; (iii) the need for a diverse and dynamic supply portfolio for specialised ICU care; (iv) inadequate monitoring of expired medication, resulting in waste; (v) the management of a large and complex stock range that exacerbates control difficulties; (vi) the direct endangerment of patient outcomes from stockouts in critical cases; (vii) the procurement of highly specific stock to meet physicians' preferences; and (viii) supplier back orders that augment the risk of shortages.

The literature on hospital inventory management offers insightful information on related problems, indicating possible underlying causes of this problem.

Bhakoo and Choi [4], for instance, note that one of the primary causes of stock discrepancies – which are mismatches between recorded and real stock levels – is insufficient inventory visibility, which often arises from inadequate record-keeping systems. Similar to this, De Vries [15] observes that order processing delays and unreliable supplier performance are one of the main causes of pharmaceutical stock-outs, especially in clinical environments with high demand, such as an ICU. Kelle *et al*. [35] allude to the negative consequence of demand variability, and note that inventory shortages are negatively affected by erratic consumption patterns in critical care units. In addition to this, stock-out risks are increased by a lack of communication between clinical and pharmacy staff, leading to incorrect or delayed orders [36]. To determine the unique underlying drivers of stockouts in the ICU under investigation, a root cause analysis was conducted by combining a qualitative categorisation of potential causes with quantitative validation. Based on the literature and on field observations, six categories of potential root causes were identified (Table 1).

Table 1: Root cause category descriptions

Root cause category	Description
Record keeping	This category encompasses all aspects of inventory record management, such as inventory management software, paper-based records, billing systems, and human errors associated with maintaining accurate inventory data.
Demand variability	This pertains to the variety that emerges from administering distinct treatments to patients, each needing individualised care.
Inventory policies	These are the existing policies of the hospital that regulate all activities related to inventory management.
Communication	This covered intradepartmental communication, interdepartmental communication, and external communication with suppliers.
Pharmacy layout	This is the physical layout of the pharmacy surgical storage area.
Suppliers	This refers to the preferred suppliers that had a contract with this private hospital.

To examine how these factors interacted, an interrelationship diagram was developed, shown in Figure 1.

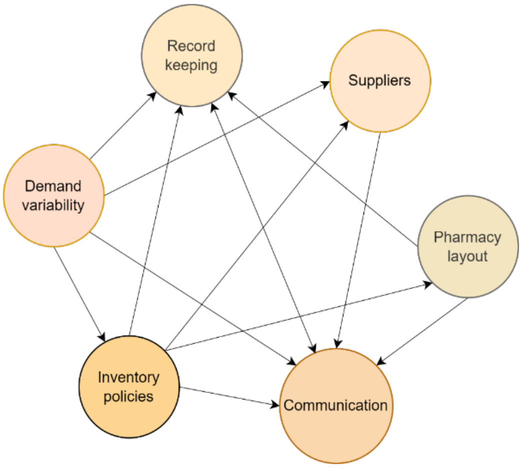

Figure 1: Interrelationship diagram

Figure 1 illustrates that variability in demand and rigid inventory policies were the principal root causes, exerting the most significant influence on other categories. These two factors engendered cascading effects, including delayed replenishment, inaccurate demand forecasting, and increased dependence on emergency workarounds, thus increasing the risk of stockouts.

To investigate further the inventory management problems of the ICU in question, a Pareto analysis was conducted to provide insights into the current inventory holding costs.

Figure 2 shows a classic Pareto distribution of the 125 ICU inventory items that were analysed. The ABC classification revealed that just 25 items (20%) accounted for R2.6 million (80%) of total inventory value, while 35 items (28%) contributed R488 000 (15%), and the remaining 65 items (52%) accounted for only R165 000 (5%) of the total value. This concentration validated the Pareto principle, and confirmed that the focused management of high-value items could yield disproportionate operational improvements.

