A COHERENT INVENTORY MANAGEMENT SOLUTION FOR SOUTH AFRICAN PUBLIC HEALTHCARE FACILITIES

Matthew Ray Higgo

Pectora roborant cultus recti

Thesis presented in fulfilment of the requirements for the degree of
Master of Engineering (Industrial Engineering)
in the Faculty of Engineering at Stellenbosch University

Supervisor: Dr Joubert van Eeden
Co-supervisor: Prof Sara Susanna Grobbelaar

March 2020

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

March 2020

Abstract

The internal pharmacy and supply chain department at healthcare facilities (hospitals and clinics) are tasked with the acquisition and distribution of stock for the entire establishment. In some cases, the physicians are responsible for quantifying the orders issued to these departments, but have little-to-no visibility of inventory data. Orders are not always delivered in full and may arrive late. Restock orders are made when inventory is still plentiful and stock regularly gets discarded due to expiring. Subject matter experts were consulted and site visits were conducted at healthcare facilities to identify the cause and effects of the inventory problems. A study was executed to define the behaviour of South Africa's healthcare supply chain and investigate the importance on minimizing cost during stock acquisition. It was found that meeting demand in healthcare is most important and the best means of diminishing cost is by reducing the number of expired items.

A systematic literature review was performed to identify inventory policies created to quantify orders for healthcare facilities. Twelve inventory policies were found, of which eleven were tested by means of a simulation modeled to behave as a small public healthcare facility reliant on one hundred products. These inventory policies performed very poorly due to the very infrequent ordering schedule (review periods) and long lead times experienced in South Africa's public healthcare supply chain. All found inventory policies used the moving average forecast technique to predict the future demand. An investigation into alternative forecast methods was conducted which found the *Holt's Linear Trend Method* (HLT) to achieve the best results in terms of accuracy and computational runtime.

The *Iterative Forecast Inventory Model* was created, which works by stepping through the predicted demand while considering the expected order arrivals, to estimate the lacking inventory required to meet upcoming demand. This model outperformed the inventory policies from literature, but decreasing demand trends caused the model to under-assume future demand. The *HLT & ND Inventory Model* was created by including a normal distribution (ND) fit of the historic demand set to calculate a minimum value for the future forecast. This model was capable of meeting all demand and minimizing the number of expired items. The target of acquiring daily inventory levels from public healthcare facilities was no longer possible, but effort is being made to capture inventory levels weekly and log all order information. The *Revised HLT & ND Inventory Model* was designed to estimate and use weekly demand given this degree of visibility. The model achieved promising results and the attention of subject matter experts (SMEs) whom would like to see the model further developed for real-world pilot testing.

Supplier unpredictability was addressed to increase the confidence of acquiring the desired order quantities and a model was created to ensure that inventory storage capacity is not exceeded. Qualitative validation for the models developed in this report was acquired from four supply chain SMEs. The feedback was covered in-depth and concluded positive towards the work done.

Opsomming

Die interne apteek- en voorsieningskettingafdeling by gesondheidsorgfasiliteite doen die verkryging en verspreiding van voorraad vir die hele onderneming. In sommige gevalle is die geneeshere verantwoordelik vir die kwantifisering van die bestellings wat aan hierdie departemente uitgereik is, maar hulle het min, of geen, sigbaarheid van die voorraaddata. Bestellings word ook nie altyd volledig afgelewer nie en kan laat kom. Aankoopbestellings word soms gemaak wanneer die voorraad nog volop is en voorraad word gereeld weggegooi omdat dit verval. Onderwerpkenners (SMEs) is geraadpleeg en besoeke by gesondheidsorgfasiliteite gehou om die oorsaak en gevolge van die voorraadprobleme te identifiseer. ’n Studie is uitgevoer om die gedrag van Suid-Afrika se gesondheidsorgverskaffingsketting te definieer en om die belangrikheid van koste tydens die verkryging van voorraad te ondersoek. Daar is gevind dat die behaling van vraag die belangrikste is, en die beste manier om koste te verminder is om die vervalde items, te verminder.

’n Sistematiese literatuuroorsig is uitgevoer om voorraadbeleide te identifiseer wat geskep is om bestellings vir gesondheidsorgfasiliteite te kwantifiseer. Twaalf voorraadbeleide is gevind, waarvan elf getoets is deur middel van ’n simulasie wat gemodelleer is om op te tree as ’n openbare gesondheidsorgfasiliteit. Hierdie voorraadbeleide het baie sleg gevaar as gevolg van die baie gereelde bestelrooster en lang leityе wat in Suid-Afrika se openbare verskaffingsketting vir openbare gesondheidsorg ervaar is. Al die voorraadbeleide wat gevind is, het die bewegende gemiddelde voorspellingstegniek gebruik om die toekomstige vraag te voorspel. ’n Ondersoek na alternatiewe voorspellingsmetodes is uitgevoer wat gevind het dat die *Holt’s Linear Trend Method* (HLT) die beste resultate ten opsigte van akkuraatheid en berekeningstyd behaal het.

Die *Iterative Forecast Inventory Model* is geskep, wat werk deur die voorspelde vraag deur te gaan terwyl die verwagte aankomste van die bestelling in ag geneem word, om die ontbrekende voorraad te skat wat benodig word om aan die opkomende vraag te voorsien. Hierdie model het beter gevaar as die voorraadbeleide uit literatuur, maar dalende vraagtendense het veroorsaak dat die model die toekomstige vraag onderskat. Die *HLT & ND Inventory Model* is gemaak met behulp van ’n normale verdeling (ND) wat ooreenstem met die historiese vraag wat gestel is om ’n minimum waarde vir die toekomstige voorspelling te bereken. Hierdie model is in staat om aan alle aanvraag te voldoen en die aantal items wat verval het, te verminder. Die verkryging van daaglikse voorraadvlakke was nie meer haalbaar nie, maar weeklikse voorraadvlakke is beskikbaar. Die *Revised HLT & ND Inventory Model* is ontwerp om die weeklikse vraag te skat en te gebruik gegewe hierdie mate van sigbaarheid. Die model het belowende resultate behaal. Onderwerpkenners wil die model in ’n loodstoets toets.

Die onvoorspelbaarheid van die verskaffer is aangespreek om die vertroue van die verkryging van die gewenste bestelhoeveelhede te verhoog, en ’n model is geskep om te verseker dat die voorraadbergingskapasiteit nie oorskry word nie. Kwalitatiewe bekragtiging vir die modelle wat in hierdie verslag ontwikkel is, is verkry van vier onderwerpkenners. Die terugvoer is in diepte bespreek en is positief ten opsigte van die werk wat gedoen is.

Acknowledgements

The author wishes to acknowledge the following people and institutions for their various contributions towards the completion of this work:

- I would first like to thank my thesis supervisors Dr. Joubert van Eeden and Prof. Sara Saartjie Grobbelaar of the Faculty of Industrial Engineering at Stellenbosch University. Dr. van Eeden always had an open door whenever I required advice about my research or writing. Prof. Grobbelaar has tirelessly created opportunities and relationships with subject matter experts incredibly helpful throughout the project. Both supervisors have consistently allowed this project to be my own work, while carefully steering me away from troublesome situations.

- A lot of gratitude goes out to the Bill & Melinda Gates Foundation for providing the funding which supported this full-time research project, 2018–2019.

- A special mention to the Stellenbosch Unit for Operations Research and Engineering (SUnORE) research group for accommodating me in their circle and providing the facilities to work, despite not having an affiliated supervisor. The experience gained, connections formed and memories gathered have been invaluable. I will be sure to continue the pursuit of excellence I have been inspired by during my time with you all.

- I am forever grateful to my parents for their unfaltering support through each of my somewhat reckless and selfish endeavours. Your faith in me has provided additional motivation and been uplifting, and heartening.

- Lastly, I would like to thank my best friend and flatmate, Keeran, for not strangling me in my sleep after hours-upon-hours of thesis related deliberation.

Table of Contents

Glossary

Alpha test: The testing phase for a product before being released into the public. This is most often performed by the developers themselves.

Beta test: The testing phase for a product at the start of release into the public. This is most often less than a month long, but in some cases will exist through the entire development phase – often seen in select computer games. The product is released to a focus group who can provide feedback for any final improvements before the official release.

Developed country: Also called a *developed country*, this term refers to a country that is stable, industrialized and has a capitalist economy. The country provides a high standard of living, long life expectancies and good literacy rates.

Developing country: This term refers to a country that is not as developed as other counties and faces economic, social and political issues.

First-In, First-Out The process of prioritising items based on the order of their arrival. Items which arrived first must be used first.

Healthcare facilities: Hospitals and clinics.

Lead time: The period of time that a facility must wait between placing and receiving an order for some specified product from a given supplier.

Panopticon gaze: A metaphor describing the concept of a powerful entity which quietly watches over and controls the majority.

Prescription stock: Any stock which can only be acquired through the pharmaceutical department inside a dispensary, such as medicine.

Review period: The period of time between placing orders.

Shelf life: The period of time that an item may be held in stock. This is the time between the arrival of the item from the supplier and the item's expiry date.

Stakeholder: Any person or institution which has an interest in the company. This includes both those that can affect the business and may be affected by the business, such as customers, suppliers, staff and capitalists.

The Cloud: Storing, sharing and accessing information via the Internet.

List of Reserved Symbols

Section 4.6: Inventory Policies

Symbol	Meaning
α	service level [%]
α_c	cycle service level [%]
μ	average daily demand [units/day]
σ	standard deviation of daily demand
σ_{L+1}	standard deviation of demand from last $L+1$ days
σ_n	standard deviation of demand from last n days
C_{BI}	cost of inventory levels at beginning of year [$]
C_{EI}	cost of inventory levels at end of year [$]
C_P	cost of inventory purchases during year [$]
$COGS$	cost of goods sold [$]
CSL	cycle service level [%]
D_A	annual demand [units/year]
D_L	expected demand during lead time [units]
D_R	expected demand during review period [units]
DOH	days on hand [days]
E_L	expected number of expired items during lead time [units]
EOQ	economic order quantity [units]
ESC	expected shortage per replenishment cycle [units]
$f_s(\dots)$	normal distribution function
$F_s(\dots)$	cumulative normal distribution function
$F^{-1}(\dots)$	inverse of the cumulative normal distribution function
I	current inventory level [units]
I_A	average annual inventory [units]
I_L	expected inventory level one lead time away [units]
L	lead time [days]
O_L	expected orders to arrive during the lead time
OL	operational leveling factor
Q	order quantity (lot size) [units]
R	review period [days]
s	reorder point [units]
S	par level [units]
SS	safety stock [units]
TO	inventory turnover ratio
U	number of units below the reorder point (undershoot) [units]
z	z-score (normal distribution)

Section 5.1: Forecast Methods

Symbol	Meaning
α	level smoothing parameter, $\in (0; 1)$
β	trend smoothing parameter, $\in (0; 1)$
γ	seasonality smoothing parameter, $\in [0, (1 - \alpha)]$
ϕ	damping parameter, $\in (0; 1)$
b_t	the trend estimate at time t
h	time step forward for the HLT, DT, HWA and HWM forecast methods
l_t	the level of the series at time t
p	window size (finite number of the most recent historic data values)
s_t	the seasonality of the data at time t
t	period [day]
w_j	weight assigned to the j^{th} most recent historic data value
x	number of data points in the historic set
y_t	most recent historic value
$\hat{y}_t$	forecast value for the current period
$\hat{y}_{t+1}$	forecast value for the next period

Section 5.2: Testing the forecasts

Symbol	Meaning
a_t	actual demand at period t
$\bar{a}$	average actual demand across the n periods
$Bias$	Bias [= 0.0 (perfect), > 0.0 (overshoot), < 0.0 (undershoot)]
MAE	Mean-Absolute-Error [smaller is better]
n	number of future periods being predicted
$RMSE$	Root-Mean-Square Error [smaller is better]
t	period [day]
$\hat{y}_t$	forecast demand of period t

Chapter 6: Creating an Order Policy

Symbol	Meaning
α	service level
μ	average daily demand
σ	daily demand standard deviation
D_{SO}	number of days which experienced a stock-out during the 365 day simulation
DOH	days on hand
E_L	expected number of expired items
EOQ	economic order quantity
F_t	forecast value at time t
$F^{-1}(\alpha, \mu, \sigma)$	inverse of the cumulative normally distributed value
I	current inventory level
I_{Max}	maximum inventory level during the 365 day simulation
L	lead time
Q	order quantity

R	review period
s	reorder point
S	Par value
ss	safety stock
TE	total expired items during the 365 day simulation
UD	total unmet demand during the 365 day simulation

Chapter 7: Improving Order Confidence for Real World behaviour

Symbol	Meaning
α	significance level, $\in [0, 1]$
$\mu_{p,s}$	average historic lead time of product p with supplier s
$\sigma_{p,s}$	historic lead time standard deviation of product p with supplier s
$\xi_{x,y}$	efficiency measure of supplier x relative to supplier y
AD	amount delivered
$B_{x,y}$	measure of benefit meeting demand supplier x has over supplier y
$c_{p,s}$	price set by the supplier per batch size
Cp, s	cost of the order for product p with supplier s
$d_{tt,p}$	expected demand on day tt for product p
D_{tt}	forecast demand for day tt
$E_{p,s}$	total expected value for unmet demand of product p with supplier s
ES_{tt}	estimated storage space used on day tt
$f_{p,s}(x_L)$	the frequency of a lead time, x_L, occurring in the historic set
$F_{p,s}(x_L)$	the CDED of a lead time, x_L, occurring in the historic set
$g_{p,s}(Q, AD)$	a second degree polynomial equation fitted to the historic set of orders
h_p	measure of space for one item of product p
H	total holding space available for storage
$i_{tt,p}$	remaining inventory for product p at the end of day tt
I_{tt}	expected inventory level at the start of day tt
$K_{x,y}$	measure of cost with supplier x relative to supplier y
L	lead time
$\mathcal{L}_{p,s}$	historic set of actual lead times
$LC_{p,s}$	lower confidence lead time value of product p with supplier s
m	number of entries in the historic lead time set
$MOQ_{p,s}$	minimum order quantity of product p for supplier s
$o_{tt,p}$	expected order quantity to be delivered on day tt for product p
p	product number
$P_{\max}$	number of products
$\mathcal{P}_p$	supplier priority list for product p
PL	priority level
$POQ_{p,s}$	paramount order quantity of product p for supplier s
$Pr[PL, p]$	preference value for the category of product p given PL, from Table 7.6
q	base order quantity before scaling with batch sizes
$Q_{p,s}$	order quantity after scaling for batch sizes based on product p and supplier s
R	review period
s	supplier number
$t_{m-1,1-\alpha/2}$	t-value for m entries in the historic lead time set at service level α
tt	future time steps in the forecast

U_p	max forecast step for determining the supplier expected values
$UC_{p,s}$	upper confidence lead time value of product p with supplier s
UD_{tt}	expected unmet demand for the forecast day tt

Section 8.1: Inventory Model Validation

Symbol	Meaning
α	service level
$\mu_{\mathcal{D}}$	average weekly demand from $\mathcal{D}$
$\sigma_{\mathcal{D}}$	weekly demand standard deviation from $\mathcal{D}$
$\triangle t$	the number of days between an order's arrival and I_T
d_t	daily historic demand for day t
$\mathcal{D}$	the historic set of weekly demands
$D_{T-1\|T}$	total weekly demand to occur between inventory level recordings I_{T-1} and I_T
$F^{-1}(\alpha, \mu_{\mathcal{D}}, \sigma_{\mathcal{I}}$	inverse of the cumulative normally distributed value resulting from $\mathcal{D}$
I_T	current inventory level (this week)
I_{T-1}	previous inventory level (one week prior)
n	number of historic weekly demands in $\mathcal{D}$
$O_{T-1\|T}$	order quantity that arrived between inventory level recordings I_{T-1} and I_T
R	review period
ss	safety stock

List of Acronyms

5WH: Who; What; When; Why; Where; How.

BOD: Board of Directors

CDED: Cumulative Discrete Exponential Distribution

CEO: Chief Executive Officer

CMD: Cape Medical Depot

DC: Distribution Centre

DSN: Digital Supply Networks

DSS: Decision Support System

EOQ: Economic Order Quantity

FIFO: First-In, First-Out

HLT: Holt's Linear Trend

HWA: Holt's Winter Additive

HWM: Holt's Winter Multiplicative

MA: Moving Average

MAE: Mean Absolute Error

MSES: Moving Simple Exponential Smoothing

NA: Naive Approach

ND: Normal Distribution

NDoH: National Department of Health

POQ: Paramount (maximum) Order Quantity

RMSE: Root Mean Square Error

SA: Simple Average

SC: Supply Chain

SDLC: Software Development Life Cycle

SES: Simple Exponential Smoothing

SLR: Systematic Literature Review

SME: Subject-matter expert

SOH: Stock on Hand

SVS: Stock Visibility Solution

WMA: Weighted Moving Average

List of Figures

List of Tables

CHAPTER 1

Introduction

Contents

This chapter will provide context on the purpose and scope of this project. The problem which needs to be addressed will be defined and the objectives which strive to solve this problem will be established.

1.1 Background

According to Gibson [33] South Africans still suffer unequal treatment when receiving primary health care. Since the *Constitution of the Republic of South Africa* was signed in 1996 [62] all South Africans have had the right to equal health care services provided at either state facilities or private dispensaries. In 1992/3, the South African private health sector was responsible for more than 60% of the total healthcare expenses, yet only tended to roughly 20% of the country's population [9, 33]. During 2013, 37% of the annual medical scheme expenses was for the private health sector [23]. The total healthcare expenses was reduced to roughly 50% by 2016 [21]. However, the portion of South Africans tended to by the private sector has remained fairly constant at roughly 20% [9, 19, 82].

Coovadia et al. [17] described that South African health care institutions had to face sudden policy changes during the end of the apartheid era. All prior restrictions to land, political- and economic-positions were done away with, causing great strain on health services which were used to a biased system in terms of age, race and gender. Simply, Coovadia et al. states that post-apartheid South African government had been placed with a daunting task and expected to deal with these changes quickly. Such changes take time to implement correctly, however the immediate pressure from the masses placed the new-found government in the spotlight.

This links up with Gibson's contemplation on how health care environments within South Africa appear to resemble Kinyon's 'panopticon gaze' [44]. The panopticon gaze expresses that a hidden, quiet power has full reign over the lives of a majority. Foucault [28] explains this power by use of the watchtower inside Stateville Correctional Centre. The tower is placed at the centre of the cell house with all existing cells facing inwards towards it. The guards stationed inside the tower can easily view the contents of each cell. Notably, each cell will thus have an equally clear view of the watchtower, as shown in Figure 1.1 [18, 68].

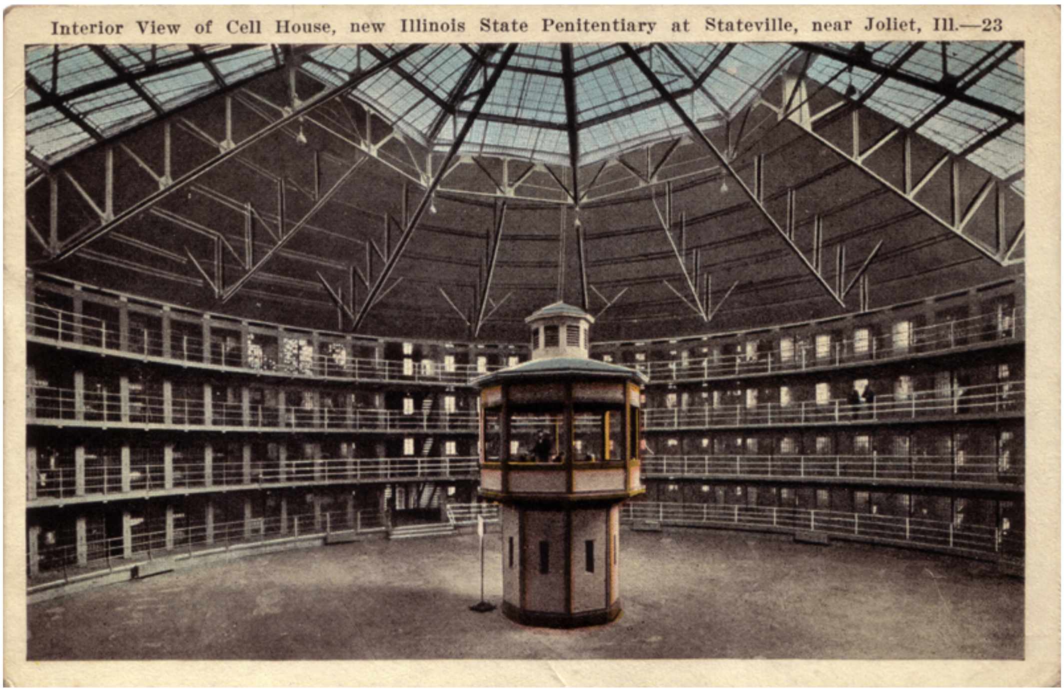

FIGURE 1.1: *Stateville Correctional Center's watchtower: The panopticon concept, from [18]*

To summarise, government receives a large amount of attention when discussing the health care situations in South Africa. Gibson proclaims that this panopticon gaze towards government is a flaw within the system and that government is not the only entity with critical decision making power. Each individual working within each healthcare facility (hospital or clinic) across South Africa have a level of impact to this panopticon gaze, and yet remain invisible to both the community and primary decision makers [33]. Such individuals are not only tasked with saving the lives of patients, but also have a responsibility to uphold the stability of the facility through wholesome business practices.

A joint attempt was made in 2015 by the National Department of Health (NDoH), Vodacom and Mezzanine to assist these individuals with their day-to-day struggles. Together the companies developed the Stock Visibility System (SVS). The SVS uses a mobile application to provide real-time visibility of stock levels in public clinics [14]. Additionally, the SVS will allow the NDoH to monitor the stock received and issued to clinics across South Africa. By July 15, 2016 the SVS had already been deployed to 3 126 clinics across South Africa [13].

1.2 Informal problem description

Strides are being made to put systems in place, such as the SVS, capable of providing stock visibility data to key players such as the NDoH. Data scientists can analyse and interpret the data to better manage the critical supply chain processes at a very high level. However, this data is not yet shared and made visible to the public healthcare facilities from which the original demand occurs. This rekindles Gibson's argument of the panopticon gaze. A means of 'power' is being created for a silent minority, which although helpful, fails to fully acknowledge and assist the individuals positioned within the healthcare facilities who issue orders in the first place.

1.3 Problem statement

A healthcare facility's internal pharmacy and supply chain department are tasked with acquiring and distributing stock for the rest of the establishment. However, in some cases it is the physicians who have the decision making power to quantify and place orders through these departments. Physicians have little-to-no visibility of stock levels nor information on where to find it. Restock orders are made when stock is still plentiful. These orders are not always delivered in full by the suppliers and may arrive late. Items go missing due to poor tracking and organisational systems, while visible stock regularly gets discarded due to expiring and priority mismanagement.