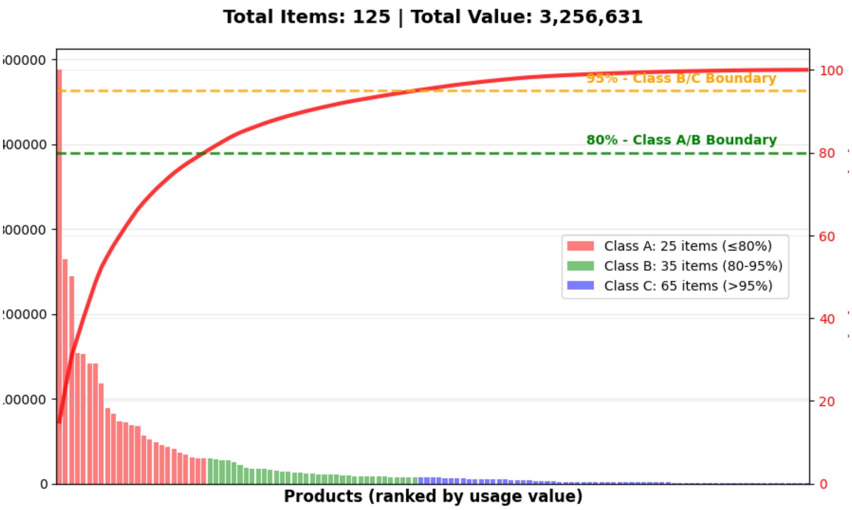

ABC Class	Item Count	% of Items	Total Value	% of Value
Class A	25	20.0%	2,603,952	80.0%
Class B	35	28.0%	487,987	15.0%
Class C	65	52.0%	164,692	5.1%

Figure 2: ABC Pareto distribution of ICU inventory

The pronounced decline following the initial 25 items is visually striking, accentuating the concentration of risk in a limited subset of products. The policy implication is clear: Class A items need rigorous monitoring and precise forecasting, whereas Class C items could be managed with simpler, less resource-intensive strategies.

Our analysis explained the inventory management system's sensitivity to variation. Given that 80% of the value was concentrated in merely 25 items, even minor demand fluctuations for these products could have disproportionate financial and operational repercussions. The prevailing static reorder policy did not account for this differentiation, and treated high-impact Class A items the same as low-value Class C items.

3.2. Phase 2: System requirements

Through a comprehensive examination of the current ICU inventory processes, supplemented by a review of the pertinent literature, numerous essential requirements were identified to guide the system design. These requirements were intended to ensure that the proposed solution effectively mitigated the operational inefficiencies that were detected, while remaining viable in the context of a private South African hospital.

Figure 3 delineates the overarching architecture of the proposed solution, which consists of three interrelated components: input requirements, the inventory management model, and the user interface. This figure shows the systematic flow of data deriving from ICU consumption patterns into the model and the ensuing production of outputs that enhance inventory visibility and inform decision-making.

Fundamentally, this framework guarantees that the system would be data-efficient, capable of adapting to fluctuations in demand, and accessible to both pharmacy and ICU stakeholders with minimal technical expertise. Consequently, it should achieve a balance between analytical rigour and operational practicality.

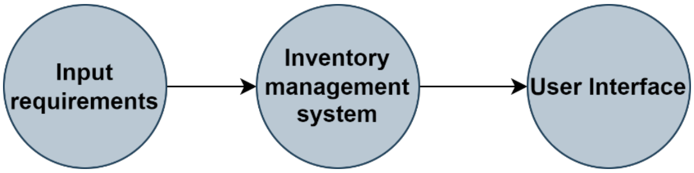

Figure 3: Conceptual design components

3.2.1. Solution requirements

The requirements guiding the conceptual design were derived directly from the root cause analysis and validated against findings from the literature review.

These can be summarised as follows:

- The solution must be able to adapt to demand variability.
- The solution should enhance inventory visibility.
- The solution should have a user-friendly interface.
- The solution must reduce stockouts.
- The solution must be able to work with minimal historical data.
- The solution should adhere to the applicable inventory policies.
- The proposed solution must be operationally feasible by adhering to the hospital's existing infrastructure.

These requirements formed the benchmark for evaluating alternative system designs and forecasting methodologies for the proposed inventory management system.

3.2.2. Justification for using a Markov decision process

In evaluating potential inventory management approaches, three broad categories of solution were shortlisted: **(i)** MDP; **(ii)** KPI-linked forecasting models; and **(iii)** conventional statistical methods such as moving averages or exponential smoothing. A structured decision matrix was applied to compare these approaches against the requirements.

The MDP model achieved the highest overall suitability score owing to its ability to:

- Operate effectively with sparse historical data while still providing probabilistic insights.
- Explicitly incorporate decision-making under uncertainty, enabling not only demand forecasting but also optimised replenishment actions.
- Adapt dynamically to variability by updating transition probabilities between demand states.
- Balance short-term service levels (avoiding stockouts) against long-term cost objectives – an advantage over memoryless Markov chains.