1.4 Objectives

This project will focus on delivering an inventory policy model suitable for South Africa's public hospitals and clinics, improving the confidence of placing stock orders. The following objectives will be pursued during the course of this thesis:

I *Refine* the mentality and focus which is necessary to engage in this research project.

II *Understand*:

- (a) the challenges faced by South African healthcare facilities,
- (b) the organisational structure of a modern healthcare facility,
- (c) how stock is procured,
- (d) the existing relationships between local healthcare facilities,
- (e) South Africa's healthcare supply chain, and
- (f) the importance of cost in an inventory model.

III *Conduct* a systematic literature review in order to find existing inventory policies used in healthcare to conduct stock orders.

IV *Investigate* if better methods exist for performing recurring elements of the found inventory policies.

V *Test* the inventory policies found in literature.

VI Explore ways to *improve* the inventory policies found in literature.

VII *Consider*:

- (a) any real-world phenomenon which may affect orders, and
- (b) a plausible solution to these cases.

VIII *Validate* the final model.

IX *Suggest* any future work which may further improve the system.

1.5 Scope

The aim of this project is to improve orders issued from healthcare facilities. The investigation will include research on the relationships between all healthcare facilities, but the focus of this project lies in the public sector where the complications are most prominent. This is illustrated in Figure 1.2. This project will not attempt to influence the behaviour nor performance of the suppliers and distribution centres upstream of the healthcare supply chain, but will still investigate the effect suppliers have on acquiring the desired stock.

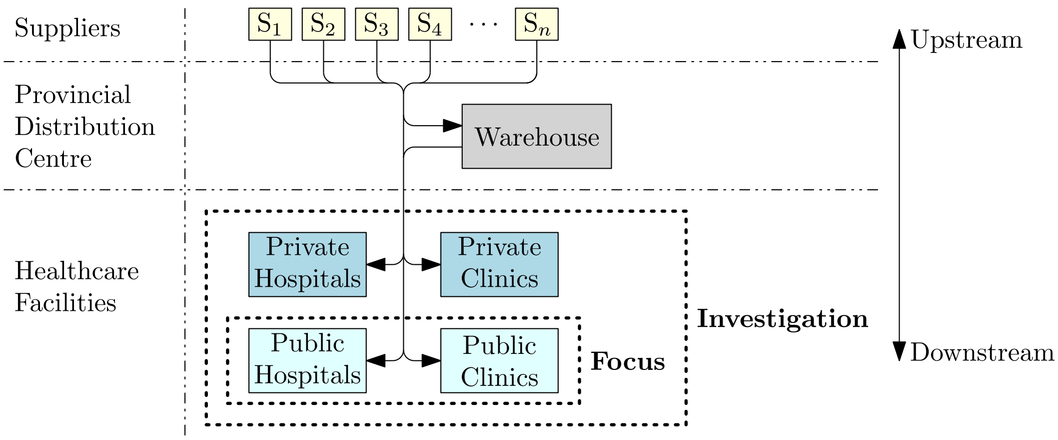

FIGURE 1.2: *Project scope with regards to the supply chain.*

There are over 4 450 public healthcare facilities (hospitals and clinics) across South Africa [54]. It is not possible to perform a site visit to each of these institutions. A combination of private and public healthcare facilities will be selected to visit in order to identify the common attributes of each. This includes, both small and large, hospitals and clinics. What is learnt from the site visits and subject matter experts will be the primary form of guidance towards solving the problem. Literature studies will be used to investigate possible solutions.

CHAPTER 2

Methodology

Contents

This chapter will develop a suitable strategy for conducting this research project in order to achieve the desired objectives that have been described in § 1.4. Several popular research methodologies will be investigated to provide inspiration towards the final report methodology, which will be defined in detail at the end of the chapter.

2.1 Methodology frameworks in literature

This project is attempting to research and develop a feasible inventory policy model capable of improving stock orders in public healthcare. Three popular methodologies that are regularly used in literature are the Agile method, SCRUM method and Software Development Life Cycle [78]. Although these methodologies are most commonly used for software creation, the approach towards research, design and development are fundamentally relevant. Each of these methodologies will be described before ultimately forming the final report methodology framework.

2.1.1 Software Development Life Cycle

Variations of the *Software Development Life Cycle* (SDLC) have emerged over time to accommodate different work, but the methodology ultimately follows several key steps: *Planning*, *Analysis*, *Design* and *Implementation*. The SDLC uses a step-by-step approach to specify exactly what needs to be done before beginning to construct the software. Developers are pressed to plan how elements of the software will interact before starting to code. This makes it easier for future developers to understand the decision making process which was followed [20, 53].

The earliest form of the SDLC is the waterfall model shown in Figure 2.1. Developed by Dr. Winston Royce in 1970 [70], the waterfall model was very popular due to its simplicity. However, Royce noticed that the testing phase often encountered unforeseen circumstances. This is the first phase which implements actual data transfer and timing. Any complications that occurred during this phase would result in a prolonged redesign.

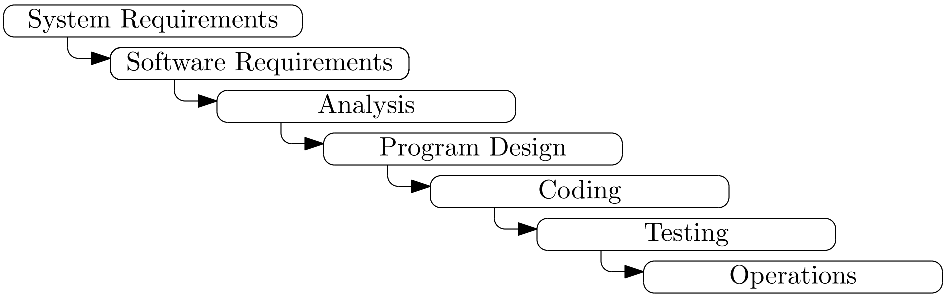

FIGURE 2.1: *Royce's original SDLC waterfall model, from [70].*

Royce suggested creating two step-back loops; one from the testing phase to the program design phase, and another from the program design phase to the software requirements phase. This allows any new-found realisations to be evaluated in accordance with the original software requirements before the program design may be changed for testing. This reworked waterfall model is shown in Figure 2.2.

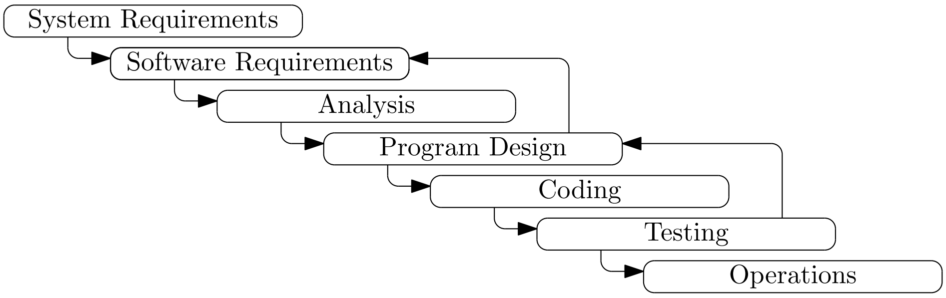

FIGURE 2.2: *Royce's reworked SDLC waterfall model, from [70].*

Er Parag Verma [87] created a variation of the SDLC which grouped the system requirements and software requirements phases together to create a more diverse phase called *investigation*. Additionally, the *program design* and *coding* phases are merged to create a single *design* phase. Verma concludes the new model by breaking down the *operations* phase into two new phases called *implementation* and *maintenance*. This model is shown in Figure 2.3.

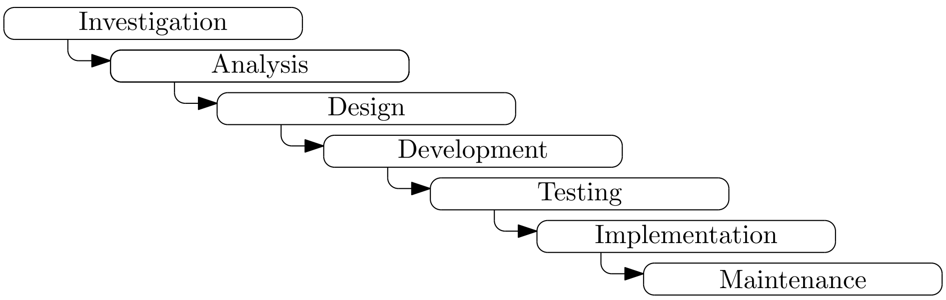

FIGURE 2.3: *Verma's SDLC waterfall model, from [87].*

Verma's waterfall model allows the user to scope their investigation outside of the system and software requirements. Users can place focus on design aspects that are essential to the project and ignore extensive discussions on coding.

2.1.2 Agile method

The *Agile method* uses an incremental, iterative approach towards completing a project. Users can avoid the timely planning stage at the start of a project and begin with the design. This is very helpful for software developers who are issued client designs. Iterations are used to review the work and make necessary changes or improvements throughout the project. Each iteration must achieve a working product before attempting to perform a review [47, 78]. Although the Agile method is known for its production speed, it can easily cause a project to finish well beyond the due date. This is most often caused by a lack of client participation. When a client is easily accessible, the regular feedback allows for quick advancements in the project. Similarly, a lack access of to client feedback increases the length of each iteration period. This makes it difficult to meet the client's expectations [47].

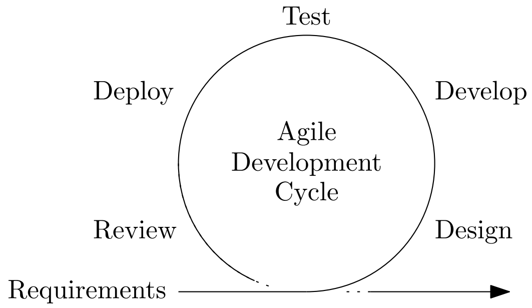

FIGURE 2.4: *Agile Development Cycle, from [78]*

The Agile method contains six phases which form the *Agile Development Cycle*, shown in Figure 2.4 [78]. In the first phase the client meets with the developers to define the *requirements* of the project. The developers then begin with the initial *design* of the product beginning to *develop* it. Once developed, the product is internally *tested* (alpha test) before getting *deployed* for real-world testing (beta test). The product is monitored and *reviewed*. If changes are required, then the development team once again enters the design stage. This loop continues until no further changes are needed [37].

2.1.3 Scrum method

The *Scrum method* was created to industrialise the Agile method [53]. Introduced by Hirotaka Takeuchi and Ikujiro Nonakain in 1986 [84], the Scrum method is executed by working in teams through a relay of short periodic phases called "sprints". Each sprint is no longer than a week or two and ends with a collective meeting of developers and project stakeholders. This meeting is used to discuss progress define the goals of the next sprint [78]. Figure 2.5 demonstrates the Scrum methodology [78].

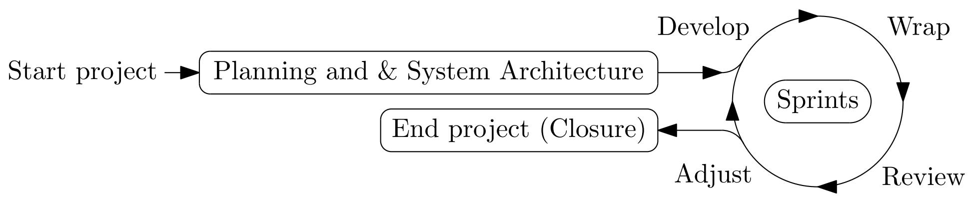

FIGURE 2.5: *Scrum methodology, from [75]*

The sprint loop is not conditioned nor limited to the *Develop*, *Wrap*, *Review* and *Adjust* steps. The loop acts as a "black box" applying the appropriate controls to meet the client's expectations. The project continues to monitor real-world elements such as existing competitors, expected quality and financial pressures until the product is concluded and handed over to the client during the *Closure* phase [75].

2.2 Concluded Project Methodology Framework

The Scrum method is dependent on regular feedback from clients or subject-matter experts (SMEs). This will not be possible during this project. Meetings with SMEs will be very rare and will be primarily used to better understand the workings of the current system. The Agile method can be useful when developing the new inventory policy model, however it lacks the research time allocation that is required for a thesis. Verma's SDLC waterfall model creates a foundation of research to initiate the project. The final methodology framework has been created with the project objectives (§ 1.4) in mind, as shown in Figure 2.6. This framework will be revisited at the end of each foreseeable chapter in order to track progress and identify the achieved objectives.

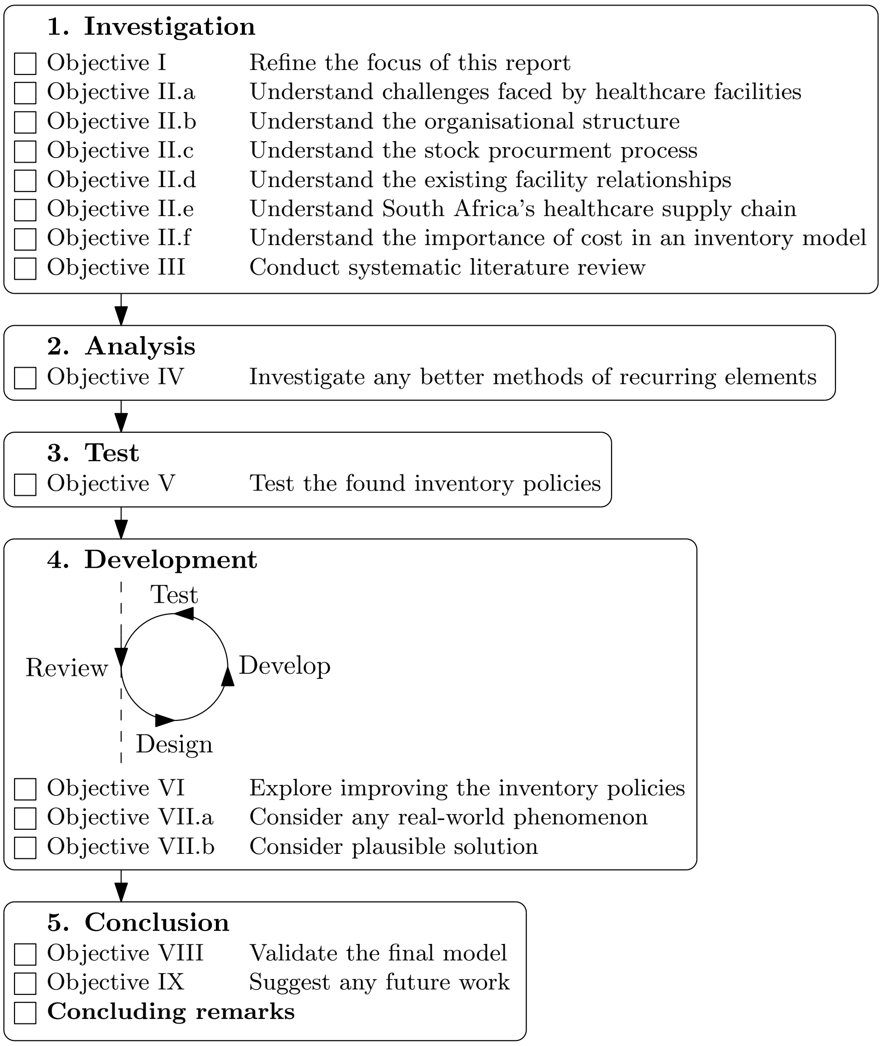

FIGURE 2.6: *Project methodology framework: Design.*

CHAPTER 3

Investigation

Contents

This chapter will start stage one of the project methodology framework (§ 2.2), the *Investigation phase*. Objectives I and II from the project objectives list (§ 1.4) will be the focus of this chapter.

3.1 The Three Worlds Framework

When starting a research project the author(s) may want to first consider the impact that their research will have on both the scientific community and in the everyday lives of the working man or woman. Once research begins, focus can easily shift away from the original objectives of the project. It is important to understand and remember the purpose of the study, what must be delivered and for whom the work is for.

The *Three Worlds Framework* [56] is used to identify where work exists with respect to the world of scientific research and real world problems. As the name suggests, this concept is comprised of three "worlds" which are used to categorize levels of knowledge. The first world, *World I*, refers to information which can be encountered on any given day by any individual. Also known as "lay knowledge", some examples include; riding a bicycle, making coffee or throwing a ball. *World II* refers to information which is obtained through research, such as; scientific studies or large work projects. *World III* describes the range of knowledge surrounding metascience[1]. This means that the third world observes and critiques work that has been done in World II. The *5WH* brainstorming technique [34] was applied to the Three Worlds Framework to make it easier to understand, creating Table 3.1.

[1]The scientific study of science itself [74].

TABLE 3.1: *Three Worlds Framework explained by means of 5WH.*

5WH	World I	World II	World III
Who	Everyone	Scientists; Engineers; Analysts.	Scientists; Engineers; Philosophers.
What	Common sense; Wisdom; Insight; Practical experience.	Searching for facts/truths from a World I problem.	Philosophy, history and ethics of science.
When	Every day.	In the present and future, taking note of the past.	During and after a World II study.
Why	To solve remedial tasks and gain insight into everyday challenges.	In order to advance society and achieve truthful results.	By placing scientific descisions through review, valid answers may be obtained.
Where	Everywhere: home; work; school; shops; etc.	Scientific research facilities, labs and work environments.	Anywhere with access to scientific decisions
How	Observing; Experience; Reflecting on oneself.	Systematic, rigorous studies.	Critical reflection of the work (theory; methodology, research design).

This paper is research-based and thus exists in World II. However, the problem driving this research emerges from World I where physicians and patients experience the effects of poor inventory management. For this reason it is important to consult more than just existing literature. Actual site visits should be conducted to fully comprehend the severity of the problem at South Africa's public hospitals and clinics. The final solution should have the best interests of the physicians and patients whom may depend on it at heart.

3.2 Fact-finding: Site visits

Section 3.1 described the importance of understanding the World I environment. In order to ensure that this awareness is achieved six site visits were conducted in 2018. This section will primarily discuss observations and important insight derived from these visits. Additional learnings from the site visits will be described when necessary in the report. A few final year medical students were willing to identify some key concerns from their experience in different hospitals. No names nor locations will be disclosed as each entity has requested to remain anonymous. This agreement made it easier for the mangers, physicians, staff and students to open up and contribute their honest opinions.

3.2.1 Public Hospital 1

This public hospital encounters very high patient demand. The main entrance to the facility has, over time, rotated to the back of the hospital to provide a larger waiting area for the abundance of patients visiting the facility. Patients are forced to sit on the floor in the corridors with blankets due to the lack of seating. Several occupied hospital beds are stationed in the lobby area due to the lack of available rooms.

The state of the facility is not a result of the abnormally high demand of patients. According to the hospital's general manager, physicians are tasked with determining the necessary order quantities for stock. These requests are sent to the supply chain and pharmacy departments for procurement. However, the physicians have no visibility of the actual stock levels in the hospital nor where products are currently being held. This lack of visibility has lead to additional, unnecessary orders which cause overstocking.

This hospital is also acting as a small local distribution centre (DC) for the smaller clinics in the area. The clinics do not have the necessary storage space to hold the stock required to meet their demand. This is another cause for overstocking taking place at the public hospital.

A few of the operating rooms had become additional storage space to accommodate the stockpile. Stock was poorly organised due to this lack of planning, making it difficult to find the desired products. This facility does not experience many stock-outs due to the overstocking. However, the lack of structure makes it difficult to find anything which leads to expired items and storage wastage. Physicians are hoarding essential medication and tools in their offices to ensure that they had what was necessary to help patients.

The physicians refer to the pharmacy department as the "the accountants". The reason being that the pharmacy department claim the electronic system is showing good supply to demand. However, this system considers demand in terms of the physicians' orders and not the actual patient demand. This makes it appear like all demand is being met when it is not. There is tracking nor count of expired items. This stock will remain unattended and taking up crucial storage space. Despite the chaos from overstocking, the general manger stated that their main concern is determining accurate stock orders to satisfy demand.

3.2.2 Public Hospital 2

This facility does not experience the same high levels of patient demand that the first public hospital did. The hospital was clean and there were no signs of overstocking. The general manager of the facility described their distribution network as a "social relationship" with the local clinics. This means that each healthcare facility is still responsible for acquiring its own stock from suppliers. However, when one of the facilities is experiencing an unexpectedly high demand and needs additional stock, the other facilites will share inventory to assist.

This relationship was created to ensure the success of each healthcare facility and support the well-being of the local community. However, the general manager stated that they were reverting to an independent system which does not share inventory. This is due to the current systems which manage the inventory at the facilities. Stock that is shared to facilities leaves the system in the same fashion as issuing it to a patient. Stock that enters the healthcare facility through sharing does not get registered. This corrupts both the demand and inventory levels at each facility.

It is far easier to revert to an independent system than to create and implement a new system capable of tracking inventory better. The general manager stated that their largest concern for the time being is to reduce or eliminate the stock which expires. Disposing of expired stock incurs high cost.

3.2.3 Private Hospital

This large multidisciplinary private hospital already has a means of tracking inventory. Each ward is provided with its own small storage room holding essential products common to the tasks performed there. The stock on hand is reviewed at the end of every month and redistributed across the hospital according to the remaining shelf life. Items with low remaining shelf life are moved to locations of higher demand. These reviews are coupled with FIFO[2] ('first-in, first-out') to reduce the number of expired items.

The hospital is in the process of improving its systems. They want the new system to provide real-time stock levels and tracking. This system will be paired with hand-held devices (computing tablets) to make it easy to locate desired medication and tools. This will also increase the simplicity of stock takes.

Orders get placed every 1–3 days directly to the suppliers and not through a provincial distribution centre. Orders take roughly two weeks (14 days) to arrive and products with less than 60 days remaining shelf life are sent back to the supplier to be replaced.

3.2.4 Public Clinic

This public healthcare facility is located in the center of a developing township. The patient demand at this facility is incredibly high. On the day of the visit, a queue of patients extended out of the building and into the parking lot area. The staff at this facility work hard to keep the establishment neat and sanitary. The storage room is kept neat and stock is organised alphabetically on shelving around the room. Physicians have to request for items from the inventory manager and are not allowed to simply take stock when necessary. The only electronic system in place is two computers for printing stickers. The inventory manager can print the dosage and instructions that has been prescribed to the patient by the doctor onto the sticker. The sticker is stuck on the medication.

All stock acquisition, distribution and disposal is tracked by pen and paper. There are months of back-logged receipts and forms that need to be processed. The stock kept on the shelves have a drawn line drawn on each of their respective boxes. When inventory level drops below this line a new order should be made. However, these reorder points were assigned many years earlier and are no longer trusted by the inventory manager. With around 200 different products to keep track of, the inventory manager no longer counts inventory levels on a daily basis and orders are missed.

3.2.5 Private Clinic

This is a small clinical pharmacist that mostly assists walk-in patients who need prescription medication. Patients may make appointments to have check-ups or small procedures, such as mole removals. The patient demand at this facility is small and managed though scheduling. There were no concerns with regards to over- or understocking. Expired items are rare due to the low holding stock levels. When a stock-out does occur, alternative medication can be proscribed to the patients. Alternatively, the patient can visit another local pharmacist to acquire the medication.