At this point, several solutions were disqualified for the following reasons:

- **Classification-based inventory control** was considered more appropriate as an auxiliary prioritisation instrument rather than as an independent forecasting technique.
- Owing to the expense, infrastructure, and training requirements**, IoT-based tracking (RFID, barcode systems)** and integrated inventory management systems were disqualified.
- **Advance hybrid models** (forecast combinations, segmentation-based forecasting, and error distribution-based techniques) were not evaluated further since they exceeded the data and expertise level available in this instance.

This comparative analysis highlights that an MDP presents the most equitable and balanced approach, combining analytical robustness with practical feasibility for addressing the ICU inventory problem. Unlike more rudimentary models, the MDP framework combines forecasting, optimisation, and policy evaluation in a unified system.

3.3. Phase 3: System design

As illustrated in Figure 3, the system design follows from the inventory management system's input requirements. The proposed ICU inventory management system integrates three components: structured data input, an MDP model, and a user interface.

Applying an MDP requires the explicit definition of states, actions, and rewards in order to model ICU inventory dynamics under uncertainty. Unlike simple Markov chains, which capture only probabilistic transitions between demand states, the MDP incorporates decision-making by assigning costs or rewards to different outcomes and identifying the optimal policy for replenishment.

To represent variability in ICU consumption, demand was categorised into three discrete states, derived from historical usage data and adjusted to reflect clinical conditions:

- **Low demand state:**

  Characterised by routine baseline consumption of medications and consumables, typically observed during periods of stable patient volumes.
- **Moderate demand state:**
  Represents fluctuations caused by variable patient inflow, specialised procedures, or rotational physician practices.
- **High demand state:**

  Reflects peaks in consumption driven by emergencies, seasonal outbreaks, or sudden increases in admissions. Transition probabilities between these states are estimated from observed ICU usage patterns, capturing the stochastic nature of demand variability.
- **Actions**
  For each state, the system evaluates possible replenishment decisions:
  - **Hold**: Maintain current stock without reordering.

- o **Reorder**: Trigger replenishment at the standard threshold.
- o **Emergency replenish**: Expedite restocking outside of normal cycles to address critical shortages.

- **Rewards**

  Each action–state combination is associated with a reward function that reflects operational objectives:
  - o **Stockout penalty**: High negative cost to represent risks to patient care.
  - o **Holding cost**: Moderate penalty for excess inventory, including expiration risks.
  - o **Service level reward**: Positive value for maintaining adequate availability of high-priority items.

The MDP framework uses these states, actions, and rewards to compute an optimal policy that prescribes the best action for each demand state, balancing the trade-off between minimising shortages and controlling costs. For instance, in a high demand state the optimal policy may favour emergency replenishment despite its higher cost, whereas in a low demand state, maintaining current stock may be preferred.

By embedding clinical realities in this state–action–reward structure, the MDP provides a practical decision-support tool that not only forecasts demand but also prescribes context-appropriate actions to improve ICU inventory resilience.

i. Data input

Inventory data was systematically gathered using Excel spreadsheets, aligning with the established reporting protocols of the hospital. Each entry contained product codes, demand statistics (mean and variance), storage limitations, lead times, and present stock quantities. This approach guaranteed complete compatibility with current workflows while supplying the structured inputs necessary for optimisation.

ii. MDP formulation

State space representation

The system models inventory control decisions using a three-dimensional state space:

$$s = (I_h, Q_o, \tau) \tag{1}$$

where:

I_h= inventory on hand

Q_o= outstanding order quantity

τ = days remaining until order arrival (zero if there is no outstanding order)

This model explicitly encapsulates lead time and outstanding orders, thereby addressing the limitations inherent in more rudimentary models that overlook order pipelines.

Action space

At each decision point, the available action A (order quantity) depends on the presence of outstanding orders:

$$A(s) = \begin{cases} \{0, q1, q2, \ldots, q_{max}\}, & if\ \tau = 0 \\ \{0\}, & if\ \tau > 0 \end{cases} \tag{2}$$

Transition Dynamics

State transitions incorporate three sequential processes:

- **Order arrival (if τ = 1):**

$$I'_{temp} = min(I_h + Q_o, I_{max}) \tag{3}$$

$$Q'_{temp} = 0,\ \tau'_{temp} = 0 \tag{4}$$

- **Demand fulfilment:**

$$I'_h = max\left(0, I'_{temp} - d\right) \tag{5}$$

- **New order placement (if $\boldsymbol{\tau'_{temp} = 0}$ $\boldsymbol{and\ a > 0}$):**

$$Q'_o = min(a, I_{max} - I'_h) \tag{6}$$

$$\tau' = \begin{cases} L, & Q'_o > 0 \\ 0, & Q'_o > 0 \end{cases} \tag{7}$$

where L is the lead time and Imax is the storage capacity.

- **Cost structure:**

The immediate cost function balances holding, shortage and ordering considerations:

$$c(s, a, d) = h \cdot I'_h + p \cdot max\left(0, d - I_{temp}\right) + f \cdot 1(a > 0) \tag{8}$$

where h = holding cost per unit, p = stockout penalty, f = fixed ordering cost, and 1(a > 0) is an indicator for order placement.

Value function and policy

The optimal replenishment policy solves the Bellman equation [37]:

$$V(s) = min_{a \in A(s)} \{\textstyle\sum_d P(d)[c(s, a, d) + \gamma V(s')]\} \tag{9}$$

where ϒ is the discount factor and P(d) is the demand distribution, modelled as either:

- Poisson when the variance ≈ mean, or
- Negative binomial when demand is over-dispersed.

Value iteration is used to compute policies, terminating when convergence falls below 10^{-4}.

Implementation features:

- **Demand modelling**: The system automatically selects the appropriate demand distribution (poisson vs negative binomial) based on variance-to-mean ratios.
- **State space optimisation**: Adaptive discretisation focuses computational effort on frequently visited states while coarsening extreme values, ensuring tractability.
- **Policy execution**: The model directly prescribes replenishment actions, balancing service level guarantees with cost efficiency.

iii. **User interface:**

A streamlined web-based interface was developed using Gradio, enabling hospital personnel to upload the Excel input file, execute the optimisation process, and obtain the results in a categorised format:

- **Monitor** – stock levels are sufficient.

- **Low** – prepare to order – approaching safety thresholds.
- **Critical** – order now – immediate replenishment required.

This translation of complex optimisation results into accessible decision categories ensures usability for pharmacy and ICU staff without needing them to have technical expertise.

4. DEMONSTRATION OF SYSTEM FUNCTIONALITY

The designed system is operationalised using a web-based interface developed with Gradio, a Python framework that automatically generates interactive user interfaces from function definitions. The interface enables file upload functionality and real-time results display through HTML rendering, eliminating the need for complex web development while maintaining professional presentation standards.

This is illustrated in Figure 4.

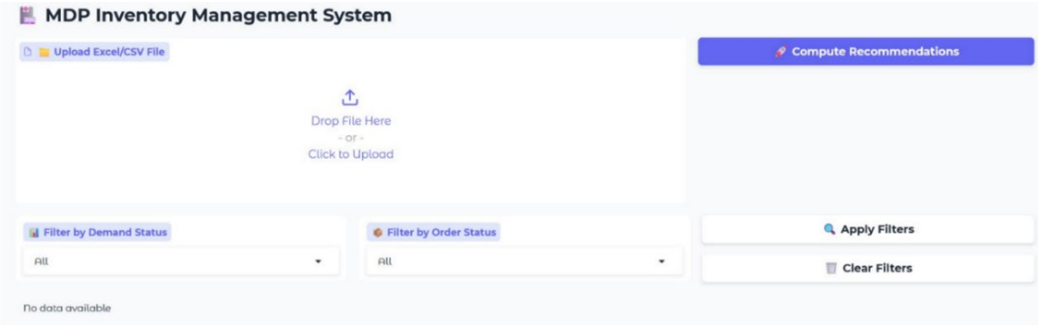

Figure 4: Initial interface of the inventory management system

To operate the system, pharmacy staff responsible for inventory control upload the most recent ICU inventory data into the system. The MDP then models state transitions under conditions of uncertain demand, generating recommendations for replenishment. Items anticipated to descend below established safety thresholds are flagged, with critical shortages initiating an "Order now" alert accompanied by an "Order now" message demonstrating this workflow.

Figure 5 displays the interface of the MDP-generated inventory recommendations in a tabular format with colour-coded alerts: green "Sufficient stock" indicators show items with adequate inventory, while red "Order now" flags identify critical items that require immediate replenishment based on current stock levels, demand patterns, and days of supply remaining.