3.2.6 Medical Students

Three final year medical students, graduating at the end of 2019, heard of this project and offered to share their experiences. Students work in several different public hospitals during the course

[2]Demand is satisfied using the oldest available item which has not yet expired [61].

of their studies. The problem of greatest concern to each of the students were stock-outs. Not having the necessary medication or tools can put lives at risk. Tools as common as syringes and needles come in different tip forms and sizes, each designed for a particular purpose, and should not be used in different ways. Items are incredibly difficult to find because the arrangement of shelving products changes frequently and items are not tracked. Found items are often expired. These items are more often returned to the shelf and ignored, than disposed of.

3.3 Healthcare Facilities

Very often researchers dive into research without taking time to observe the bigger picture. Such observations allow the researcher to break down the problem into smaller, manageable problems. This section will conduct research on how hospitals are structured and run.

3.3.1 Healthcare facility organisational structure

Just like any functioning business, a healthcare facility requires a structure to management called the organisational structure. Larger facilities, like hospitals and public clinics, require very complicated organisational systems. Smaller facilities require simpler systems. There are five common services to any healthcare organisational structure. Table 3.2 sub-categories the departments and jobs found in large healthcare facilities into these five services [66, 67].

TABLE 3.2: *Organisational structure within a large healthcare facility, adapted from [66]*

#	Services	Sub-categories	
1	Administrative Services	CEO BOD Hospital president	Vice president Department heads
2	Informational Services	Admissions Billing Collections Health education	Human resources Information systems Medical records
3	Therapeutic Services	Dietary Medical Psychology Nursing Occupational Therapy Pharmacy	Physical Theraphy Respiratory Therapy Social Services Speech Pathology Sports Medicine
4	Diagnostic Services	Emergency Medicine Medical Imaging	Medical Laboratory
5	Support Services	Biomedical technology Central supply Housekeeping	Maintenance Biomedical technology

The organisational structure in Table 3.2 has existed for many years and can be depicted by the hierarchical chart shown in Figure 3.1. The concept of managing through a hierarchy has become very discouraged. This is because information moves much slower through a hierarchy and employees are not able to be innovative due to strict procedures. Levit argues that this creates a dejected environment with distrust and a lack of talent recognition [48].

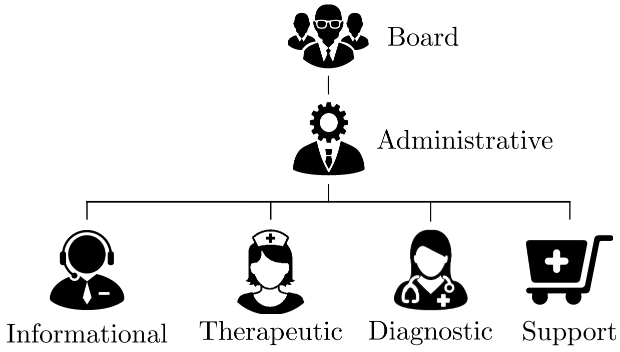

FIGURE 3.1: *Traditional healthcare facility organisational chart, adapted from [66, 67].*

A more modern approach to the traditional organisational chart is the organisational pyramid, shown in Figure 3.2. The pyramid is used to symbolise the importance of each service throughout the business. Without a strong foundation in place the rest of the business is sure to fail. The base of the pyramid is the largest, which implies that successful business requires a strong foundation. This foundation is in the hands of the employees which make the business run. These employees deserve to be heard and respected for their efforts. The chief executive officer (CEO) and board of directors (BOD) appear at the top of the pyramid to watch over the business, but do not form part of the foundation. They must be mindful of the importance of every employee in order to achieve their greatest desires. This ties in with Gibson's discussion of the panopticon gaze described in the report background, § 1.1.

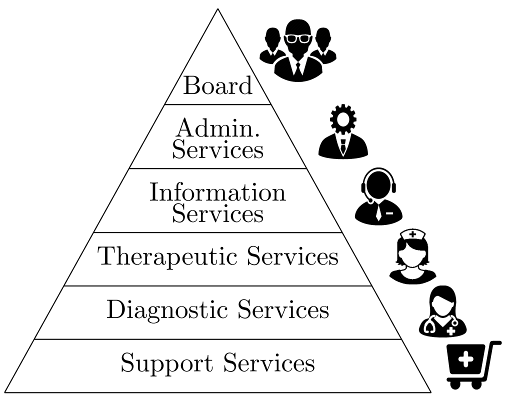

FIGURE 3.2: *Symbolic healthcare facility organisational pyramid, adapted from [66, 67].*

3.3.2 Healthcare facility procurement process

As specified in § 1.5, this project will not reconfigure nor optimise any of the existing supply chain structures in place. However, it is important to understand the procurement process at South African healthcare facilities. The supply chain structures of the visited healthcare facilities (§ 3.2) were examined. Additionally, interviews with subject-matter experts (SMEs) from a provincial distribution centre were conducted. Seven key players were found to influence the procurement process:

1. **Supplier**: Each supplier provides an assortment of medication to chose from. Each healthcare facility will have its own arrangements as to which stock is directly ordered from the supplier, and which stock is ordered through the provincial distribution centre.

2. **Provincial distribution centre**: A large, organised storage facility which distributes stock to hospitals and clinics in the provincial region. This is often an attempt to provide a lower lead time than the original supplier can offer. In South Africa, there is only one large distribution centre per province.

3. **Hospital Warehouse**: The section of the hospital responsible for managing inventory storage and distributing stock appropriately to the wards. Some smaller healthcare facilities lack the necessary storage space to hold all their inventory and will rent an external, nearby storage location. This is evident from the first public hospital which was visited (§ 3.2.1) and was acting as the warehouse for several local, small healthcare facilities.

4. **Wards**: A ward is a division of the hospital which provides a specific form care to patients. Some larger hospitals have additional storage rooms assigned to each individual ward to increase productivity. Clinics do not typically have wards.

5. **Physicians**: These are the doctors and nurses operating within each ward. The physicians at the first visited public hospital (§ 3.2.1) were responsible for requesting stock orders from the pharmacy and supply chain departments. This is irregular. Physicians are not trained to manage stock orders. Inventory acquisition should be fully managed by the pharmacy and supply chain departments.

6. **Hospital Pharmacy**: In charge of ordering medication for the facility.

7. **Hospital Supply Chain**: In charge of ordering hardware, such as beds and scalpels.

Figure 3.3 is a cross-functional flowchart describing the systematic steps of procurement as experienced by a public hospital. This flowchart describes the intentional design of the current public procurement process, but does not always reflect the reality of how public hospitals are functioning. The first public hospital visited (§ 3.2.1) shifted the responsibility of deciding when more stock needed to be ordered onto the physicians in the wards rather than entrusting the task to trained material handlers.

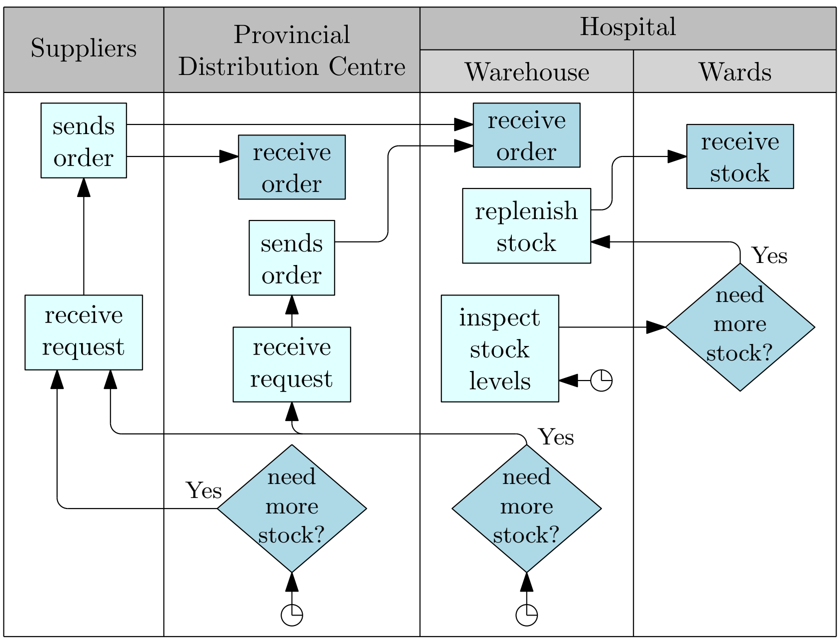

FIGURE 3.3: *Hospital inventory procurement process.*

Similarly, the cross-functional flowchart has been created describing the systematic steps of procurement for public clinics. This flowchart, shown in Figure 3.4, is very similar to that of the public hospital. Hospitals need to ensure that wards remain well-enough is stocked that operations may continue uninterrupted. Clinics only need to maintain the main warehouse (storage room). Small clinics, like the private clinical pharmacist (§ 3.2.5), may have a shopping area with non-prescription medication available for easy purchase by clients. The shopping area will only require small amounts of resupply using stock on hand from the warehouse.

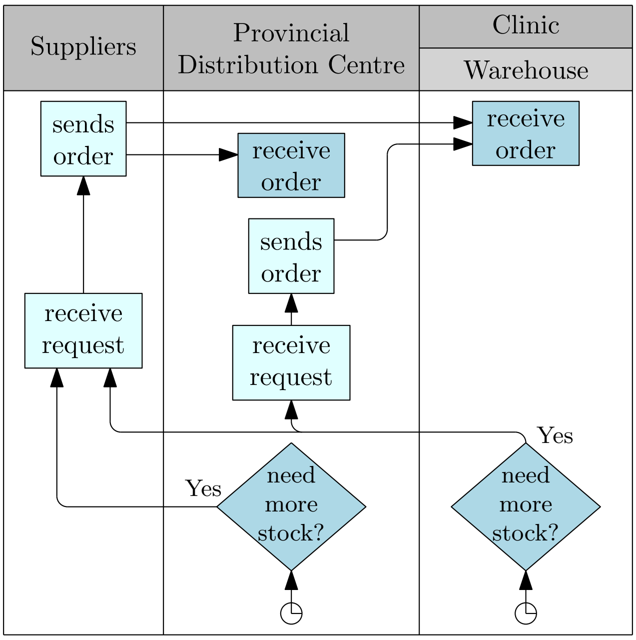

FIGURE 3.4: *Clinic inventory procurement process.*

3.3.3 Healthcare facility relationships

Supply chain structures are designed independently and can differ from one another. The procurement process, described in § 3.3.2, identified key players responsible for managing inventory at a healthcare facility. This section will describe three relationships that were found to exist between local healthcare facilities and the suppliers.

Scenario 1: Classic relationship

These healthcare facilities operate independently and have no association with other local healthcare facilities. The relationship structure exists solely between the facility itself and the suppliers. Stock is either acquired directly from the suppliers or through the local (provincial) distribution centre. Figure 3.5 is a representation of this relationship.

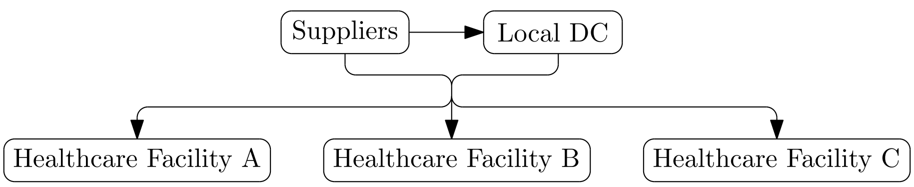

FIGURE 3.5: *Classic relationship supply chain structure.*

Scenario 2: Social relationship

These healthcare facilities place stock orders independently in the same fashion as the *classic relationship*. However, these healthcare facilities share inventory data with other local facilities and support one-another when needed. Figure 3.6 shows this supply chain structure.

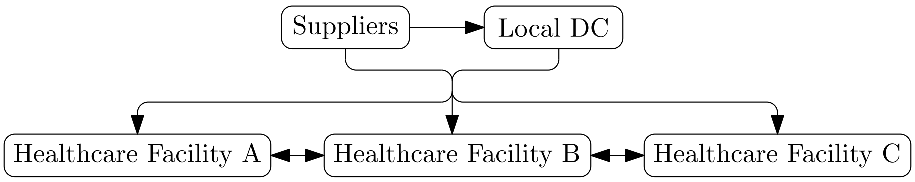

FIGURE 3.6: *Social relationship supply chain structure.*

The *social relationship* theoretically provides support to all parties involved. Being able to acquire essential stock quickly during unforeseen circumstances can remove the stress of waiting for new orders over possibly long lead times[3]. However, the second public hospital site visit (§ 3.2.2) proved that poor inventory management systems can cause more issues than it solves. Careful planning is necessary to ensure that inventory gets tracked accurately without confusing the order policy at each healthcare facility.

Scenario 3: Big-brother relationship

This relationship exists to assist small healthcare facilities whom lack the necessary holding (warehouse) space. As depicted in Figure 3.7, one larger healthcare facility (the 'big-brother') acts as a centralized distribution centre for smaller healthcare facilities incapable of carrying their entire stock in-house. This larger facility is responsible for placing the orders of all supported facilities. Stock is distributed to the smaller facilities in manageable quantities.

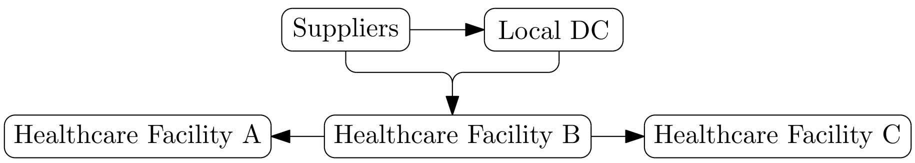

FIGURE 3.7: *Acting centralized DC supply chain structure.*

The smaller healthcare facilities should theoretically experience very short lead times. However, the first public hospital which was visited (§ 3.2.1) claimed to be supporting some facilities in neighbouring towns roughly 1–2 hours away. With an ever-growing population, the hospital cannot sustain such a relationship.

3.4 Conditions of South Africa's Healthcare Supply Chain

The supply chain that drives the public healthcare in South Africa, a developing country, may vary from that of a developed country. This section will investigate the conditions of South Africa's public healthcare supply chain.

South Africa's *National Department of Health* (NDoH) periodically releases information on their contractual agreements with healthcare suppliers. This documentation, known as the *Master*

[3] The period of time between placing an order and receiving the delivery [7]

Procurement Catalogue, is freely available to the public and describes any chances to full agreement. On July 13th, 2018 a complete list of all contractual agreements (1 174 individual contracts) between the NDoH and its suppliers, commissioned 2013–2018, was released [80]. These contracts describe 1 062 medical unique product descriptions which are provided by 78 suppliers. Information on products include the lead time, batch sizes[4] and cost.

3.4.1 Lead times

The 1 174 contracts of the Master Procurement Catalogue were examined and the frequency of lead times have been plotted in Figure 3.8. Only 25 contracts (2.13%) promise to deliver products within one week (≤ 7 days). It was found that 1 134 contracts (96.59%) promised to deliver within two weeks (10–14 days). Only 13 contracts (1.11%) will deliver in three or four weeks (21–28 days). One contract stipulated delivering a product after 3 months (90 days), but this can be classified as an outlier. The public healthcare in developed countries experience lead times of primarily 1–7 days. This means that only 2.13% of South Africa's public healthcare contracts are on par with developed countries.

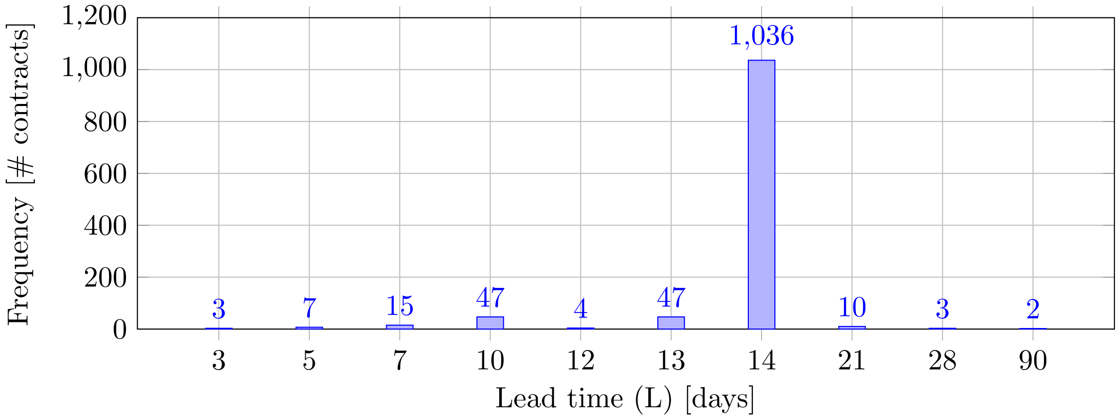

FIGURE 3.8: *Lead time frequency for medical products.*

3.4.2 Minimum order quantities and Batch sizes

Both minimum order quantities (MOQ) and batch sizes are provided in the *Master Procurement Catalogue*. In all cases the batch size of the product was equal to the MOQ. Therefore, the MOQ will be considered as the batch size going forward in the report. Batch sizes are required for determining final order quantity values.

3.4.3 Purchase and delivery cost

Purchase cost is the price associated with buying one batch of a specific product. The purchase cost listed in the Master Procurement Catalogue includes the 15% value added tax. According to subject matter experts (SMEs) from the NDoH, the delivery cost has been included in the purchase cost for each contract. This cost is primarily refers to the delivery from the supplier to the provincial medical depot, but in most cases remains unchanged for direct deliveries to the healthcare facility.

[4]The number of items in a pallet. This amount is specified by the supplier and order quantities must be a multiple of this size.

3.4.4 Review periods

An obtainable list of healthcare facility review periods could not be found. However, during each of the site visits (§ 3.2) the material handler was asked how frequently orders are placed. Private healthcare facilities place orders every 1–3 days, while public healthcare facilities appear to order either once a week or once a month. For example, assume that the healthcare facility is to place an order by 12:00am every Friday. The order may be issued earlier, such as Wednesday, but will still be treated as the agreed upon Friday order. Placing orders earlier than the agreed upon deadline can result in poorer order quantities, due to unpredicted demand behaviour. Because orders are structurally every seven or thirty days in the public sector and can be as low as one day in the private sector, the following review periods will be investigated: $\mathcal{R} = \{1; 7; 30\}$.

3.4.5 Shelf life (expiry dates)

According to the NDoH SMEs, products across the public healthcare supply chain have very long shelf lives ranging from 2 months (60 days) to 3 years (1095 days). This same information was conveyed during a later site visit to a private hospital. Despite the differences between the private and public healthcare sectors, the products are ultimately the same, as are the life expectancies.

3.5 Inventory Cost in Public Healthcare Facilities

This section will research the cause and effect of inventory cost at public healthcare facilities. The types of cost associated with inventory control will be described and a test will be conducted on the relationship between the expenditure of money and meeting demand.

3.5.1 Other types of cost

The process to acquire and store inventory is a costly one. Purchase and delivery costs have already been covered in § 3.4.3. This section will describe the other costs associated with managing inventory in a South African public healthcare facility. This will only consider the general inventory costs experienced by all facilities and will not include facility-specific costs, such as employee salaries, renting additional storage space, cleaning and utility fees.

Carrying cost

Items that are stored in inventory take up space, require security and, in some cases, need maintenance. In healthcare, medicine, drugs, donor blood and vaccine need to be stored under specific conditions, such as cool temperatures. this requires additional utility fees to run necessities like refrigeration or air conditioning. It was mentioned before that utility fees will not be included because each facility has its own unique set-up in this regard. However, these costs may be collectively grouped into what is called the *carrying cost*.

The carrying cost is the annual cost associated with keeping one item and maintaining it in inventory. Because of the complexity that can make up such a measure of cost, it is often expressed as a percentage of the item's original purchase cost. According to SMEs from the NDoH, the carrying cost can theoretically be ignored as it is a "necessary expense". However, in the case of overstocking, the carrying cost would noticeably increase and will thus be considered. The carrying cost (cc_p) will be assumed as 10% per annum of the singular item's purchase cost.

$$cc_p = \frac{10\%}{365 \text{ days}} \times \frac{pc_p}{MOQ_p}$$

Disposal cost

Disposal cost (dc_p) is the cost associated with discarding or destroying waste products. There are many products in healthcare that, due to the risk of their chemical nature, have to be destroyed in the correct, sanitary method. This form of disposal is outsourced and very costly. This cost will vary based on each item's composition and the quantity being disposed. SMEs at a private hospital claimed to spend roughly R1 000.00 to destroy about 20 items.

3.5.2 Multi-objective model: Cost vs Unmet Demand

All the individual variable costs have been defined. However, simply producing an array of order quantities that perfectly meet the predicted future demand is not enough to describe how cost is affected by orders of alternative sizes. By creating an objective to ***minimize* the unmet demand**, rather than to eliminate it, a model can be developed which determines many combinations of order quantities. Each combination of order quantities would then result in some measure of total cost. This brings forth the second objective which is to ***minimize* the total cost**.

A full set of empty orders would result in zero cost, completely minimizing the cost (R0.00). However, with no stock arriving to satisfy the demand the unmet demand would be at its maximum. Similarly, larger orders will result in higher costs and lower unmet demand. At some point the orders will be large enough to satisfy all the demand. Anything ordered beyond that would result in higher total costs, but no longer improving the unmet demand measure.

The *Pareto front* refers to a set of non-dominated solutions. This means that one objective value cannot be improved without worsening the other objective value [46]. The results of a min-min multi-objective model will either appear linear (Figure 3.9a) or in the form of exponential decay (Figure 3.9b).

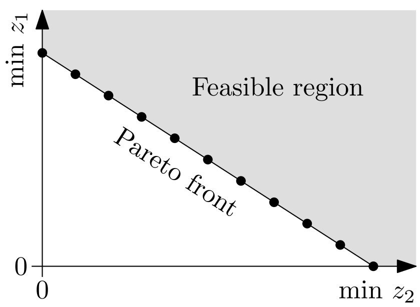

(a) *Linear*

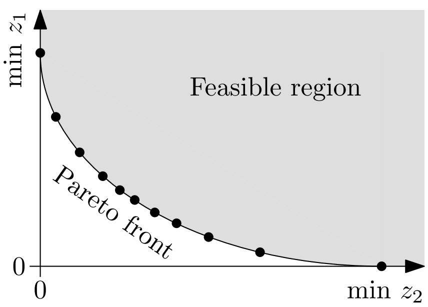

(b) *Negative exponential*

FIGURE 3.9: *Possible relationships for a min-min multi-objective model.*

Decision variables

Decision variables are the values which the model may change. In this model, only the order quantity Q_p of each product p may be changed.

$Q_p =$ order quantity of product p to be placed at time $t = 0$ already scaled according to the batch size MOQ_p

Parameters

Parameters are used to define the characteristics of an object. A review period of 7 days will be used for the test which will requires less computational time to solve than the larger 30 day review period, while adding more complexity than the simplistic one day review period.