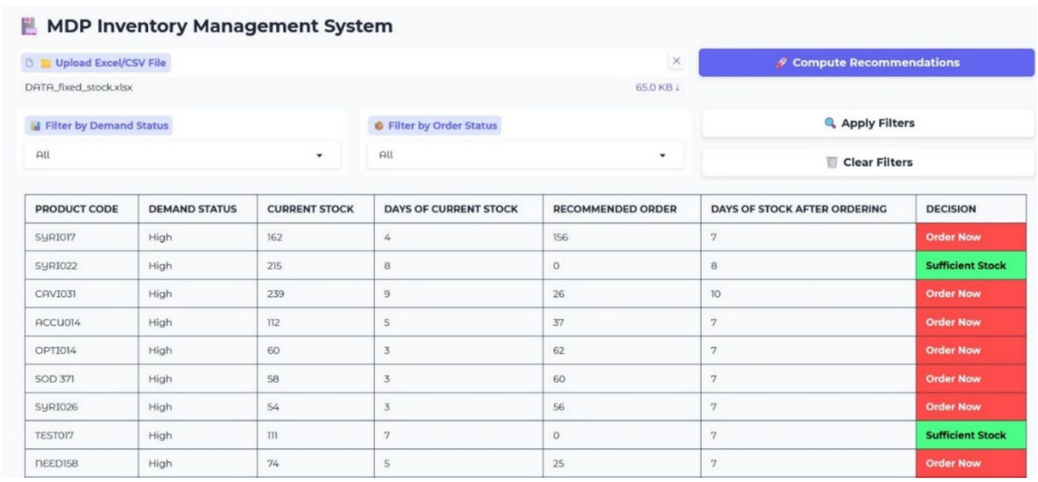

Figure 5: Interface with data uploaded and order flagging

5. SYSTEM VALIDATION AND RESULTS

Through the integration of structured Excel inputs, probabilistic modelling, and an intuitive interface, the ICU inventory management system converts raw data into an actionable decision-support tool. The MDP accounts for demand variability, balancing the risks associated with overstocking versus stockouts. The interface also facilitates the translation of complex outputs into accessible guidance for pharmacy and clinical personnel. This integration effectively mitigates shortages, reduces waste caused by expiration, and enhances the continuity of ICU services through proactive, data-driven inventory management.

5.1. Verification and validation

The private hospital under study relied on a simple heuristic, such as a fixed coverage rule of maintaining seven days of stock, to guide pharmaceutical inventory replenishment decisions. While intuitive, these policies struggled with demand variability and frequently resulted in either excess inventory or dangerous stockouts. To address this problem, an MDP inventory system was developed that dynamically adapts ordering decisions based on current inventory state, outstanding orders, and demand uncertainty.

Verification and validation were conducted against seven predefined requirements. A controlled evaluation compared the MDP approach with the hospital's seven-day coverage heuristic using identical demand scenarios to ensure a fair evaluation. The methodology included four representative pharmaceutical products spanning diverse demand patterns, evaluated over 60-day periods using 500 paired Monte Carlo simulations per product, with bootstrap hypothesis testing ($\alpha = 0.05$) providing statistical validation.

Requirement verification

System verification confirmed compliance with all functional requirements through systematic testing and empirical evidence, as shown in Table 2. The service level, defined as the percentage of total demand successfully met from available stock, served as the primary measure of inventory system effectiveness in meeting demand. Stockout reduction was measured using a paired Monte Carlo simulation (500 runs x 60 days per product), tracking the proportion of days when inventory reached zero and could not meet demand, with bootstrap hypothesis testing ($n= 2000$, $\alpha=0.05$) confirming the statistical significance of the observed reductions. The verification process validated both algorithmic performance and practical implementation capabilities for diverse demand patterns and operational scenarios.

Table 2: Requirement verification

Requirement	Verification evidence
Demand variability adaptation	Service level improvements: 10.8–41.2% for diverse demand patterns (high-volume, high-variability, intermittent).
Inventory visibility enhancement	Real-time stock tracking implemented; system provides current levels, in-transit quantities, and expiration status.
User-Friendly interface	Intuitive policy lookup function: ordering decisions are computed automatically from the current state.
Stockout reduction	2.6–39.2% reduction in stockout rates for all test products ($p<0.05$), statistical significance.
Minimal historical data	The system operates with only the mean daily demand and standard deviation; no complex forecasting is required.
Policy adherence	Respects maximum inventory limits (100% compliance), lead time constraints, and weekly review cycles.
Operational feasibility	Compatible with existing infrastructure; MDP state spaces range from 20 to 16 999 states with sub-second computation.

Validation results

Performance across demand patterns

Table 3 demonstrates the MDP's superior adaptation to demand variability, with service level improvements ranging from 10.8% to 41.2%. Performance gains were most pronounced for high-variability products such as CAVI031, indicating effective handling of uncertain demand patterns. All improvements achieved statistical significance ($p<0.05$).

Table 3: Validation results

Product	Demand profile	Hospital service level	MDP service level	Service improvement	Stockout reduction
SYRI017	High-volume (46 units/day)	83.7%	94.5%	+10.8%	12.6%
CAVI031	High- variability (27 units/day)	45.5%	86.6%	+41.2%	39.2%
OPTI014	Standard profile (18 units/day)	77.6%	96.9%	+19.3%	18.8%
ENDO064	Intermittent (0.1units/day)	68.6%	93.8%	+25.2%	2.6%

6. CONCLUSION

This research addressed critical inventory management problems in the ICU of a 197-bed private hospital in South Africa, where frequent stock-outs threatened patient care and operational efficiency. Through a systematic root cause analysis, requirements engineering, and system design, an MDP-based inventory management system was developed to enhance stock availability while minimising costs.

6.1. Key findings and contributions

Root cause analysis identified demand variability and inflexible replenishment policies as primary drivers of inventory difficulties. Pareto analysis revealed that 80% of inventory value was concentrated in just 25 items (20% of the stock), validating differentiated management approaches.

The MDP-based system achieved an increase in service level of 10.8% to 41.2% and stockout reductions of 2.6% to 39.2% for diverse demand patterns, compared with the hospital's seven-day coverage heuristic. The system proved operationally feasible for resource-constrained settings, requiring minimal historical data and using an intuitive web-based interface. This research demonstrates practical MDP application in healthcare environments with limited data availability, combining probabilistic modelling with classification-based inventory control.

6.2. Research limitations

The analysis relied on simplifying assumptions about demand behaviour and deterministic lead times. Empirical validation used data from a single private hospital with limited historical scope (60-day periods), potentially limiting generalisability to public facilities or different contexts. The model focused on operational efficiency without explicitly incorporating patient-centred performance measures.

6.3. Future research

The limitations of this study provide several opportunities for further research. First, the analysis relied on a set of simplifying assumptions about demand behaviour and replenishment policies. Future work should therefore investigate more complex demand environments and stochastic supply processes, potentially through hybrid models that integrate machine learning with classical inventory control frameworks. Second, the empirical validation of the proposed approach was constrained by the availability and quality of the data. Further research could examine the integration of real-time data acquisition technologies, such as RFID and IoT-based monitoring, to improve inventory accuracy and system responsiveness. Third, the study was conducted in a specific healthcare context, which may limit the generalisability of the findings. Comparative studies of diverse healthcare systems – public and private, urban and rural – would provide insight into the scalability and adaptability of the methodology. Finally, while the model focused primarily on operational efficiency, future studies could extend the framework to incorporate patient-centred performance metrics, such as service levels and treatment continuity, thereby balancing efficiency with quality-of-care outcomes.

By addressing these limitations, future research could contribute to the design of more robust, adaptive, and resilient inventory management systems for healthcare supply chains.

REFERENCES

[1] **H. Coovadia, R. Jewkes, P. Barron, D. Sanders, and D. McIntyre**, "The health and health system of South Africa: Historical roots of current public health challenges," *The Lancet*, vol. 374, no. 9692, pp. 817–834, Sep. 2009. doi: 10.1016/S0140-6736(09)60951-X

[2] **S. Basu, J. Andrews, S. Kishore, R. Panjabi, and D. Stuckler**, "Comparative performance of private and public healthcare systems in low and middle-income countries: A systematic review," *PLoS Medicine*, vol. 9, no. 6, e1001244, Jun. 2012. doi: 10.1371/journal.pmed.1001244

[3] **M. van Zyl-Cillié, D. Demirtas, and E. Hans**, "Wait! What does that mean? Eliminating ambiguity of delays in healthcare from an OR/MS perspective," *Health Systems*, vol. 12, no.1, pp. 3–21, Jan. 2023. doi:1080/20476965.2021.2018362

[4] **V. Bhakoo and T. Choi**, "The iron cage exposed: Institutional pressures and heterogeneity across the healthcare supply chain," *Journal of Operations Management*, vol. 31, no. 6, pp. 432–449, Sep. 2013. doi: 10.1016/j.jom.2013.07.016

[5] **S. Chopra and P. Meindl**, *Supply chain management: Strategy, planning, and operation*. Boston, MA: Pearson, 2016.