$P =$	number of products	$= 100$
$R =$	review period	$= 7$ days
$p =$	product number	$= \{1, 2, \ldots, P\}$
$t =$	time (day)	$= \{0, 1, 2, \ldots\}$
$L_p =$	lead time of product p	$= \{L_1, L_2, \ldots, L_P\}$
$MOQ_p =$	batch size/min order quantity of product p	$= \{MOQ_1, MOQ_2, \ldots, MOQ_P\}$
$T_p =$	max time step for product	$= L_p + R$

One hundred sample products were created in accordance with the conditions of South Africa's healthcare supply chain, investigated in § 3.4 and § 3.5. Shelf lives were randomly generated between 2 and 36 months (60 and 1080 days) in intervals of 30 days, in accordance with § 3.4.5. Figure 3.8 was used to choose the distribution of lead times for the 100 sample products, as shown in Table 3.3. A random selection of contracts were chosen from the Master Procurement Catalogue (§ 3.4.2) that describe different products and achieve the desired lead time distribution in Table 3.3. The respective purchase costs (pc_p) and MOQ of these products were also collected. The sample products list is shown in Tables C.1–C.3, Appendix C.

TABLE 3.3: *Assignment of lead times to the 100 sample products.*

Lead time	L	Count	Percentage	#Sample products
3–7 days	7	25	2.13%	2
10–14 days	14	1134	96.59%	96
21 days	21	10	0.85%	1
28 days	28	3	0.26%	1
90 days	90	0	0.17%	0

Variables

Variables are used to store results that occur in the model. The results that need to be captured in this model is the daily inventory level and unmet demand of each product.

$i_{p,t} =$ inventory level of product p at the end of day t
$ud_{p,t} =$ unmet demand of product p at the end of day t

Objective functions

An objective function describes the intention of the model. This particular model has two objectives. This first objective is to minimize the total cost. The first half of this function solves the purchase cost of the product by determining the number of batches per order and multiplying that by the cost per batch. The second half of the function calculates the carrying cost of the product by multiplying positive inventory levels (from the time the order arrives) by the per-item daily carrying cost:

$$\min\ z_1 = \sum_{p=1}^{P} \left(\left(\frac{Q_p}{MOQ_p} \times pc_p \right) + \sum_{t=L_p}^{T_p} \begin{cases} i_{p,t} \times cc_p & , i_{p,t} > 0 \\ 0 & , \text{else} \end{cases} \right) \qquad \text{[Total cost]}$$

The second objective is to minimize the total unmet demand that occurs:

$$\min\ z_2 = \sum_{p=1}^{P} \sum_{t=L_p}^{T_p} ud_{p,t} \qquad \text{[Total unmet demand]}$$

Constraints

Constraints are the rules that must be adhered to. These rules are constructed to make the model behave realistically. The first constraint makes it that the starting inventory for all products are zero to ensure that the orders are the only thing that has an effect the objective functions:

$$i_{p,0} = 0 \qquad \text{[Starting inventory]}$$

Orders will only get placed on the first day ($t = 0$):

$$o_{p,t} = \begin{cases} Q_p & , t = 0 \\ 0 & , \text{else} \end{cases} \qquad \text{[Placing the order]}$$

Every order quantity must be a whole number, $\mathbb{W} = \{0, 1, 2, \dots\}$, and a multiple of the product's batch size.

$$Q_p \in \mathbb{W} \qquad \text{[Type logic contraint]}$$

$$Q_p \% MOQ_p = 0 \qquad \text{[Batch size logic contraint]}$$

The inventory level at the end of day t is equal to the amount of inventory remaining from the previous day ($i_{p,t-1}$), plus any order placed one lead time ago ($o_{p,t-L_p}$), minus the demand experienced on that day ($d_{p,t}$). The time steps will only continue until the next theoretical order arrival, which would be one review period after the current order has arrived ($t = L_p + R = T_p$).

$$i_{p,t} = \begin{cases} i_{p,t-1} + o_{p,t-L_p} - d_{p,t} & , i_{p,t-1} > 0 \\ o_{p,t-L_p} - d_{p,t} & , \text{else} \end{cases} \quad , t = \{1, 2, \dots, T_p\} \qquad \text{[Daily inventory level]}$$

A negative inventory level indicates the amount of unmet demand ($ud_{p,t}$). The unmet demand will only be counted over the period starting from when the order would arrive ($t = L_p$) until the next theoretical order arrival, which would be one review period later ($t = L_p + R = T_p$).

$$ud_{p,t} = \sum_{t=L_p}^{max} \begin{cases} (-1) \times i_{p,t} & , i_{p,t} < 0 \\ 0 & , \text{else} \end{cases} \quad , t = \{L_p, L_p + 1, \dots, T_p\} \qquad \text{[Unmet demand]}$$

One more constraint will be added to the model. This model is aware of the upcoming daily demand, making it is possible to know the minimum order of each product capable of meeting the full demand once the order arrives. Ensuring that the order quantities never exceed this amount will prevent the decision variables from becoming too large and make sure that solutions do not venture far away from the true, optimal Pareto front.

$$Q_p \leq \left\lceil \frac{\sum_{t=L_p}^{T_p} d_{p,t}}{MOQ_p} \right\rceil \times MOQ_p$$

3.5.3 Solving the Multi-objective model: DBMOSA

The *Dominance-based Multi-objective Simulated Annealing* (DBMOSA) algorithm [79] was selected to solve the model. This stochastic metaheuristic is single solution-based, only retaining solutions that have dominated the feasible set. The fundamental constructs of DBMOSA are:

- An *archive* ($\mathcal{A}$) to store all non-dominated solutions.
- *Stochastic selection* of a *neighbouring solution* for performance review.
- A *reheating parameter*, $\beta \in (1, \frac{5}{4}]$, which encourages accepting worsening solutions.
- A *cooling parameter*, $\alpha \in [\frac{3}{4}, 1)$, which discourages accepting worsening solutions.
- *Epoch* counts for procedure management.

The DBMOSA metaheuristic will operate as follows in order to construct the archive of best solutions (Pareto front) and will continue to iterate until some pre-specified maximum iteration, or some large number of poor epochs [79]:

1. Iteration $\longleftarrow$ 1, Epoch $\longleftarrow$ 1, PoorEpochs $\longleftarrow$ 0, Attempts $\longleftarrow$ 0, Accepts $\longleftarrow$ 0
2. A single feasible solution, $\underline{x}$ is selected as the *current solution* to begin the search. This solution is tested to observe the resulting objective values, $f_1(\underline{x})$ and $f_2(\underline{x})$. The solution $\underline{x}$ is added as an initial entry into the archive, $\mathcal{A}$, along with a record of the resulting $f_1(\underline{x})$ and $f_2(\underline{x})$ values.
3. If PoorEpochs > MaxPoorEpochs, then **stop**.
4. Iteration $\longleftarrow$ Iteration +1. If Iteration > MaxIteration, then **stop**.
5. Another single feasible solution $\underline{x}'$ neighbouring $\underline{x}$ is selected. This solution is also tested for objective values, $f_1(\underline{x}')$ and $f_2(\underline{x}')$.
6. If solution $\underline{x}'$ does not dominate any of the archived solutions in $\mathcal{A}$:
   (a) Reject $\underline{x}'$. Attempts $\longleftarrow$ Attempts +1.
   (b) if Attempts > MaxAttempts, then:
      - PoorEpochs $\longleftarrow$ PoorEpochs +1
      - Epochs $\longleftarrow$ Epochs +1
      - $T = T \times \beta$
      - Attempts $\longleftarrow$ 0, Accepts $\longleftarrow$ 0

(c) Return to Step 3.

7. Else, if solution $\underline{x}'$ does dominate solutions in the archive ($\mathcal{A}$):

(a) $|\tilde{\mathcal{A}}_{\underline{x}}|$ = number of solutionsin $\mathcal{A}$ dominated by $\underline{x}$.

(b) $|\tilde{\mathcal{A}}_{\underline{x}'}|$ = number of solutionsin $\mathcal{A}$ dominated by $\underline{x}'$.

(c) $\triangle E = \frac{|\tilde{\mathcal{A}}_{\underline{x}'}|-|\tilde{\mathcal{A}}_{\underline{x}}|}{|\mathcal{A}|+1}$

(d) $\text{Pr} = \min\{1, e^{-\triangle E/T}\}$

(e) If rand(0,1) $>$ Pr, then reject $\underline{x}'$ as becoming the current solution. Attempts $\longleftarrow$ Attempts +1. If Attempts $>$ MaxAttempts, then:

- PoorEpochs $\longleftarrow$ PoorEpochs +1
- Epochs $\longleftarrow$ Epochs +1
- $T = T \times \beta$
- Attempts $\longleftarrow$ 0, Attempts $\longleftarrow$ 0

(f) Else, if rand(0,1) $\leq$ Pr, then accept $\underline{x}'$ as the current solution ($\underline{x} \longleftarrow \underline{x}'$). Accepts $\longleftarrow$ Accepts +1. If Accepts $>$ MaxAccepts, then:

- Epochs $\longleftarrow$ Epochs +1
- $T = T \times \alpha$
- Accepts $\longleftarrow$ 0, Accepts $\longleftarrow$ 0

(g) Return to Step 3.

3.5.4 Multi-objective results

The DMBOSA metaheuristic described in § 3.5.3 was run on the multi-objective model created in § 3.5.2. The maximum iterations were set to 30 million with an epoch size of 50 accepts/attempts. A maximum of 300 000 poor epochs were allowed to occur. The resulting Pareto front of 353 solutions is shown in Figure 3.10. The multi-objective results form a negative exponential curve as predicted and demonstrated in Figure 3.9b.

The minimum cost to achieve zero unmet demand is R1 134 070.36 and the maximum unmet demand at zero cost is 65 592 units. The next best solution resulted in 528 units of unmet demand (0.805% of total demand) and cost R1 020 971.10 (9.97% cheaper than at zero unmet demand). This may appear as a fair trade-off to a decision maker. However, this large change in cost was caused by drastically reducing the order quantity of one expensive product, Product 14 (R508.30 per item — see Table C.1, Appendix C). To meet the full demand of Product 14 501 items were ordered, but in the cheaper solution only 14 items were ordered. This reduced the total cost to the facility by R247 542.10, but failed to meet 98.6% of Product 14's demand.

Inventory policies want to achieve as little unmet demand as possible. This is very significant in healthcare where the well-being of patients are at stake. The cost of a life is difficult to quantify, but should always be placed above attempts to save money. Another factor, which has not yet been considered, is the confidence of orders being placed. It is not uncommon for suppliers to occasionally deliver orders later than expected or fail to provide the full order quantity. These issues will be addressed separately in Chapter 7 to increase the confidence of orders.

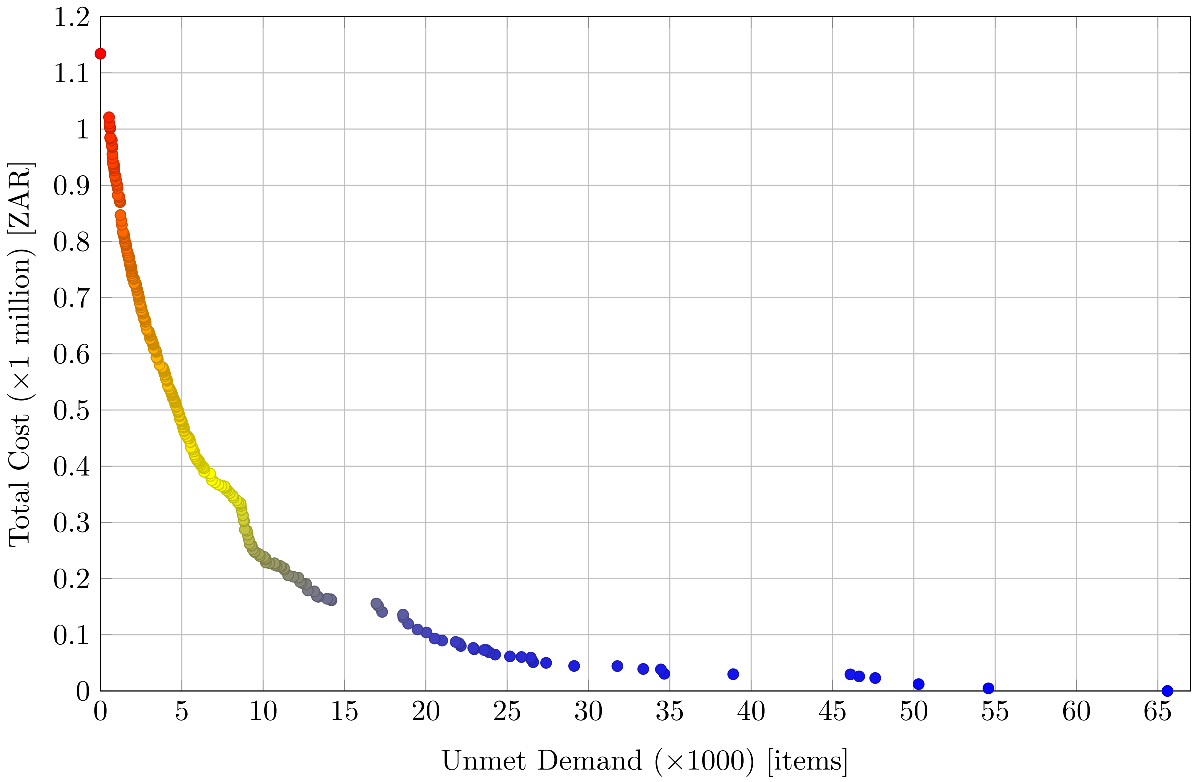

FIGURE 3.10: *Multi-Objective results.*

3.5.5 Concluding remarks

The Pareto front has shown that order sets can be created that produce much better cost margins by only slightly reducing the total unmet demand. However, in the healthcare sector, saving money by reducing order quantities is ethically shunned. All physicians prioritise treating their patients over saving money. There already exists a level of uncertainty that suppliers may not be able to meet full orders. For these reasons, all inventory policy models will be tested for their ability to meet the full demand. Excessive ordering may meet the full demand, but would also cause unnecessarily high carrying (from overstocking) and disposal costs (from expired items). Therefore, the final inventory model should be accurate enough to meet the full demand, prevent overstocking and successfully eliminate expired items.

3.6 Key Focus Efforts

One subject matter expert (SME) from Belgium, with research experience addressing foreign healthcare problems, described three primary concerns for any healthcare environment. Effort must be taken to ensure that focus is given to all three criteria. If one of these three efforts were to fail, then health issues will not get resolved. This research project will be addressing points 1 and 2 of the three focus efforts.

1. Ensure *medicine availability*,
2. Protect physicians' *time availability* to care for patients, and
3. Make sure that patients *get to the healthcare facility.*

3.7 Project Progress

Objectives I and II were achieved during this chapter. The first seven targets in the *Investigation phase* of the project methodology framework have been completed. The updated framework is shown in Figure 3.11.

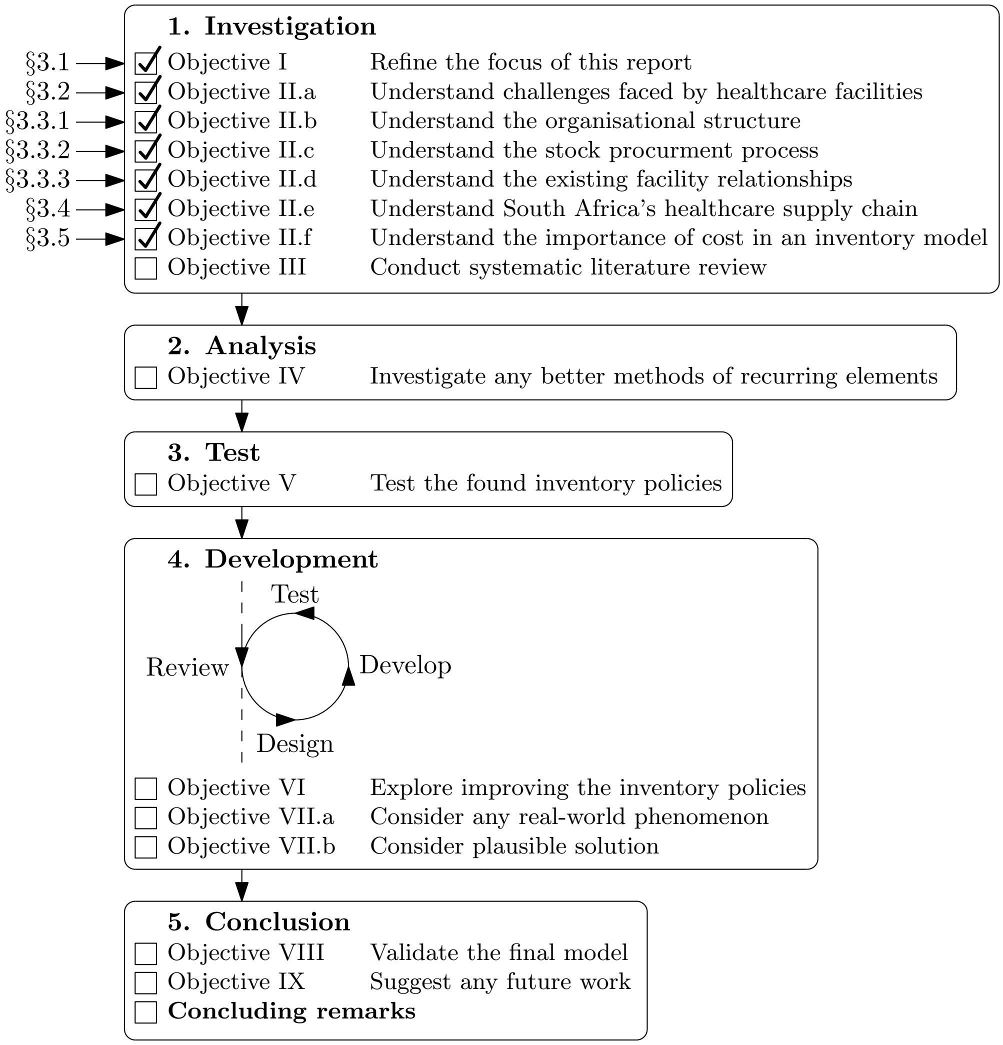

FIGURE 3.11: *Project methodology framework: Chapter 3.*

CHAPTER 4

Systematic Literature Review

Contents

This chapter will address the final target of the *Investigation phase* in the project methodology framework, § 2.2, completing Objective III from the project objectives list, § 1.4.

The *Systematic Literature Review* (SLR) is a powerful tool for users to acquire relevant research papers befitting their topic with reproducible results [10, 42]. In other words, anyone conducting research on the same topic at that same instance in time should ultimately arrive at the same findings. Studies which contribute original content towards a SLR are known as *primary studies*. Studies which make reference to primary studies, such as the SLR, are known as *secondary studies* [10, 45]. Siddaway [76] states that a good SLR should aspire to achieve the following:

- Determine the current extent of progress made with regards to that particular problem,
- Locate any existing gaps in the available literature,
- Be aware of any contradictions and identify why they exist,
- Review the content of the chosen literature,
- Attempt to contribute to and advance the existing research, and
- Describe possible future work.

The SLR is performed sequentially by following a five step process, described below [76]. Additionally, this five step process will form the basis of the SLR which will hereafter be conducted.

I **Scoping**: An introduction to the topic. The research objectives and target audience should be clearly defined. Do an initial search for anyone who might have already conducted a SLR on the defined topic.

II **Planning**: Identify the topic interests and create a list of primary keywords which will later be used as search terms during the *Searching* step.

III **Searching**: Select at least one appropriate search engine from which the literature will be acquired. Using the search terms defined in the *Planning* step, perform several filtered searches using the selected search engine database(s). For more confident results, make use of as many databases as the research time will permit. Be aware of the search fields when using different search engines. Simply put, do not perform searches which merely filter through literature titles, authour keywords or abstracts, but rather explores all of these fields.

IV **Screening**: The found literature must be exported for further review. Some important information to export would include the record title, authour name(s), abstract, year and citation count. The final selection process now involves reading through the title and abstract of each record in order to identify which records is actually applicable to the topic. If time permits, document the reason for rejecting any record and keep a count thereof.

V **Eligibility**: Acquire the real records of the chosen (non-rejected) literature and attentively read through each source.

4.1 Scoping

This is the first step in the SLR. The problem being investigated will be introduced and the objectives of the SLR will be defined.

4.1.1 Introduction to topic

In 2014 there were 4 473 hospitals and clinics (public and private sector) within South Africa providing 119 155 beds [54]. Western Cape News Online [89] (also cited in [40]) explains that 2012 Health Minister, Aaron Motsoaledi's, announced that approximately R14.2 million worth of expired medication had been destroyed over a 12 month period (between April 2011 and April 2012). The greatest cause for expiries is overstocking [26]. During 2012, South Africa experienced shortages for antiretroviral products, used to treat HIV/AIDS, in six of its nine provinces [72]. In 2016 South Africa quantified to a third of southern Africa's new HIV/AIDS infections (270 000 people) and experienced 110 000 deaths due to the infection. Subsequently 7.1 million people in the country were still left living with HIV/AIDS [5].

The need for antiretrovirals in South Africa, as well as other medical supplements, coaxes healthcare facilities into over-ordering medication. This leads to overstocking and expired products. This wastage of both money and resources must be addressed by taking another look at the existing inventory policies. Tools such as decision support systems (DSS) can assist with the inventory management process. Managers may even consider re-evaluating the organisational structure of the hospital to optimise storage space, capture real-time demand values, prevent loosing stock and reduce the number of expiries.

4.1.2 Objectives

This SLR will attempt to review the existing literature which discusses medical dispensary inventory policies, the use of DSSs within hospitals and any indication of noteworthy healthcare organisational structures.

4.1.3 Initial search

After performing a search across search engines *Google*, *Google Scholar*, *Scopus* and *Web of Science*, no similarly conducted SLR was found.

4.2 Planning

Filters will be used to reduce the search until only the most appropriate, topic-related literature remains. The keywords which will later be used during the Searching step of the SLR, § 4.3, can be grouped to form three main searches. The **first search** will look for "inventory management" and "decision support systems". The **second search** will look for "pharmacy", "inventory" and at least one of the following terms: "policy"; "lead time"; "order quantity"; "lot size". The **third search** will look for the terms "organizational structure" and "ward". All healthcare facilities are of interest. Therefore, all three searches will need to either include the keyword "hospital" or "clinic". This also helps to scope the search to relevant literature. A detailed examination at the found records will done in § 4.4 to help identify:

I The *research dates* of conducted.

II The *industry category* to which the records are focussed.

III The most frequently used *author keywords.*

IV The *research locations, observation or site testing.*

4.3 Searching

Scopus (70+ mil records) [24] and *Web of Science* (150+ mil records) [15] were selected as appropriate search engines. Both sources use reliable academic databases and allow the user to export all record information (author, title, abstract, citations, year, type, etc.) which will be required during the Screening stage, § 4.4.

Table 4.1 identifies the number of records found on each search engine during the three searchs. Filters limited results to articles, technical reports, journals, book chapters and conference proceedings. No restrictions were set on dates. Records which were found in both search engines (Duplicates) were found and removed. No duplicates occurred between the different searches.

4.4 Screening

All remaining 188 record abstracts were collected from either Scopus or Web of Science and meticulously read through. Although timely, this step is important to ensure that only literature relevant to the topic gets selected for full review. Only 41 records (21.8%) were relevant to the topic and chosen for acquisition. The selection process is described by Figure 4.1. Here are several examples of abstract topics which were frequently encountered but did not have relevance to the topic:

TABLE 4.1: *Systematic Literature Review: Search terms and findings, conducted 25 May, 2018.*

#	Search terms	Scopus	Web of Science	Duplicates
1	("hospital" OR "clinic") AND "inventory management" AND "decision support system"	10	3	2
2	("hospital" OR "clinic") AND "inventory" AND "pharmacy" AND ("policy" OR "lead time" OR "order quantity" OR "lot size")	68	8	7
3	("hospital" OR "clinic") AND "ward" AND "organizational structure"	108	6	6

- Determining practical solutions to paediatric pain control.
- Manufacturing policies for drug production.
- Nurse experience surveys.
- Organisational requirements for mass burn incidents.
- Patient experience surveys and negative feedback.
- Patient safety questionnaires.
- Training of advanced nurse practitioners.
- Trauma centre characteristics effect on patient outcomes.

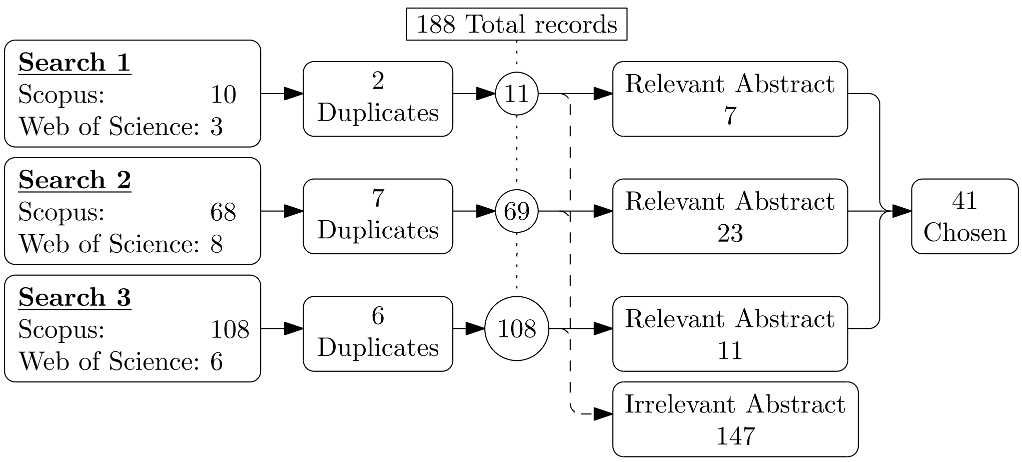

FIGURE 4.1: *SLR: Initial records selection.*

The acquisition process for the chosen records is outlined in Figure 4.2. Only 18 of the records were openly accessible in English. Two records were only available in a foreign language. The twenty-one remaining records had to be requested. Requests were sent to both the Stellenbosch University Library and the authors via ResearchGate. Eight of these records were made available. This means only 26 of the 41 chosen records (63.4%) could be physically acquired for study.

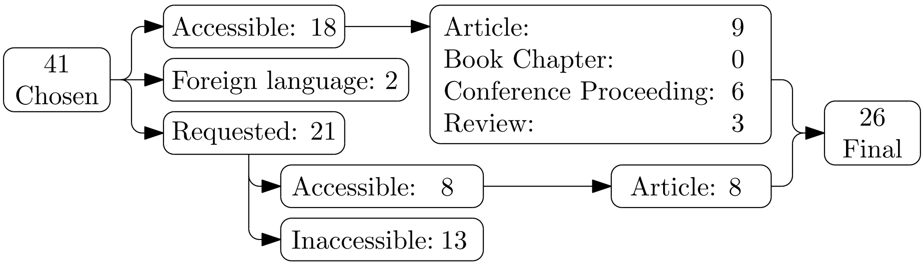

FIGURE 4.2: *SLR: Final records selection.*

Table A.2 in Appendix A on page 108 contains all the source titles for both the chosen and acquired records. The majority of records are articles, making up 73% of the found literature and 65% of the acquired literature. Conference proceedings form the second largest portion of records contributing 17% towards the found literature and 23% of the acquired literature. The final four records consist of three reviews and only a single book source. Only one source came up more than once, the "American Journal of Hospital Pharmacy", and provided five of the chosen records and four of the acquired records.

4.4.1 Research date

Publication dates can be used to identify when a topic was most discussed. Figure 4.3 shows the number of records published in order of date. Publications first appeared in the early 1970s, but only took off at the start of the 1980s. Interestingly, the 1990s had very few publications on this topic. The twenty-first century launched the bulk of this research.

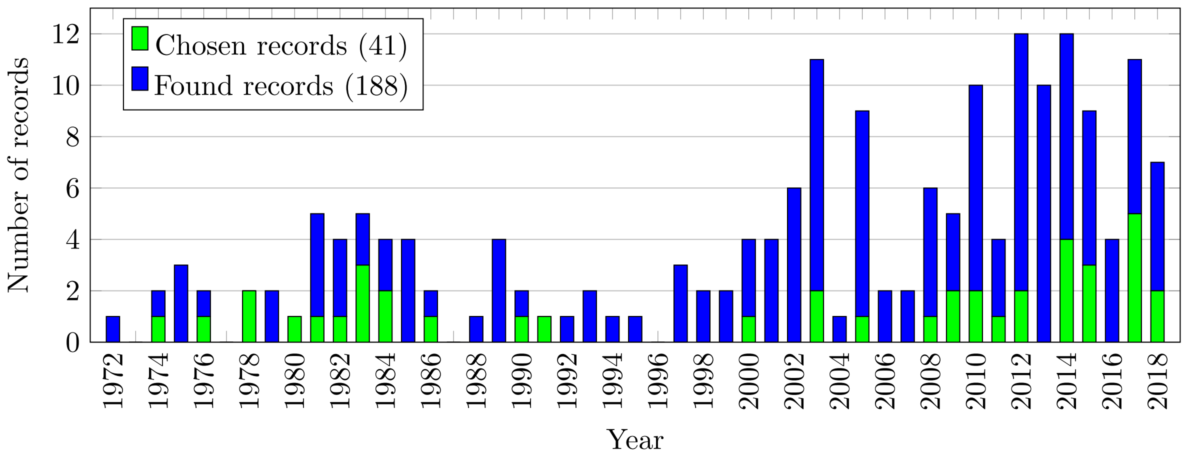

FIGURE 4.3: *SLR: Number of found and chosen records, sorted by date of publication.*

Figure 4.3 as shows that the chosen records are as equally distributed over the timeline as that of the found records. This means that the literature may be able to provide insight into both modern and old systems. However, only 26 of the 41 chosen records were obtainable. Figure 4.4 shows the chosen records on the same time axis, categorized by their accessibility. The acquired records appear to remain well distributed across the timeline, but experience an unfortunate gap in literature from 1985 to 2002. Six of the earlier records were still obtainable and will provide some insight from past studies. It makes sense that most of acquired records were published during the twenty-first century, due to modern forms of publishing. The younger authors are also more likely to respond to a request.

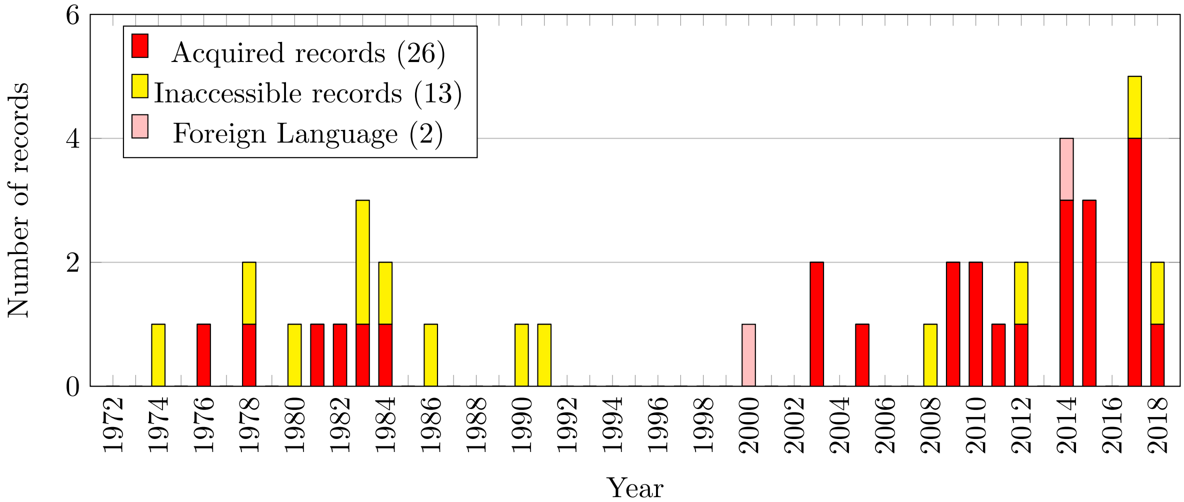

FIGURE 4.4: *SLR: Number of acquired and inaccessible records, sorted by date of publication.*

4.4.2 Industry category

Five different industry types were identified by reviewing the abstracts and author-defined keywords. Each record may encompass more than one industry type. Figure 4.5a shows the distribution of the chosen 41 records based on these five industry types. Similarly, Figure 4.5b classifies shows the distribution of the acquired 26 records. The two polar pie charts appear to have an almost identical distribution of the five industry types. This indicates that the reduced (acquired) literature still embodies the same proportion of industries with that of the original (chosen) literature. The two most prominent industries are Supply Chain and Healthcare. There is also a fair amount of coverage in Information Systems and Distribution. Human Resources makes up the minority of the literature.

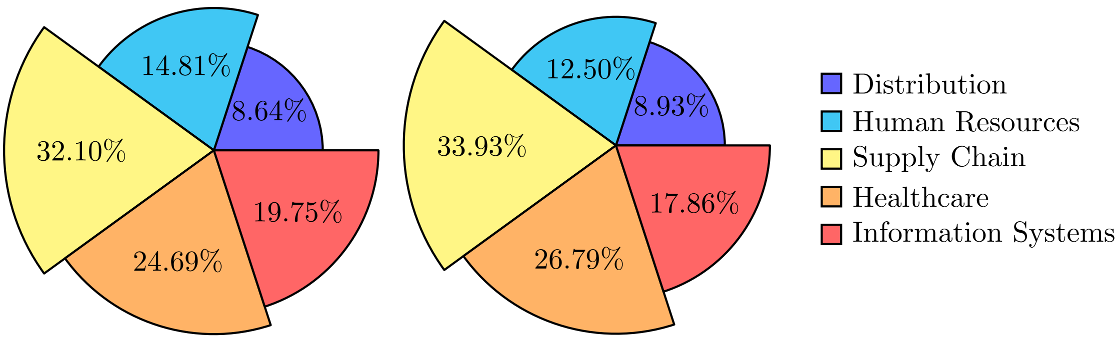

(a) *Chosen records* (b) *Acquired records*

FIGURE 4.5: *SLR: Classification of records by industry type.*

4.4.3 Author keywords

Authors assign relevant keywords to their work which assists researchers with finding it. Figure 4.6 shows all the keywords of the chosen 41 records. Similar terms were grouped together. For example,"Medical car" was grouped with "Healthcare", and "Analysing" was grouped with "Analysis".

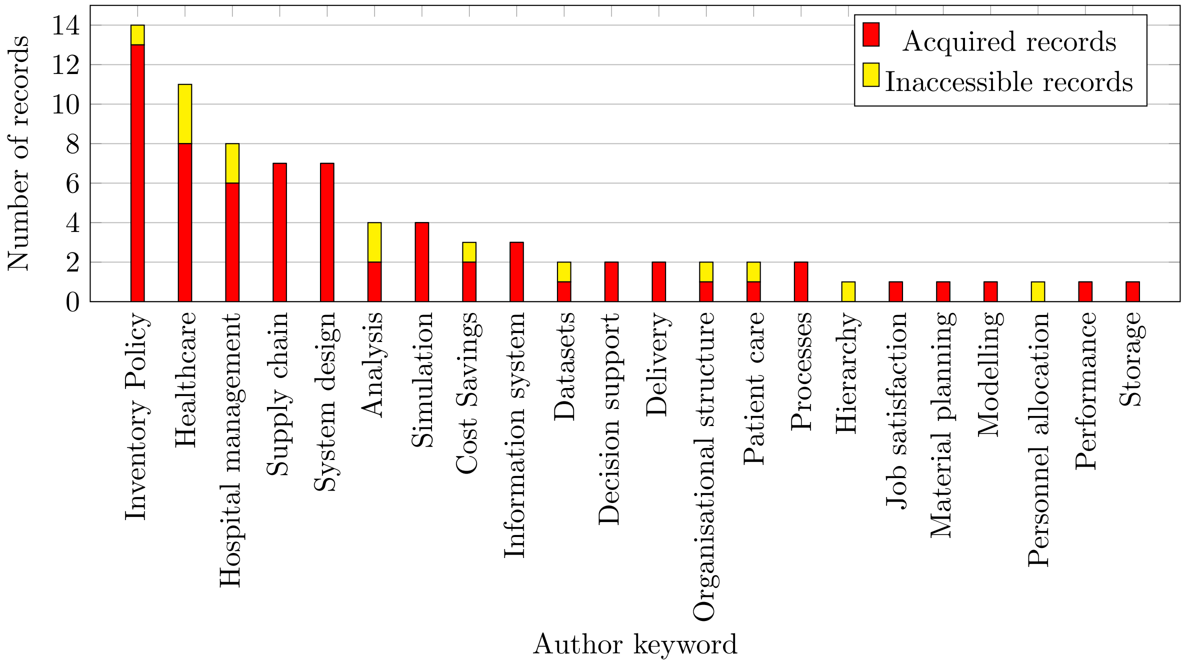

FIGURE 4.6: *SLR: Number of author defined keywords in the chosen literature.*

Some authors will provide several keywords to their work while others may provide none. It is for this reason that the difference in the number of keywords between the acquired and inaccessible records in Figure 4.6 are so small. As such, the bar chart can only provide a rough idea of the discussions covered in the literature. However, it would appear that these records primarily discuss inventory policies, healthcare, hospital management, supply chain and system design.

4.4.4 Research locations

Every location in the world differs to some degree. Infrastructure, law, religion(s), warfare, population density and economics can drastically effect the outcomes of any research project. Studies performed in one region might not be applicable in another region. For example, a wealthy institution in a developed country might have used expensive robotics to solve a particular problem. This would not be helpful to a poor institution in an underdeveloped country trying to solve a similar problem. Table A.1 in Appendix A (page 107) lists the cities in which the acquired records conducted their research.

Figure 4.7 visually represents these regions on the world map (shaded areas). Most of the research (50% of the acquired records) was conducted in the United States of America (USA). Additionally, 38% of the acquired records was conducted in Europe and 12% in Asia. None of this research was conducted near South Africa, which is a developing country. With exception to the two records obtained from Indonesia and Thailand, the acquired literature is dominantly developed country studies. Any observation and finding in these records should be carefully evaluated under South African conditions.

FIGURE 4.7: *SLR: Geographic locations of acquired records, defined by shaded regions.*

4.5 Eligibility

All the acquired records were thoroughly read through. Table 4.2 identifies the topics found with reference to the record(s) which discuss it. The most prominently discussed topic was "Inventory Policies" which featured in fourteen of the twenty-six records (54%) of the acquired records. This corresponds with the bar chart shown in Figure 4.6 proving the expectations of the literature. "Hospital Database Design" was the second most discussed topic, but only consisted of three records. The remaining nine records each explored their own unique topic.

TABLE 4.2: *SLR: Acquired literature topic list.*

Topic	# Records	Reference list
Inventory Policies	14	[1, 8, 22, 29, 32, 39, 41, 50, 52, 59, 36, 86, 90, 91].
Hospital Database Design	3	[6, 71, 92].
Blood Donation Data	1	[4].
Patient Assessment Wards	1	[16].
Computerised Order Entry	1	[18].
Correctional Facility Drug Levels Review	1	[43].
Short Stay Unit Data Review	1	[49].
eHealth Service Delivery	1	[55].
Standard Times of Nursing Activities	1	[57].
Agent-based Simulation for Patient Scheduling	1	[64].
Ward Managers Survey Results	1	[65].
Total	**26**	

Only the topics which are relevant to this paper will be investigated further in the proceeding section(s). The "Inventory Policies" will definitely be revised to explore stock order management. The records on "Hospital Database Design" will also be studied for any helpful insight. In order to understand the benefits of electronic order placements, "Computerised Order Entry" will be reviewed. The concept of patient scheduling sounds promising for a DSS, but this particular record describes solving the problem with complicated agent-based simulations. This will not be achievable in a real-time DSS and hence will not be discussed.

4.6 Inventory Policies

Fourteen records obtained during the SLR, § 4.5, provide insight into inventory policies used for healthcare facilities in developed countries. This section will describe two popular inventory management concepts, namely the 'ABC Inventory Control' and the 'Economic Order Quantity', as well as the unique inventory policies discovered during the SLR. Thereafter, more traditional continuous and period review policies, described by Chopra and Meindl [12], will be explained.

4.6.1 ABC Inventory Control

The *ABC Inventory Control* is used to assist managers with prioritising their focus on the few products that ultimately make up the majority of the inventory cost. As the name "ABC" suggests, there are three product categories [59]. Figure 4.8 provides a graphic representation of the concept. Category "A" refers to the number of products (10–15%) which make up the majority of the inventory costs (70–80%). These are the products which managers should direct most of their attention towards. Category "B" accounts for 20–25% of the number of products and 15–20% of the inventory costs. Category "C" refers to the majority of the number of products (60–70%) that make up the minority of inventory costs (10–15%) and require the least attention [1, 86].

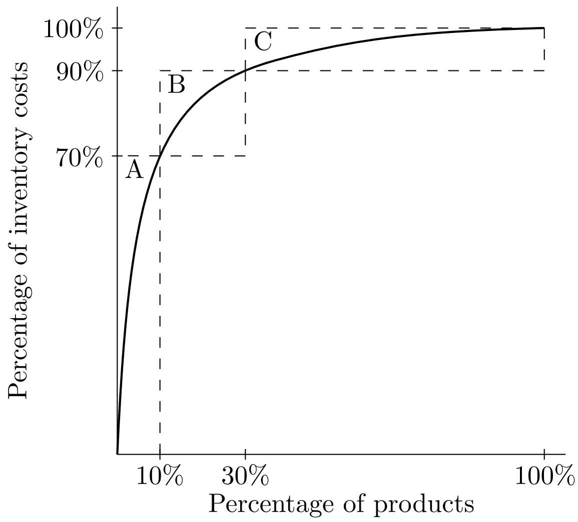

FIGURE 4.8: *ABC inventory control classification.*

4.6.2 Economic Order Quantity

The *Economic Order Quantity* (*EOQ*) refers to the order quantity of a given product that will yield the lowest total annual costs. This is achieved by finding the order quantity which results in equal carrying and ordering costs, as shown in Figure 4.9. The *EOQ* is an ideal number and is often not close to the batch size defined by the supplier. However, due to the small deviation in total cost that occurs as a result of adjusting the order quantity, the change is often deemed acceptable [29, 8].

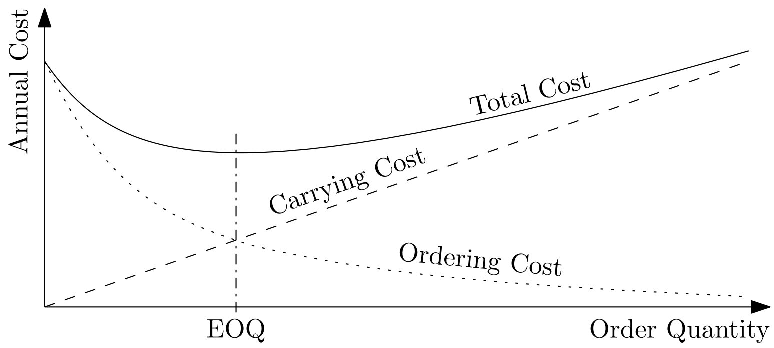

FIGURE 4.9: *Economic Order Quantity curve, adapted from [59].*

The EOQ, the expected total annual cost, and the number of orders required per year can be calculated using Equations 4.1, 4.2 and 4.3 respectively, where: A = fixed ordering cost; D = annual demand; P = purchase cost of each unit; H = annual inventory holding cost expressed as a percentage; Q = order quantity, which is often the EOQ rounded up to the nearest specified batch size [1, 8, 32].

$$EOQ = \sqrt{\frac{2DA}{HP}} \tag{4.1}$$

$$Total\ Annual\ Cost = PD + \frac{DA}{Q} + \frac{QHP}{2} \tag{4.2}$$

$$Number\ of\ orders\ per\ annum = \frac{D}{Q} \tag{4.3}$$

4.6.3 Types of Inventory Policies

According to Wilson et al. [90] all inventory control models can be divided into one of two primary categories, based on their review system. *Continuous Reviews* refer to systems that perform unceasing, real-time inventory level checks. When inventory levels reach the predetermined reorder point value (s) an order gets placed for new stock. *Periodic Reviews* perform inventory level checks at certain intervals, known as the review period (R). The review period varies between facilities, but are often set to either ‘daily’ ($R = 1$ day), ‘weekly’ ($R = 7$ days) or ‘monthly’ ($R = 30$ days). One benefit to this is achieving alignment with delivery schedules across all the stock keeping units for a single supplier. This creates predictability which assists with distribution planning and delivery cost reduction. Noel [59] explains how the periodic inventory method provides simplicity and lower costs, but creates a lack of control and easily permits shortages to occur unnoticed.

Wilson et al. [90] explains that both review methods can be subdivided based on how the orders are quantified. The *Fixed Order Quantity Model* will always use the same order quantity. The *Order-Up-To Model* changes the order quantity with the intent of restocking inventory back to a pre-determined, maximum Par level, S. Figure 4.10 is a flow diagram describing this classification of the four possible inventory policies.