[6] **L.H. Lee, S. Teng, E.P. Chew, K.W. Lye, P. Lendermann, and I.A. Karimi**, "Application of multi objective simulation optimization techniques to inventory management problems," in *Proceedings of the Winter Simulation Conference*, Orlando, FL, Dec. 2005, pp. 1684–1691. doi: 10.1109/WSC.2005.1574439

[7] **South African Government**, "National Health Act 61 of 2003," Government Printer, Pretoria, South Africa, 2004. Available: https://www.gov.za/documents/national-health-act

[8] **F.R. Johnston and C. Beard**, "Forecasting systems for production and inventory control," *International Journal of Operations & Production Management*, vol. 12, no. 5, pp. 4–27, May. 1992. doi: 10.1108/01443579210011381

[9] **C. Lewis**, *Demand forecasting, and inventory control*. London, UK: Routledge, 2012.

[10] **P.C. Bell**, "Forecasting demand variation when there are stockouts," *Journal of the Operational Research Society*, vol. 51, no. 3, pp. 358–363, Mar. 2000. doi:10.1057/palgrave.jors.2600877

[11] **D.K. Barrow and N. Kourentzes**, "Distributions of forecasting errors of forecast combinations: Implications for inventory management," *International Journal of Production Economics*, vol. 177, pp. 24–33, Jul. 2016. doi:10.1016/j.ijpe.2016.03.017

[12] **L. Tiacci and S. Saetta**, "An approach to evaluate the impact of interaction between demand forecasting method and stock control policy on the inventory system performances," *International Journal of Production Economics*, vol. 118, no. 1, pp. 63–71, Mar. 2009. doi: 10.1016/ j.ijpe.2008.08.010

[13] **M. Kalchschmidt, R. Verganti, and G. Zotteri**, "Forecasting demand from heterogeneous customers," *International Journal of Operations & Production Management*, vol. 26, no. 6, pp. 619–638, Jun. 2006. doi: 10.1108/01443570610666975

[14] **M. Lei, S. Li, and S. Yu**, "Demand forecasting approaches based on associated relationships for multiple products," *Entropy*, vol. 21, no. 10, 974, Oct. 2019. doi: 10.3390/e21100974

[15] **V. Karamshetty, H. de Vries, L.N. van Wassenhove, S. Dewilde, W. Minnaard, D. Ongarora, K. Abuga, and P. Yadav**, "Inventory management practices in private healthcare facilities in Nairobi County," *Production and Operations Management*, vol. 31, no. 2, pp. 828–846, Feb. 2022. doi: 10.1111/poms.13445

[16] **M.L. Puterman**, *Markov decision processes: Discrete stochastic dynamic programming*. Hoboken, NJ: Wiley, 2014.

[17] **A.R. Vila-Parrish, J.S. Ivy, R.E. King, and S.R. Abel**, "Patient based pharmaceutical inventory management: A two stage inventory and production model for perishable products with Markovian demand," *Health Systems*, vol. 1, no. 1, pp. 69–83, Jun. 2012. doi: 10.1057/hs.2012.2

[18] **Z. He and W. Jiang**, "A new belief Markov chain model and its application in inventory prediction," *International Journal of Production Research*, vol. 56, no. 8, pp. 2800–2817, Apr. 2018. doi: 10.1080/00207543.2017.1405166

[19] **P.K. Santhi**, “Markov decision process in supply chain management,” PhD dissertation, Madurai Kamaraj University, Madurai, Tamil Nadu, India, 2019.

[20] **G.O. Ferreira , E.F. Arruda, and L.G. Marujo**, “Inventory management of perishable items in long-term humanitarian operations using Markov decision processes,” *International Journal of Disaster Risk Reduction*, vol. 31, pp. 460–469, Oct. 2018. doi: 10.1016/j.ijdrr.2018.05.010

[21] **M. Devnani, A.K. Gupta, and R. Nigah**, “ABC and VED analysis of the pharmacy store of a tertiary care teaching, research and referral healthcare institute of India,” *Journal of Young Pharmacists*, vol. 2, no. 2, pp. 201–205, Apr. 2010. doi: 10.4103/0975-1483.63170