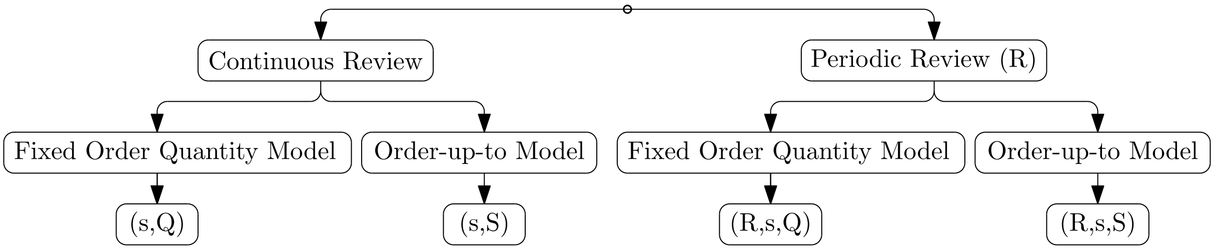

FIGURE 4.10: *Types of inventory policies.*

4.6.4 Inventory Policy Models found in the SLR

The literature found during the SLR, § 4, discussed many variations of determining the reorder point (s), order quantity (Q) and Par value (S). The subsections to follow will describe these policies.

Policy 1 – (R,s,S) Par inventory model

The *Par Inventory Model* appeared often in literature. This order-up-to model operates with periodic reviews, classifying it as a (R,s,S) inventory model [25]. During each review the material handler must perform an item count of the remaining stock on hand (SOH) to determine the inventory levels (I). Orders are placed to suppliers for the number of items which will bring the inventory levels back up to Par. The order quantity must be rounded up to the nearest batch size [81].

Because orders are made whenever stock levels are lower than the Par value (S), the reorder point can be described as $s = S - 1$ and the order quantity is calculated as $Q = S - I$ [90]. In reality, healthcare facilities can have over 200 products. Most of the public healthcare facilities do not have accurate electronic inventory systems in pace. The material handler must manually count the inventory level of each product, look up the unique Par level and calculate the order. Because of this, material handlers often resort to 'scoping' inventory rooms for product bins that stand out as low on stock and only count those items. This means that not all necessary orders get placed or order quantities are made up ignoring the Par Inventory Method. This can result in both stock-outs and overstocking [51].

Policy 2 – (R,s,Q) Kanban Two-Bin Inventory Model

The *Kanban Two-Bin Inventory Model* is designed to be user-friendly and reduce the time for stock management, the chance of human error and the number of expired items [81]. Just as the name 'Two-Bin' suggests, each product is assigned two identical containers (bins) which are packed into shelving one behind the other. Stock is taken from the front bin. When the last item gets picked, the empty bin must be moved to a reorder pile. The material handler will place an order to replace the used items. In the meantime the second bin is pulled forward to supply stock for the workforce [25]. Ordered items are packed into the empty bin and placed behind the current front bin. This enforces FIFO to reduce the number of expired products [51].

Each container is sized to hold one fixed order quantity (Q). This means that the Kanban Two-Bin model is a (R,s,Q) inventory policy with a reorder point of $s = Q$. This system requires a very short lead time (1–2 days) to reduce the need to hold a lot of safety stock. This means the Kanban Two-Bin model requires efficient, nearby distributors. This model allows busy physicians the freedom to pick stock quickly when needed [90, 69].

The Kanban Two-Bin model lacks accuracy of real-time stock visibility. Stock orders are only updated when a bin gets depleted This may cause misinterpretations of real-time inventory levels. For example; if one bin, of some product, is able to hold 50 items and becomes empty, then the system may determine the demand was 50 items. However, more items may have been used from the second bin. This means that the real demand could be anything between 50 and 99 items (up to roughly twice the estimated demand). This lack of visibility may cause problems for analysts attempting to learn from stock data.

Policy 3 – (R,s,Q) by van Der Linde

This model, described by van Der Linde [86], was tested on A-category products, according to the ABC inventory control (see page 35), at St. Luke's Memorial Hospital in 1980. The expected demand to occur over the lead time (D_L) is calculated using the average daily demand (μ) and the product's lead time (L). Thereafter, the safety stock (SS) and reorder point (s) are calculated. The order quantity (Q) is determined using the reorder point, the current inventory level of the product (I) and the Economic Order Quantity (EOQ), described in § 4.6.2.

$$D_L = \mu \times L$$

$$SS = 2.3\sqrt{D_L}$$

$$s = D_L + SS$$

$$Q = s + EOQ - I$$

Policy 4 – (s,Q) by Wilson, Hodge and Bivens

Wilson, Hodge and Bivens [90] describe a continuous review inventory policy model which offers the decision maker two choices for calculating the reorder point value (s) using the average daily demand (μ) and standard deviation of the daily demand (σ). The choice specifies the probability at which the demand may exceed the inventory level. Using $\mu + 3\sigma$ will yield the lowest probability of stock-out, but may be more costly than using $\mu + 2\sigma$. The authors acknowledged that these risk values are very unlikely to hold in a periodic review system.

$$s = \begin{cases} \mu + 2\sigma & ,2.27\% \text{ prob. demand will exceed the inventory level} \\ \mu + 3\sigma & ,0.135\% \text{ prob. demand will exceed the inventory level} \end{cases}$$

The Par value (S) is set to twice the reorder point. The order quantity (Q) takes into account how far the inventory level has dropped below the reorder point value, called the undershoot (U). The undershoot is calculated using the expected demand to occur over the review period (D_R), where R is the review period in days.

$$S = 2s$$

$$D_R = \mu \times R$$

$$U = \frac{(D_R)^2}{2(D_R)} - \frac{1}{2}$$

$$Q = S - s + U$$

Policy 5 – (s,S) by Kelle, Woosley and Schneider

Kelle, Woosley and Schneider [41] created a continuous review order-up-to model designed to find an optimal pre-determined schedule that will minimize the overall inventory costs to the hospital. The model drivers (a_i), model loads (b_i) and fixed A, B, C, D variables are all constants determined by Schneider's Approximation formulas [73]. Unfortunately this process was not explained in the literature and cannot be recreated. The constants provided by the article are specific to their own system and will deliver poor results when used elsewhere. Therefore, this model will not be considered due to the lack of necessary information required to replicate it.

The order quantity (Q) is set to the Economic Order Quantity (EOQ), described in § 4.6.2. The undershoot value (U) is calculated using the average daily demand (μ) and standard deviation of the daily demand (σ). The standard deviation of the last $L+1$ days (σ_{L+1}) and service level (α) were used to calculate the values y, w and $p(y)$, from Schneider's Approximation formulas. Thereafter, the reorder point (s) and Par value (S) could be calculated.

$$Q = EOQ$$

$$U = \frac{\mu^2 + \sigma^2}{2 \times \mu}$$

$$y = \frac{Q \times (1 - \alpha)}{\sigma_{L+1}}$$

$$w = \sqrt{\ln\left(\frac{E}{y^2}\right)}$$

$$p(y) = \frac{a_0 + a_1 w + a_2 w^2 + a_3 w^3}{b_0 + b_1 w + b_2 w^2 + b_3 w^3 + b_4 w^4}$$

$$s = \mu(L+1) + (p(y) \times \sigma_{L+1}) - MAX\left(\frac{\sigma^2}{\mu} - 1; 0\right) \times \frac{A + By}{C + Dy}$$

$$S = s + Q - U$$

Policy 6 – (R,s,S) by Wilson, Hodge and Bivens

Wilson, Hodge and Bivens [90] modelled a periodic order-up-to policy to be used within a cancer centre's ambulatory care clinic which had originally been using the Par inventory model (see page 37). The new model uses the *days on hand* (DOH) to determine an operational levelling factor (OL) which scales the reorder point value (s). The Par value (S) is defined as twice the reorder point value. The expected demand over the review period (D_R) is used to calculate the undershoot value (U) and then the order quantity (Q) can be determined.

$$OL = \frac{DOH}{2}$$

$$s = (\mu + 2\sigma)OL$$

$$S = 2 \times s$$

$$D_R = \mu \times R$$

$$U = \frac{(D_R)^2}{2(D_R)} - \frac{1}{2}$$

$$Q = S - s + U$$

The authors never specified the DOH equation used, therefore the equations were obtained from *Accounting Explained* [2] and *Finance Train* [27]. The DOH is calculated using the inventory turnover ratio (TO), which can be determined with the cost of goods sold ($COGS$) and the average annual inventory (I_A). The $COGS$ is the sum of; the cost of inventory levels at the beginning of the year (C_{BI}), the cost of inventory purchased during the year (C_P), and the cost of inventory levels at the end of the year (C_{EI}).

$$COGS = C_{BI} + C_P - C_{EI}$$

$$TO = \frac{COGS}{I_A}$$

$$DOH = \frac{365}{TO}$$

Policy 7, 8, and 9 – (R,s,S) by Gebicki *et* al.

Gebicki *et* al. [32] performed a comparison of four different inventory policies. All these policies use the expected demand to occur over the lead time (D_L), where μ is the average daily demand and L is the expected lead time. The reorder point of each policy factors in the standard deviation of the daily demand (σ). The order quantities (Q) are determined by taking into account the current inventory level (I).

The first three of these four policies merely test different alternatives to defining the z-score (z). The first policy, **Policy 7**, set the z-score (normal distribution) for all products to 1.96, which delivers a 97.5% service level (α). The Par value (S) is the sum of the reorder point value (s) and the Economic Order Quantity (EOQ), described in § 4.6.2.

$$z = 1.96, \quad \alpha = 97.50\%$$

$$D_L = \mu \times L$$

$$s = D_L + (\sigma \times z \times \sqrt{L})$$

$$S = s + EOQ$$

$$Q = S - I$$

The second inventory policy, **Policy 8**, made use of the ABC inventory control, described in § 4.6.1, to assign different z-scores to products based on their category classification. This model is expected to perform slightly better than the previous policy in terms of meeting demand, but increases the risk of overstocking.

$$z = \begin{cases} 1.96 & ,\alpha = 97.50\% \quad \text{(C-category)} \\ 2.33 & ,\alpha = 99.01\% \quad \text{(B-category)} \\ 3.09 & ,\alpha = 99.90\% \quad \text{(A-category)} \end{cases}$$

$$D_L = \mu \times L$$

$$s = D_L + (\sigma \times z \times \sqrt{L})$$

$$S = s + EOQ$$

$$Q = S - I$$

The third inventory policy conceptualizes how to optimally stock medical dispensing machines based on product criticality values (1–3), and do not focus on the placing of orders that will fully satisfy the facility's demand. This policy will not be considered.

The fourth inventory policy, **Policy 9**, achieved the lowest average cost and stock-out values when tested by Gebicki *et* al. [32]. This makes it the most promising inventory policy proposed by the paper. This model fixed the z-score values for all products to 3.09 (99.9% service level) to optimise the service level. Additionally, Gebicki *et* al. introduce the expected number of items to expire during the lead time (E_L) and the expected orders to arrive during the lead time (O_L) into the model.

$$z = 3.09, \quad \alpha = 99.90\%$$

$$D_L = \mu \times L$$

$$s = D_L + (\sigma \times z \times \sqrt{L})$$

$$I_L = I - s - E_L$$

$$q = MAX(D_L - O_L - I, 0)$$

$$S = s + q$$

$$Q = \begin{cases} q + (S - I_L) & \text{,if } I_L < s \\ q & \text{,else} \end{cases}$$

4.6.5 Inventory policy model described by Jensen and Bard

In order to get a comparative model for the (s,Q) model described in Policy 4, the 'general solution' for the (s,Q) model described by Jensen and Bard [38] will be considered.

Policy 10 – (s,Q) by Jensen and Bard

This inventory model uses a z-score value (normal distribution) to determine the safety stock value (SS). The safe stock increases the reorder point (s) which begins as the expected demand to occur over the lead time (D_L). The actual model uses inventory cost variables to calculate an appropriate z-score, however to ensure best performance the z-score will be set to 2.33, resulting in a 99.01% service level (α).

$$D_L = \mu \times L$$

$$Q = EOQ$$

$$SS = z \times \sigma$$

$$s = D_L + SS$$

4.6.6 Inventory policy models described by Chopra and Meindl

Chopra and Meindl [12] present two universal inventory models used for supply chains, one continuous review model and one periodic review model. A cycle service level (CSL) is used to calculate the safety stock value (SS). The safety stock is determined using the inverse of the cumulative normal distribution function, $F^{-1}(CSL, D_{L+R}, \sigma_{L+R})$. The programming code to calculate this is shown in § B.2, Appendix B. The CSL is the probability that no stock-out will occur during the period between two successive replenishment deliveries. For testing these models, the CSL will be set to 99.0% for optimum results.

Policy 11 – (s,Q) Continuous Review Model by Chopra and Meindl

The continuous review policy described by Chopra and Meindl only requires the lead time (L), daily demand average (μ) and standard deviation (σ) to compute a solution. The order quantity (Q) is set to the expected demand to occur over the lead time (D_L).

$$CSL = 99.0\%$$

$$D_L = \mu \times L$$

$$\sigma_L = \sigma\sqrt{L}$$

$$Q = D_L$$

$$s = D_L + SS$$

$$SS = F^{-1}(CSL, D_L, \sigma_L) - D_L$$

Policy 12 – (R,s,Q) Periodic Review Model by Chopra and Meindl

Chopra and Meindl do not use a reorder point for the periodic review model. Orders get placed at each review unless the current inventory level (I) is larger than the Par level (S). Unlike the continuous model, which only requires enough inventory over the lead time period (L), the periodic model has to store enough stock for both the lead time and review period ($L + R$).

$$CSL = 99.0\%$$

$$D_{L+R} = \mu \times (L + R)$$

$$\sigma_{L+R} = \sigma\sqrt{L+R}$$

$$S = \mu + SS$$

$$Q = \begin{cases} S - I & \text{,if } I < S \\ 0 & \text{,else} \end{cases}$$

$$SS = F^{-1}(CSL, D_{L+R}, \sigma_{L+R}) - D_{L+R}$$

Expected Shortage per replenishment Cycle

Chopra and Meindl provide a means of predicting the expected average unmet demand to occur during a replenishment period when using their models. This value, the *Expected Shortage per replenishment Cycle* (ESC), is calculated by using the cumulative normal distribution function, $F_s\left(\frac{SS}{\sigma_L}\right)$, and the normal distribution function, $f_s\left(\frac{SS}{\sigma_L}\right)$, where σ_L and SS were calculated during either Policy 11 or 12.

$$\begin{aligned} ESC &= -SS\left[1 - F_s\left(\frac{SS}{\sigma_L}\right)\right] + \sigma_L\left[f_s\left(\frac{SS}{\sigma_L}\right)\right] \\ \text{[Excel]} &= -SS\left[1 - NORM.DIST\left(\frac{SS}{\sigma_L}, 0, 1, 1\right)\right] + \sigma_L\left[NORM.DIST\left(\frac{SS}{\sigma_L}, 0, 1, 0\right)\right] \end{aligned}$$

Lead time uncertainty

Real-world deliveries do not always behave systematically. Lead times may be longer or shorter than contractually agreed upon. For scenarios which show a lot of variance in lead times, Chopra and Meindl suggest using the following equations for their two inventory models:

$$L_{ave} = \text{average of historic lead times}$$

$$s_L = \text{standard deviation of historic lead times}$$

Continuous Review Model	**Periodic Review Model**
$D_L = \mu \times L_{ave}$	$D_{L+R} = \mu \times (L_{ave} + R)$
$\sigma_L = \sqrt{L_{ave}\sigma^2 + \mu^2 s_L}$	$\sigma_{L+R} = \sigma\sqrt{L_{ave} + R}$

4.6.7 Remarks

All twelve inventory policies, with exception to Policy 5, will be tested in Chapter 6 and used to compare with the new inventory model(s) that will be designed. The only inventory policy that considered expiring items or expected orders to arrive during the lead time of the current order was Policy 9. These events can have a significant impact on the behaviour of inventory levels and should, in some form, be implemented into the final inventory model for this paper.

For almost every inventory policy, the reorder point value (s) was the main contributor for implementing the facility's lead time and/or review period into the formula. This is because orders need to be issued a sufficient amount of time before a stock-out occurs, allowing the delivery at least the agreed upon lead time to arrive. This works for continuous review systems, which can place orders the moment inventory levels reach the reorder point. It should even perform acceptably for short review periods (1–3 days). However, the long review periods experienced in South Africa's healthcare supply chain may cause these inventory policies to fail.

The economic order quantity (EOQ) was used in seven of the twelve inventory policies. However, this order quantity does not take the lead time, nor review period, into account. In fact, the EOQ assumes that the number of orders placed during the year are flexible and is calculated after the order quantity using Equation 4.3 (page 36). The EOQ is not designed to scale with lead times and review periods. This makes the EOQ a fixed value, which may cause over- or understocking.

The inventory policies most often attempt to estimate the total future demand by multiplying the average demand (μ) by the necessary number of days. This was expressed by the equations shown below.

$$D_L = \mu \times L$$

$$D_R = \mu \times R$$

$$D_{L+R} = \mu \times (L + R)$$

This is evidently the application of the Moving Average forecast method. There exists many other forecast methods used to predict future events. Better predictions of upcoming demand can improve the performance and accuracy of an inventory policy. Therefore, an investigation will be conducted on the use of other popular forecast methods in Chapter 5. The optimal forcast method will be used in the design of any new inventory policy model during Chapter 6. Chapter 6 will begin by testing the inventory policies found throughout the course of this chapter.

4.7 Project Progress

Objective III was achieved during this chapter. All eight targets in the *Investigation phase* of the project methodology framework have now been accomplished. Figure 4.11 shows the updated framework.

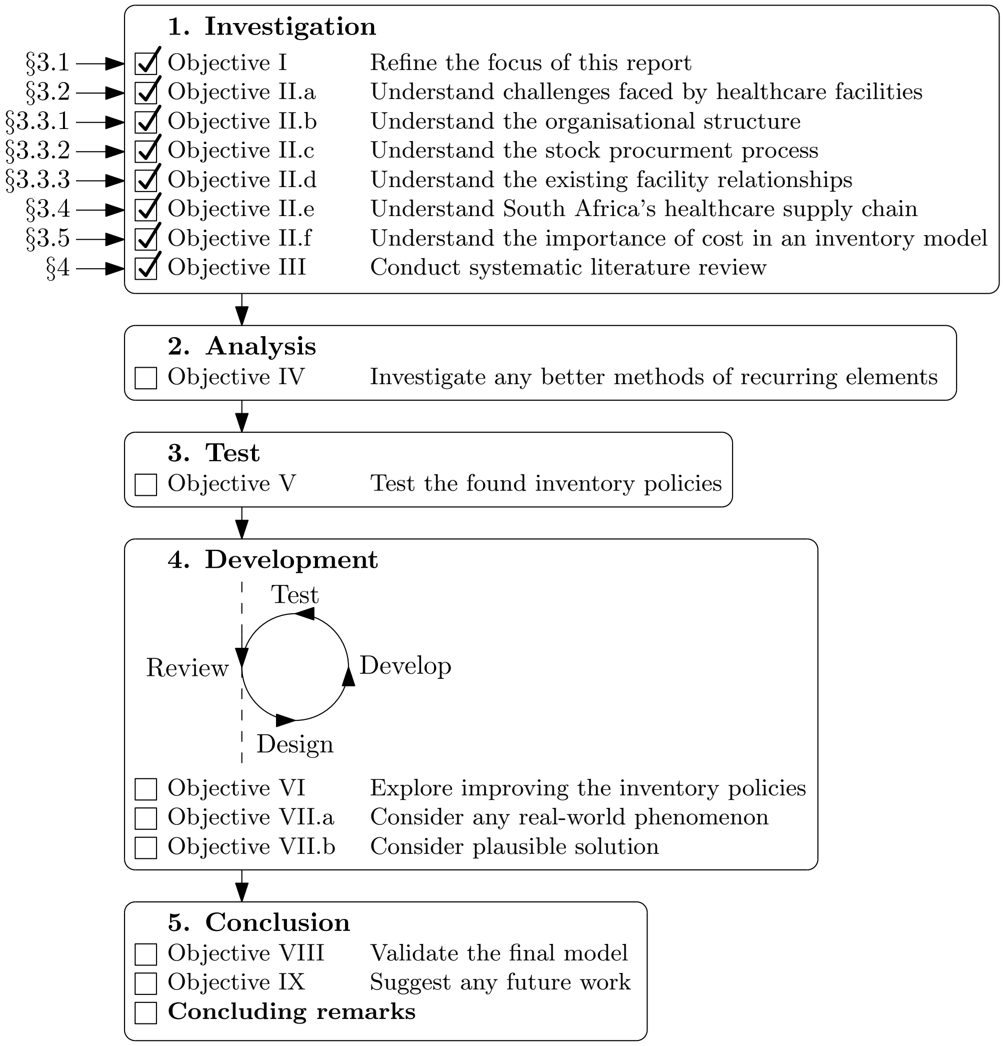

FIGURE 4.11: *Project methodology framework: Chapter 4.*

CHAPTER 5

Forecasting

Contents

This chapter will perform the *Test phase* of the project methodology framework, § 2.2, addressing Objective IV from the project objectives list, § 1.4.

The inventory policies discovered in § 4.6 use of the average daily demand to predict the expected future demand which may result in inaccurate predictions. In the case of an increasing demand trend the predicted demand would be too low. This means there is understocking, which could result in stock-outs. In the case of a decreasing demand trend the predicted demand would be too large, which may result in overstocking.

5.1 Forecast Methods

Forecasting is the means of predicting future occurrences with the help of existing historic data [**Chauhan2016**]. This section will investigate ten popular forecasting techniques in order to determine the best method capable of producing accurate demand predictions for inventory policy models.

5.1.1 Naive Approach Method (NA)

The *Naive Approach* (NA) received its name from the lack of sophistication of this method. Best used to forecast steady data sets, this model assigns the very last available historic data value from the learning set as the forecast value for all the future n periods:

$$\hat{y}_{t+1} = y_t$$

where y_t is the most recent historic data value and $\hat{y}_{t+1}$ is the forecast value [77].

5.1.2 Simple Average Method (SA)

The *Simple Average* (SA) forecast method is designed to be used for demand sets where, despite small fluctuations in data values, the average of the data points remain fairly consistent throughout the entire historic data set. The forecast value for all future n periods are equal to the mean of the *entire* historic data set [58]:

$$\hat{y}_{t+1} = \frac{1}{x}\sum_{j=0}^{x-1} y_{t-j}$$

where the historic data set has x many data points [77].

5.1.3 Moving Average Method (MA)

The *Moving Average* (MA) forecast method uses a smaller, finite number of the *most recent* historic data points to determine the average rather than the entire historic data set. The MA method is good at adapting to recent changes in the demand set's behaviour [7]:

$$\hat{y}_{t+1} = \frac{1}{p}\sum_{j=0}^{p-1} y_{t-j}$$

where p is the "window size" (finite number of previous values) [77].

5.1.4 Weighted Moving Average Method (WMA)

The *Weighted Moving Average* (WMA) is an adaptation of the MA method. In this model each historic data point is weighted based on its recency. The more recent the data point, the larger the weight. The summation of the weights must always be 100% [77]:

$$\hat{y}_{t+1} = \sum_{j=0}^{p-1} w_j y_{t-j}$$

$$w_j = \left(\frac{p-j}{\sum_{k=0}^{p-1}(p-k)}\right), \quad \text{for all } j = \{0, \ldots, p-1\} \tag{5.1}$$

where w_j is the weight assigned to the historic data value based on the window size p.

5.1.5 Simple Exponential Smoothing Method (SES)

The *Simple Exponential Smoothing* (SES) also makes use of uses weights to produce a forecast. These weights are controlled by a constant called the smoothing parameter, $\alpha \in (0; 1)$. However, unlike the WMA method which only considers the most recent historic data, the SES method looks at the entire historic data set:

$$\hat{y}_{t+1} = \alpha \sum_{j=0}^{n-1} (1-\alpha)^j y_{t-j}$$

where n is the number of points in the historic data set. This equation can be adjusted to a less computational formula after having calculated the first forecast normally. The succeeding forecast values can be calculated using [77]:

$$\hat{y}_{t+1} = \alpha y_t + (1-\alpha)\hat{y}_t$$

5.1.6 Moving Simple Exponential Smoothing Method (MSES)

The SES method uses the entire historic data set to predict a forecast. For the sake of interest, the SES will be also be applied to a window size p, similarly to the MA method and for this reason will be classified as the *Moving Simple Exponential Smoothing* (MSES):

$$\hat{y}_{t+1} = \alpha \sum_{j=0}^{p-1} (1-\alpha)^j y_{t-j}$$

and thereafter the same formula described in the SES method may be used to reduce computational time:

$$\hat{y}_{t+1} = \alpha y_t + (1-\alpha)\hat{y}_t$$

5.1.7 Holt's Linear Trend Method (HLT)

Holt's Linear Trend Method (HLT) is an adaptation of the SES method and makes use of trend prediction. The forecast value is calculated using Equation 5.2, which is the summation of the *Level* of the series and the *Trend* estimates at time t, Equations 5.3 and 5.4 respectively, calculated from the historic data set [35, 77]:

$$\text{Forecast:} \quad \hat{y}_{t+h|t} = l_t + hb_t \tag{5.2}$$

$$\text{Level:} \quad l_t = \alpha y_t + (1-\alpha)(l_{t-1} + b_{t-1}) \tag{5.3}$$

$$\text{Trend:} \quad b_t = \beta(l_t - l_{t-1}) + (1-\beta)b_{t-1} \tag{5.4}$$

where $\alpha \in [0,1]$ and $\beta \in [0,1]$ are smoothing parameters optimized to fit the forecast to the historic data set and h is the step forward in the forecast. Table 5.1 is an example of the application of this method:

TABLE 5.1: *Holt's Linear Trend example.*

t	y_t	l_t	b_t	$\hat{y}_{t\|t-1}$	h	$\hat{y}_{t+h\|t}$
⋮	⋮	⋮	⋮	⋮	⋮	⋮
24	68.20	68.21	2.102	68.25	–	–
25	28.12	68.49	2.102	70.31	–	–
26	69.79	69.92	2.102	70.60	–	–
27	72.60	72.50	2.102	72.02	–	–
–	–	–	–	–	1	74.60
–	–	–	–	–	2	76.70
–	–	–	–	–	3	78.80

5.1.8 Damped Trend Method (DT)

The *Damped Trend Method* (DT) was created in 1985 by Gardner and McKenzie which uses a damping parameter, $\phi \in (0,1)$, to skew the results from HLT method to better-fit data that have trends with an odd curvature [30]. There have been several adaptations of this method since 1985, however none of these have improved the outcome of the forecasts [31]. For this reason, only the 1985 version will be considered [31, 35]:

$$\text{Forecast:} \quad \hat{y}_{t+h} = l_t + (\phi + \phi^2 + ... + \phi^h)b_t$$

$$\text{Level: } l_t = \alpha y_t + (1-\alpha)(l_{t-1} + \phi b_{t-1})$$

$$\text{Trend: } b_t = \beta(l_t - l_{t-1}) + (1-\beta)\phi b_{t-1}$$

where, if $\phi = 1$ it is identical to the HLT method, and if $\phi < 1$ the trend approaches a constant some time in the future [60].

5.1.9 Holt-Winter Additive Method (HWA)

Data values can fluctuate for many reasons, but often follow a pattern. For example, shops may experience higher demands at the end of each month when customers receive their monthly pay. These systematic fluctuations are referred to as the *seasonality* of the data. The *Holt-Winter Additive Method* (HWA) is a modification of HLT method [35, 77]:

$$\text{Forecast: } \hat{y}_{t+h} = l_t + hb_t + s_{t-m+h}$$

$$\text{Level: } l_t = \alpha(y_t - s_{t-m}) + (1-\alpha)(l_{t-1} + b_{t-1})$$

$$\text{Trend: } b_t = \beta(l_t - l_{t-1}) + (1-\beta)b_{t-1}$$

$$\text{Seasonal: } s_t = \gamma(y_t - l_{t-1} - b_{t-1}) + (1-\gamma)s_{t-m}$$

where s_t is the seasonality, using a new smoothing parameter $\gamma \in [0, (1-\alpha)]$, for controlling the fluctuations of the forecast.

5.1.10 Holt-Winter Multiplicative Method (HWM)

The *Holt-Winter Multiplicative Method* (HWM) is another take on adding seasonality to the HLT method, similarly to the before-mentioned HWA method. The equations for this method are described as follows [35]:

$$\text{Forecast: } \hat{y}_{t+h} = (l_t + hb_t)s_{t-m+h}$$

$$\text{Level: } l_t = \alpha\left(\frac{y_t}{s_{t-m}}\right) + (1-\alpha)(l_{t-1} + b_{t-1})$$

$$\text{Trend: } b_t = \beta(l_t - l_{t-1}) + (1-\beta)b_{t-1}$$

$$\text{Seasonal: } s_t = \gamma\left(\frac{y_t}{(l_{t-1} + b_{t-1})}\right) + (1-\gamma)s_{t-m}$$

5.2 Testing the forecasts

The ten forecast methods described in § 5.1 were tested for their ability to produce accurate forecasts of demands likely to be encountered at healthcare facilities.

5.2.1 Accuracy measures for testing the forecasts

Three accuracy measures will be used to compare the different forecast methods, namely the *Root-Mean-Square Error* ($RMSE$), *Mean-Absolute-Error* (MAE) and the *Bias* ($Bias$):

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(\hat{y}_t - a_t)^2}, \quad \text{smaller is better} \in [0, \infty]$$

$$MAE = \frac{1}{n}\sum_{t=1}^{n}|\hat{y}_t - a_t|, \quad \text{smaller is better} \in [0, \infty]$$

$$Bias = \sum_{t=1}^{n}(\hat{y}_t - a_t), \quad = 0.0 \text{ (perfect)}, > 0.0 \text{ (overshoot)}, < 0.0 \text{ (undershoot)}$$

where n is the number of future periods being predicted, $\hat{y}_t$ is the forecast demand of period t, a_t is the actual demand at period t and $\bar{a}$ is the average actual demand across the n periods [88, 83]. R-squared values below 70% indicate a poor ability to follow historic data, while values above 85% indicate a strong ability.

5.2.2 Demand sets for testing the forecasts

Three demand sets were used to test each forecast model's ability to predict relatively "flat" demand, "increasing" demand and "decreasing" demand. The three demand sets are shown in Figure 5.1 and consist of 365 days (1 year) historic demand values ($t = \{-1, \ldots, -365\}$) to be used as the 'learning set' for the forecast method to fit to, and 45 days known future values ($t = \{0, \ldots, 44\}$) as the 'test set' to inspect accuracy. Forty-five days were selected to represent a healthcare facility with a 30 day review period, waiting a lead time of 14 days and also predicting the demand of the current day; $30 + 14 + 1 = 45$ days.

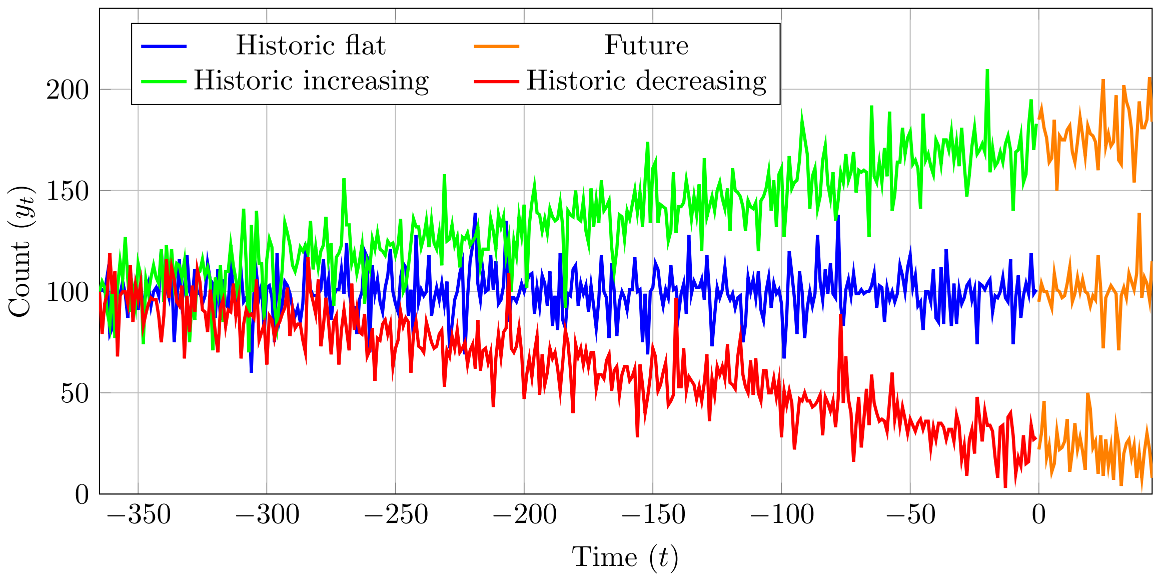

FIGURE 5.1: *The demand sets to be used for testing the forecast methods described in this section.*

5.2.3 Test results

The ten forecast methods were tested using the 'Flat Trend', 'Increasing Trend' and 'Decreasing Trend' demand sets described in § 5.2.2. Each forecast was developed in Python, making use of popular existing packages when available. The portion of code responsible for creating each forecast is provided in § B.1, Appendix B. The results from each forecast method were tested, as planned, using the accuracy measures from § 5.2.1. The results are shown in Table 5.2.

TABLE 5.2: *Forecast results: RMSE & MAE, smaller is better. Bias, the nearer to 0.00 the better.*

Forecast	Flat Trend			Increasing Trend			Decreasing Trend		
	RMSE	*MAE*	*Bias*	*RMSE*	*MAE*	*Bias*	*RMSE*	*MAE*	*Bias*
NA	10.03	6.18	2	12.91	10.13	204	11.58	9.29	252
SA	10.08	6.07	-43	44.15	42.47	-1911	43.79	42.60	1917
MA	10.08	6.07	-43	14.76	11.93	-381	10.75	8.36	162
WMA	10.08	6.07	-43	14.21	11.38	-336	10.75	8.36	162
SES	10.08	6.07	-43	13.71	10.82	-291	10.46	8.02	117
MSES	10.08	6.08	-43	14.76	11.93	-381	10.75	8.36	162
HLT	10.03	6.18	2	11.73	8.84	-30	12.03	9.62	-315
DT	10.08	6.07	-43	13.71	10.82	-291	10.46	8.02	117
HWA	10.27	6.36	-22	11.90	8.93	-22	10.69	8.31	62
HWM	13.36	10.09	370	12.30	9.64	-34	28.93	24.16	1025

When inspecting the *RMSE* results, the HWM and SA methods show the worst performance, while the other methods appear to perform very similarly. However, the HWA and HLT method show the best overall results, showing the greatest potential for predicting future demand values in the inventory policies.

5.2.4 Concluding with an acceptable forecast method

Both forecasts, HWA and HLT, have proven to produce good accuracy in their predictions. In order to determine which of the two forecast methods will be used the forecasts were again tested using the same three demand sets as before ('flat', 'increasing' and 'decreasing'). The total demand of the period was captured, as well as the computational run time on the computer, using the same Python script shown in § B.1, Appendix B. The total predicted demand should be as close to the actual demand as possible. A smaller computational time is desirable to produce good efficiency for systems with a large product database. Table 5.3 shows the results.

TABLE 5.3: *Forecast results: HWA and HLT.*

Demand set	Tests	Actual	HWA	HLT
Flat	Total predicted demand	4 543	4 521	4 545
	Computational time [sec]	–	2.103	0.480
Increasing	Total predicted demand	8 031	8 002	8 001
	Computational time [sec]	–	4.169	1.153
Decreasing	Total predicted demand	1 008	731	693
	Computational time [sec]	–	2.370	0.4992

Both the HWA and HLT method performed well at predicting the total demand of the 'flat' and 'increasing' demand sets and both experienced some trouble with predicting the total demand for the decreasing demand set. The HLT method computed its results roughly four times faster than the HWA method. Lengthy computational times can hinder the work rate of physicians and is directly in violation of point 2 of the key focus efforts (§ 3.6), to protect *time availability* required for treating patients. For this reason the HLT method will be used in future models to upcoming demand.

5.3 Project Progress

Objective IV was achieved during this chapter. This completes the *Analysis phase* of the project methodology framework and has been updated in Figure 5.2.

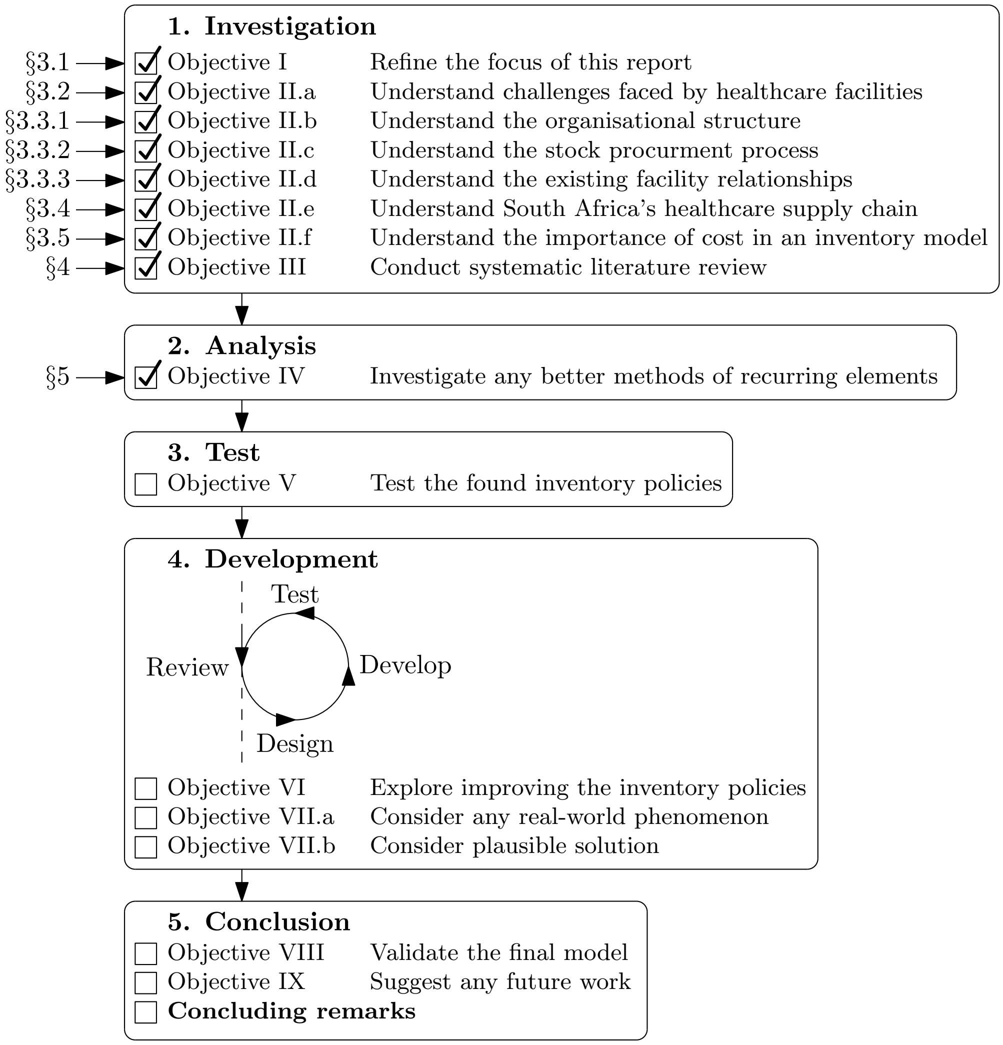

FIGURE 5.2: *Project methodology framework: Chapter 5.*

CHAPTER 6

Creating an Order Policy

Contents

This chapter will address the Test phase and the first target in the *Development phase* of the project methodology framework, § 2.2. This will achieve Objectives V and VI from the project objectives list, § 1.4.

6.1 Testing the Inventory Policy Models

Eleven inventory policies were discovered in § 4.6. This section will test these policies in order to determine their effectiveness in the South African healthcare sector and use the results for a comparative assessment of any new inventory policy models developed in this section.

6.1.1 Creating the Test

In order to ensure fair testing between the inventory policy models one large simulation will be created to capture and compare results. This test will pretend to represent a small public healthcare facility with the same 100 products created during the *Parameters* section of § 3.5.2. All products will begin with zero items in inventory (stocked-out). The sample products list is shown in Tables C.1–C.3, Appendix C. The first 100 days of the simulation will allow the inventory policy to enter the system and establish its own order routine. Thereafter a further 365 days (1 year) will be simulated. The results from this one year will be captured and compared.

Facility warehouses are limited to the amount of stock they are able to keep on hand at any given moment. Over-ordering can cause inventory levels to exceed the warehouse storage capacity and result in unstructured stock filing. The arrival of stock orders should not cause inventory levels to exceed the available storage space [63]. However, because there is no obtainable data set of product storage measures, all 100 products will receive the same unit of storage space (value of 1). Further, the test itself will not define a limit to the facility's storage space. This will allow equal opportunity for each inventory policy to perform at its best with regards to minimizing the unmet demand (maximize meeting the demand). However, the maximum inventory level to occur during the 365 day simulation period will be captured and reviewed. The simulation will be repeated for each review period (R) from the predetermined set $\mathcal{R} = \{1; 7; 30\}$ (§ 3.4.4).

South African healthcare facilities more frequently issue orders on weekly and monthly basis's, thus the simulation will be conducted in the form of a period review system. The EOQ (§ 4.6.2) does not scale in accordance with the lead time of a product, nor the review period of a facility. This means that the order quantity used for a daily review period would be exactly the same as an order placed one a month. Normally the review period would be altered to optimise the use of the EOQ, but for this simulation the following equation will be used to replace the EOQ: $Q = (\mu \times T_{max}) - I$, where μ is the average daily demand, I is the current inventory level and $T_{max} = L + R$ for each product.

The Par value in Policy 1 was set to $S = \mu \times L$. Policy 4 was modeled using the reorder point $s = \mu + 3\sigma$, causing the largest daily demand standard deviation (σ) increase, to obtain the lowest stock-out probability. The days on hand value for Policy 6 was set to $DOH = L + R$.

Policy 8 uses the ABC inventory control to assign z-scores. The ABC inventory control was thus applied to the 100 products by reviewing the $Cost\ per\ item = Price/MOQ$. The category of each product (A, B or C) is indicated in the product list, Appendix C. The number of products and the resulting cost percentage for each category are shown in Table 6.1, along with the z-scores for Policy 8.

TABLE 6.1: *Testing Policy 8: z-values assigned to each category.*

Category	#Products	%Costs	z-value	Confidence
A	11	68.61%	3.09	99.90%
B	15	21.26%	2.33	99.01%
C	69	10.13%	1.96	97.50%

Policy 9 uses a measure E_L which is the expected number of expired items during the lead time period. Expiries should be eliminated by proper use of FIFO unless there is over-ordering, therefore in this test $E_L = 0$.

6.1.2 SLR Inventory Policies' results

The final 365 simulation days for the systematic literature review (SLR) inventory policies are shown in Tables 6.2–6.4. These results include the total unmet demand [items] (UD), the total number of days which experienced any stock-out [days] (D_{SO}), the maximum inventory level to occur [items] ($I_{\max}$) and the total number of expired items (TE) during the simulated year.

TABLE 6.2: *SLR inventory policy simulation results: $R = 1$.*

	$R = 1$			
Policy	UD	D_{SO}	I_{Max}	TE
Policy 1	57 013	243	967 508	393 661
Policy 2	152 995	303	1 111 562	481 945
Policy 3	59 067	217	1 894 042	1 191 418
Policy 4	905 739	365	148 275	0
Policy 6	166 194	313	829 326	221 642
Policy 7	60 608	221	1 891 538	1 185 277
Policy 8	60 382	217	1 893 121	1 187 423
Policy 9	56 742	220	2 352 284	1 203 458
Policy 10	51 859	237	972 994	439 611
Policy 11	62 547	219	1 333 628	530 640
Policy 12	869 332	365	153 813	0

TABLE 6.3: *SLR inventory policy simulation results: $R = 7$.*

	$R = 7$			
Policy	UD	D_{SO}	I_{Max}	TE
Policy 1	565 904	342	106 808	0
Policy 2	519 495	364	182 358	442
Policy 3	206 062	312	291 125	22 634
Policy 4	1 887 956	365	41 868	0
Policy 6	367 971	361	160 101	0
Policy 7	211 858	316	284 540	22 804
Policy 8	284 017	317	284 017	22 804
Policy 9	169 833	324	282 561	78 690
Policy 10	281 763	357	158 597	0
Policy 11	192 583	339	166 729	943
Policy 12	2 420 580	365	19 253	0

Policies 4 and 12 may have experienced zero expiries in all three review period scenarios, however they resulted in the highest unmet demand levels. This indicates that there was a significant amount of understocking. All eleven policies performed poorly, causing 59–100% of the year to experience stock-outs.

6.1.3 Concluding remarks

A more detailed model is necessary capable of considering the stock on-route, and accurately predicting the future demand and expiries which may occur. This concept was highlighted at the beginning of § 4.6.7, when it was noticed that only Policy 9 of the SLR inventory policies attempted to include these events.

TABLE 6.4: *SLR inventory policy simulation results: $R = 30$.*

	$R = 30$			
Policy	UD	D_{SO}	I_{Max}	TE
Policy 1	1 640 731	365	104 987	0
Policy 2	717 533	282	346 472	0
Policy 3	724 615	363	354 630	6 598
Policy 4	1 665 459	365	117 426	0
Policy 6	348 172	305	327 035	0
Policy 7	737 441	364	357 134	6 507
Policy 8	733 460	364	349 569	6 538
Policy 9	678 638	354	373 258	13 880
Policy 10	816 114	310	352 657	0
Policy 11	1 607 912	365	106 308	0
Policy 12	2 698 672	365	22 550	0

6.2 The Iterative Forecast Inventory Model

This section will attempt to design a new inventory model capable of out-performing the eleven policies tested in § 6.1. The focus of the model is to eliminate any unmet demand experienced throughout the entire simulated year (365 days) while maintaining acceptable inventory levels.