[22] **P. Jonsson and S.A. Mattsson**, “Inventory management practices and their implications on perceived planning performance,” *International Journal of Production Research*, vol. 46, no. 7, pp. 1787–1812, Apr. 2008. doi: 10.1080/00207540600988071

[23] **A.R. Vila-Parrish and J.S. Ivy**, “Managing supply critical to patient care: An introduction to hospital inventory management for pharmaceuticals,” in *Handbook of Healthcare Operations Management: Methods and Applications*. B.T. Denton, Ed., New York, NY: Springer, 2013, pp. 447–463. doi: 10.1007/978-1-4614-5885-2_17

[24] **L. Al-Qatawneh and K. Hafeez**, “Healthcare logistics cost optimization using a multicriteria inventory classification,” in *Proceedings of the International Conference on Industrial Engineering and Operations Management*, Kuala Lumpur, 2011. Available: https://ieomsociety.org/ieom2011/pdfs/IEOM075.pdf

[25] **F.F. Flores and D.C. Whybark**, “Multiple criteria ABC analysis,” *International Journal of Operations & Production Management*, vol. 6, no. 3, pp. 38–46, 1986. doi: 10.1108/eb054765

[26] **R. Huijsman and J.M.H. Vissers**, “Supply chain management in health care: State of the art and potential,” in *The emerging world of chains and networks: Bridging theory and practice*. T. Camps, P. Diederen, G.J. Hofstede, and B. Vos, Eds., Den Haag, Netherlands: Reed Business, 2004, pp. 147–167.

[27] **D. Gupta and B. Denton**, “Appointment scheduling in health care: Challenges and opportunities,” *IIE Transactions*, vol. 40, no. 9, pp. 800–819, Sep. 2008. doi: 10.1080/07408170802165880

[28] **S. Mabizela, H.N. Nakambale, and V. Bangalee**, “Predictors of medicine redistribution at public healthcare facilities in King Cetshwayo District, KwaZuluNatal, South Africa,” *BMC Health Services Research*, vol. 23, no. 1, 1108, 2023. doi: 10.1186/s12913-023-10096-4

[29] **S. Saha and P.K. Ray**, “Inventory management and analysis of pharmaceuticals in a healthcare system,” in *Healthcare systems management: Methodologies and applications: 21st century perspectives of Asia*. P.K. Ray and J. Maiti, Eds., Singapore: Springer, 2017, pp. 71–95. doi: 10.1007/978-981-10-5631-4_7

[30] **D. Shenoy and R. Rosas**, *Problems & solutions in inventory management*. Cham, Switzerland: Springer, 2018. Available: https://link.springer.com/content/pdf/10.1007/978-3-319-65696-0.pdf

[31] **D. Kritchanchai, P. Engelseth, and S. Srisakunwan**, “Implementing and using new information technology in hospital logistics,” in *Proceedings of the 2018 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, Bangkok, Thailand, Dec. 2018, pp.1186-1190. doi: 10.1109/IEEM.2018.8607418

[32] **A. Abugabah, A.A. Smadi and L. Houghton**, “RFID in healthcare: A review of the real-world application in hospitals,” *Procedia Computer Science*, vol. 171, pp. 2637–2646, 2020. doi: 10.1016/j.procs.2023.03.004

[33] **K. Agrawal and H. Matsuo**, “Forecasting and inventory management of short lifecycle products,” *Operations Research*, vol. 44, no. 1, pp. 131–150, Jan./Feb. 1996. doi: 10.1287/opre.44.1.131

[34] **Y. Tadayonrad and A.B. Ndiaye**, “A new key performance indicator model for demand forecasting in inventory management considering supply chain reliability and seasonality,” *Supply Chain Analytics*, vol. 3, 100026, Sep. 2023. doi:10.1016/j.sca.2023.100026

[35] **P. Kelle, J. Woosley, and H. Schneider**, “Pharmaceutical supply chain specifics and inventory solutions for a hospital case,” *Operations Research for Health Care*, vol. 1, nos. 2–3, pp. 54–63, Jun./ Sep. 2012. doi: 10.1016/j.orhc.2012.07.001

[36] **N.H. Mustaffa and A. Potter**, “Healthcare supply chain management in Malaysia: A case study,” *Supply Chain Management: An International Journal*, vol. 14, no. 3, pp. 234–243, May 2009. doi: 10.1108/13598540910954575

[37] **R. Bellman**, *Dynamic programming*. Princeton, NJ: Princeton University Press, 1957.