6.2.1 Building the model

In the same manner that the other inventory policies were tested (described in § 6.1.1), assume that sufficient storage space is readily available to hold inventory. Chapter 5 investigated the best method for predicting the future demand, one review period plus lead time ($R + L$ days) from the current date. Holt's Linear Trend (HLT), page 47, proved to deliver both good predictions and computational times in § 5.2.4 and will be used for this model. In order to assist the forecast with achieving correct demand levels the predicted values will be increased by some safety stock percentage (ss).

The *Iterative Forecast Inventory Model*, described in Figure 6.1, steps through the future $R+L$ days, starting from the current day (d), and simulates the expected behaviour of each day. The objective is to determine the required amount to fulfil the total expected demand (F_t) between one lead time from now ($d + L$) and the next review's order arrival ($d + R + L$). At the start of each time step (t) the earliest remaining order batch ($O_{t,m}$) is reviewed. If the day of order (m) is older than one lead time (L) plus shelf life (SL), then the order is expected to have expired. All expired quantities are logged and the next remaining order batch is reviewed.

If the order batch has not yet expired then the day's remaining unmet demand (r), which has been increased by some safety stock percentage (ss), is subtracted from the order batch. Any remaining unmet demand for that time step is subtracted from the next earliest order batch. This continues until $r = 0$. If there are no remaining order batches, $r > 0$ and the time step has progressed by at least one lead time (when today's order would arrive), then the predicted unmet demand (q) is increased. At the end of the simulation, q is the necessary quantity required to satisfy demand until the next review's order arrives. Once all $L + R$ simulation time steps have completed, the order quantity (Q) is calculated using q and the product's minimum order quantity (MOQ).

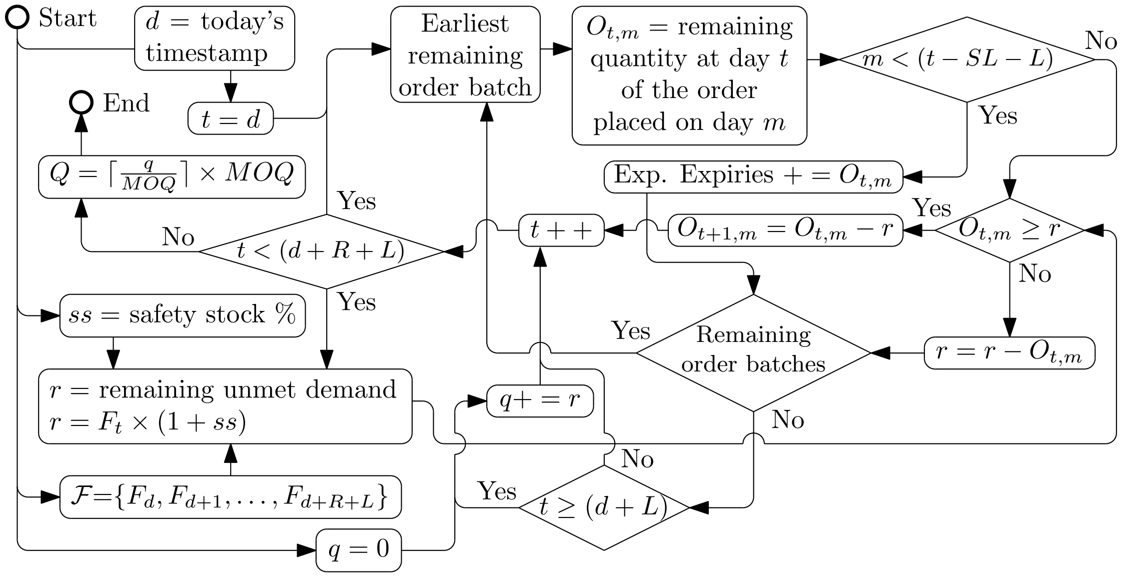

FIGURE 6.1: *Diagram depicting the Iterative Forecast Inventory Model.*

6.2.2 Testing the model

The test was conducted in the same manner as for the SLR policies during § 6.1. The tests were repeated using safety stock percentages $ss \in \{10\%, 20\%, \ldots, 90\%\}$ and the results are shown in Table 6.5. The only simulations which achieved zero stock-outs and no expired items were at $R = 1$ and $ss \geq 40\%$. The results from both the $R = 7$ and $R = 30$ simulations failed to completely eliminate any unmet demand. Very large safety stock percentages were required in order to improve results.

TABLE 6.5: *Iterative forecast model simulation results*

ss	10%	20%	30%	40%	50%	60%	70%	80%	90%
$R = 1$									
UD	103308	20266	1275	0	0	0	0	0	0
D_{SO}	365	284	35	0	0	0	0	0	0
I_{Max}	22078	34048	49688	63954	79021	92973	107427	120527	134940
TE	0	0	0	0	0	0	0	0	0
$R = 7$									
UD	149984	78148	36041	16822	7708	3226	1674	641	105
D_{SO}	319	279	207	137	73	37	20	10	3
I_{Max}	82879	99641	121817	142495	164295	184519	206669	229572	251273
TE	0	0	0	0	0	0	0	0	0
$R = 30$									
UD	232925	177051	133671	96806	67751	44550	30619	23419	18565
D_{SO}	348	330	321	300	279	243	208	187	163
I_{Max}	328361	373613	420124	466851	513252	558901	608851	658781	707408
TE	474	1324	2212	3921	6883	11996	19166	28810	39415

6.2.3 Concluding remarks

The *Iterative Forecast Inventory Model* yielded much better results than the previously tested SLR policies. However, this model was not capable of entirely eliminating unmet demand levels for weekly and monthly review periods. Further adaptations of this model will be necessary in order to find an inventory model capable of adapting to any review period. The fact that such large safety stock percentages were necessary to meet demand hint that the predicted demand values from the forecast were off to some degree. The next model will attempt to understand and rectify this problem.

6.3 The HLT & ND Inventory Model

Forecasts can sometimes over-assume the behaviour of real-world demand. The HLT method uses historic data to predict the trend of a demand set and calculate future values. In § 5.2.4, both the HLT and HWA forecasts showed some struggle with decreasing historic data trends and under-assumed the total predicted demand (recall Table 5.3). The more time periods that need to predicted, the more smaller estimated demand becomes. The *The HLT & ND Inventory Model* will attempt to prevent these occurrences.

6.3.1 Building the model

In order to ensure a more confident *total demand* value the *Normal Distribution* (ND) will be used to define a suitable lower bound for the predicted forecast values. Figure 6.2 demonstrates the use of the ND with a 30% service level ($\alpha = 0.3$) for both increasing and decreasing HLT predictions. The final forecast set will be created by taking the maximum value between the original forecast (HLT) and the normally distributed value at some service level (α). This can be expressed mathematically, where $\hat{y}_t$ is the original HLT forecast value for day t and $F^{-1}(\alpha, \mu, \sigma)$ is the inverse of the cumulative normally distributed value at a service level α with a historic mean demand (μ) and standard deviation (σ):

$$F_t = \text{Max}\{\hat{y}_t \; ; \; F^{-1}(\alpha, \mu, \sigma)\}$$

The coding used to perform this is shown in § B.2, Appendix B. The rest of the model will behave the same as the *Iterative Forecast Inventory Model*.

6.3.2 Testing the model

Large safety stock and service level values will result in very large predicted forecasts. This model will only be simulated all combinations of $ss \in \{10\%, 20\%, \ldots, 50\%\}$ and $\alpha \in \{10\%, 20\%, \ldots, 90\%\}$. The results for these simulations are shown in Tables 6.6–6.11. This model achieved good results for all three review periods. The safety stock percentage had the greatest effect on reducing the unmet demand. It can be said that all three models will perform well where $\alpha \geq 10\%$ and $ss \geq 30\%$, based on the simulation demand sets.

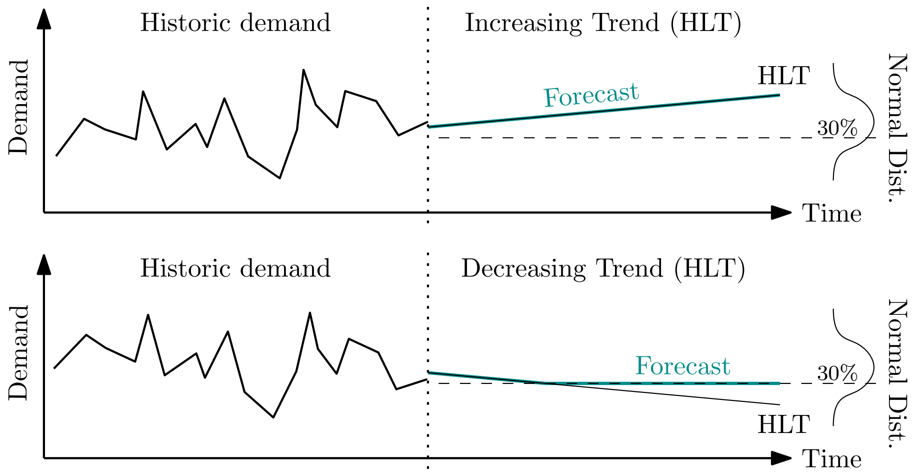

FIGURE 6.2: *Example: generating a forecast with the Normal Distribution of 30%.*

TABLE 6.6: *$R = 1$: Unmet demand and #days with a stock-out.*

$R = 1$										
	Unmet Demand (UD)					Days with a stock-out (D_{SO})				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	52 879	2 069	0	0	0	364	94	0	0	0
20%	23 054	64	0	0	0	350	6	0	0	0
30%	6 541	0	0	0	0	244	0	0	0	0
40%	1 106	0	0	0	0	94	0	0	0	0
50%	296	0	0	0	0	30	0	0	0	0
60%	162	0	0	0	0	22	0	0	0	0
70%	21	0	0	0	0	7	0	0	0	0
80%	0	0	0	0	0	0	0	0	0	0
90%	0	0	0	0	0	0	0	0	0	0

TABLE 6.8: *$R = 7$: Unmet demand and #days with a stock-out.*

$R = 7$										
	Unmet Demand (UD)					Days with a stock-out (D_{SO})				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	35 296	1 756	15	0	0	126	40	1	0	0
20%	10 790	240	15	0	0	85	9	1	0	0
30%	2 386	225	15	0	0	53	7	1	0	0
40%	577	185	15	0	0	26	5	1	0	0
50%	186	88	15	0	0	10	3	1	0	0
60%	62	14	0	0	0	2	1	0	0	0
70%	0	0	0	0	0	0	0	0	0	0
80%	0	0	0	0	0	0	0	0	0	0
90%	0	0	0	0	0	0	0	0	0	0

TABLE 6.7: $R = 1$: *max inventory level and total expiries.*

$R = 1$										
	Max Inventory (I_{Max})					Total Expiries (TE)				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	22 734	35 686	52 086	65 668	79 807	0	0	0	0	0
20%	22 912	36 841	52 898	66 321	80 356	0	0	0	0	0
30%	24 299	38 439	53 994	66 955	81 112	0	0	0	0	0
40%	26 692	40 184	55 264	67 997	81 817	0	0	0	0	0
50%	28 597	41 742	56 487	69 475	83 459	0	0	0	0	0
60%	31 324	44 893	58 047	72 713	85 842	0	0	0	0	0
70%	33 945	48 193	61 995	76 289	90 188	0	0	0	0	0
80%	38 239	52 249	67 464	81 499	95 164	0	0	0	0	0
90%	44 106	59 243	73 993	88 508	103 070	0	0	0	0	0

TABLE 6.9: $R = 7$: *max inventory level and total expiries.*

$R = 7$										
	Max Inventory (I_{Max})					Total Expiries (TE)				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	85 177	106 377	128 331	147 325	168 915	0	0	0	0	0
20%	87 841	109 083	130 547	149 583	170 859	0	0	0	0	0
30%	90 632	111 834	132 478	151 618	173 292	0	0	0	0	0
40%	93 276	113 735	134 667	153 357	175 743	0	0	0	0	0
50%	95 251	115 258	137 091	155 353	177 420	0	0	0	0	0
60%	98 161	118 507	139 799	159 716	182 088	0	0	0	0	0
70%	102 420	122 368	144 195	164 003	185 974	0	0	0	0	0
80%	108 650	128 164	150 553	170 499	192 936	0	0	0	0	0
90%	114 832	136 280	158 994	180 300	203 448	0	0	0	0	0

TABLE 6.10: $R = 30$: *Unmet demand and #days with a stock-out.*

$R = 30$										
	Unmet Demand (UD)					Days with a stock-out (D_{SO})				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	31 552	1 412	80	0	0	87	31	14	0	0
20%	7 412	0	0	0	0	37	0	0	0	0
30%	1 054	0	0	0	0	13	0	0	0	0
40%	0	0	0	0	0	0	0	0	0	0
50%	0	0	0	0	0	0	0	0	0	0
60%	0	0	0	0	0	0	0	0	0	0
70%	0	0	0	0	0	0	0	0	0	0
80%	0	0	0	0	0	0	0	0	0	0
90%	0	0	0	0	0	0	0	0	0	0

TABLE 6.11: $R = 30$: *max inventory level and total expiries.*

$R = 30$										
	Max Inventory (I_{Max})					Total Expiries (TE)				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	345 601	391 695	442 007	491 883	537 928	484	1 533	2 578	4 467	8 069
20%	349 852	399 402	448 402	496 206	541 876	539	1 573	2 758	4 467	8 269
30%	352 998	402 204	451 636	499 560	545 724	573	1 613	2 758	4 467	8 369
40%	358 620	405 661	455 427	503 474	550 517	773	1 653	2 758	4 467	8 519
50%	363 181	409 099	459 139	506 954	554 212	773	1 653	2 758	4 467	8 569
60%	368 593	415 126	464 562	512 525	561 833	813	1 653	2 758	4 467	8 569
70%	374 914	422 887	470 536	520 255	569 701	813	1 653	2 758	4 467	8 689
80%	382 943	431 763	480 471	530 364	578 929	813	1 653	2 758	4 607	9 760
90%	397 374	444 945	496 174	546 561	595 527	813	1 653	2 858	5 335	12 729

6.3.3 Concluding remarks

The HLT & ND Inventory Model hypothesised that forecasts were occasionally under-assuming the total future demands. It has been observed that implementing a normally distributed lower bound value, $F^{-1}(\alpha, \mu, \sigma)$, is capable eliminating stock-outs for service levels as low as $\alpha = 10\%$. Going forward, the combination of the Holt's Linear Trend method and Normal Distribution will remain the primary means of forecast, $F_t = \text{Max}\{\hat{y}_t\ ;\ F^{-1}(\alpha, \mu, \sigma)\}$.

Although good conditions were already achieved at $\alpha = 10\%$, a noticeable improvement was observed at the lower safety stock values when increasing from a 10% to 20% service level. Selecting a 20% service level will provide more certainty that the model will perform well. Therefore, assuming availability for accurate daily inventory levels, it is recommended to use the HLT & ND Inventory Model with a 30% safety stock (ss) and 20% service level (α) for all three review periods, $R = \{1, 7, 30\}$, as summarised in Table 6.12.

TABLE 6.12: *Summary of the recommended inventory model conditions.*

Inventory Model	Recommended Conditions			
	Review Period (R)	Safety Stock (ss)	Service Level (α)	Historic Inventory Levels
HLT & ND	1	30%	20%	30 days, daily
	7	30%	20%	
	30	30%	20%	

This inventory policy has proven to perform well under ideal conditions, such as testing with a singular supplier per product who delivers on an exact time frame (lead time) without fail. In truth, each product may have several suppliers capable of delivering in different time frames at various mixed rated (costs) and batch sizes. Real world suppliers may differ in reliability, varying in their actual lead times and ability to meet the desired order quantity. The next chapter will attempt to expand the current inventory policy to optimize orders based on these real-world behaviours.

6.4 Project Progress

Objectives V and VI were achieved during this chapter. This completed the *Test phase* and the first target of the *Development phase* of the project methodology framework. The updated framework is shown in Figure 6.3.

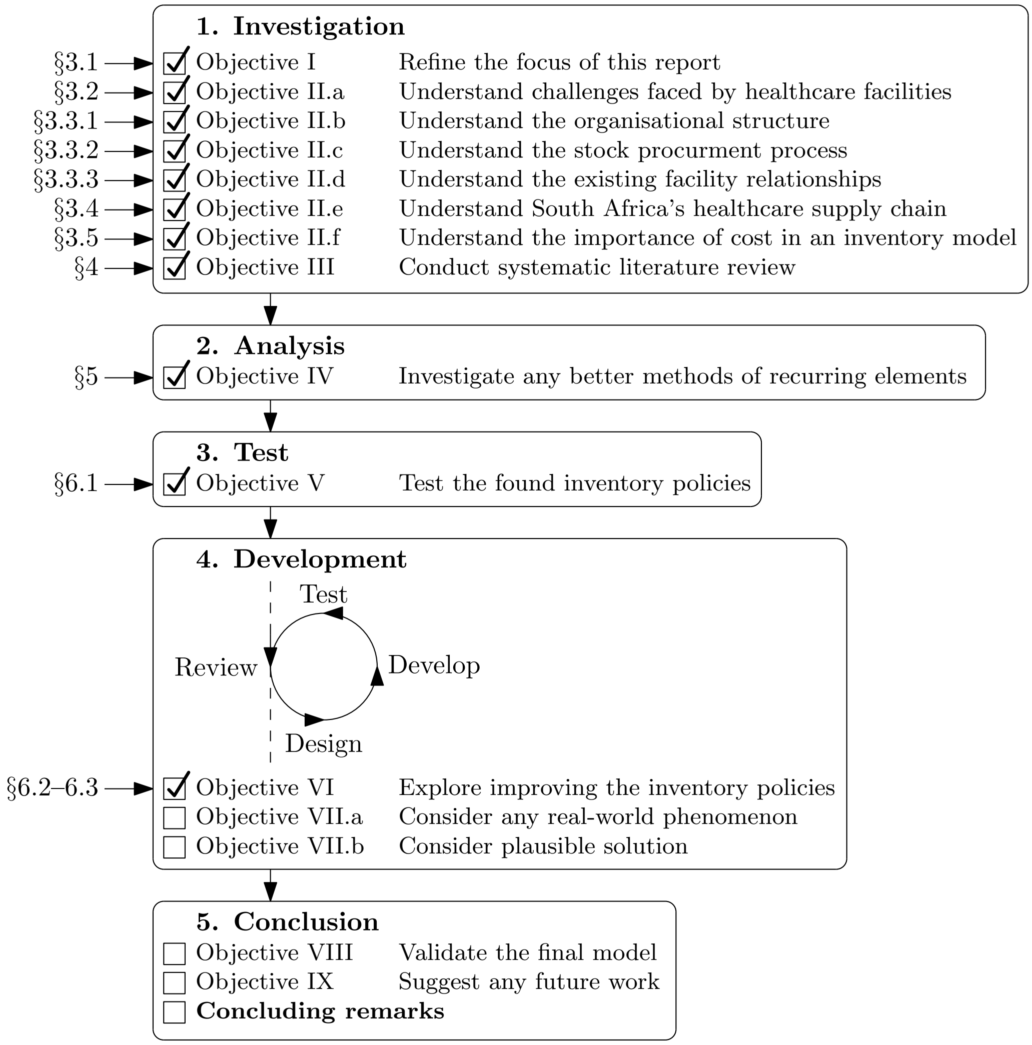

FIGURE 6.3: *Project methodology framework: Chapter 6.*

CHAPTER 7

Improving Order Confidence for Real World behaviour

Contents

A successful inventory policy was created in § 6.3 capable of producing orders suitable for South Africa's difficult lead times and review periods. However, these tests were designed to fairly test each inventory model under ideal conditions in order to observe their behaviour. The chosen inventory policy now needs to be expanded for non-ideal environments. This chapter will attempt to identify and address several real-world supply chain behaviours that may interfere with meeting the necessary demand or incur additional costs. The last two targets in the *Development phase* of the project methodology framework, § 2.2, will be addressed. Theses targets aim to achieve Objective VII from the project objectives list, § 1.4.

7.1 Supplier Prioritisation

Any healthcare facility, when issuing an order, places a level of trust in the supplier(s) to meet the requested quantities and deliver the goods within an agreed upon number of days. Suppliers are limited to what they can produce and the speed at which goods can be transported. This is outside the control of healthcare facilities issuing the orders. However, by reviewing the historic performance of suppliers, orders may be created which could increase the likelihood of acquiring the necessary amount of goods. An example of this is by Chopra and Meindl (§ 4.6.6) whom used historic lead times to improve their inventory policies to accommodate the uncertainty of real-world supplier behaviour.

7.1.1 Varying Lead times

By keeping record of when an order was issued and when it arrived, a historic set of actual lead times ($\mathcal{L}_{p,s}$) can be created for each product p delivered by supplier s. This information is very valuable as it depicts the true behaviour of a supplier with respect to a product. Using the most recent performance to create the historic sets can be used to encourage suppliers to improve their standing against competitors. Some future work may include performing tests to find the optimal number of historic points for the set.

Confidence Interval (C.I.)

One use for this data set is to determine a *confidence interval* (C.I.) for each supplier which will indicate the likely maximum and minimum lead time to occur.

Two popular methods for calculating a C.I. is the *Normal Distribution* (described before in § 6.3) and the *t-Distribution*. The t-Distribution is very popular for data sets which, ideally, should be experiencing repetitive results. The C.I. produced by a t-Distribution is always less than or equal to that of the Normal Distribution, which by comparison spans nearly all data points in the data set for large service levels ($\alpha \in [0, 1]$). The upper and lower confidence values of the t-Distribution C.I. is calculated as follows [3]:

$$\mathcal{L}_{p,s} = \{L_1, L_2, \ldots, L_m\}, \quad m = \text{number of recent historic data points}$$

$$\mu_{p,s} = \text{average}(\mathcal{L}_{p,s})$$

$$\sigma_{p,s} = \text{std.dev}(\mathcal{L}_{p,s})$$

$$\text{Lower Confidence value:} \quad LC_{p,s} = \min\left(\mu_{p,s} \pm t_{m-1,1-\alpha/2}\sqrt{\frac{\sigma^2_{p,s}}{m}}\right)$$

$$\text{Upper Confidence value:} \quad UC_{p,s} = \max\left(\mu_{p,s} \pm t_{m-1,1-\alpha/2}\sqrt{\frac{\sigma^2_{p,s}}{m}}\right)$$

Figure 7.1 demonstrates the difference between a C.I. produced using the t-Distribution from that of the Normal Distribution where $\alpha = 98\%$ for the given set of 26 historic lead times ($m = 26$). The t-Distribution will be used to determine the C.I. as it defines the most frequent behaviour of the supplier and greatly reduces the margin of choice for the next lead time to likely occur.

$$\mathcal{L}_{p,s} = \{14, 14, 11, 17, 10, 18, 10, 14, 14, 14, 14, 14, 14, 21, 14, 21, 14, 16, 14, 15, 19, 14, 11, 12, 15, 12\}$$

$\alpha = 0.98$

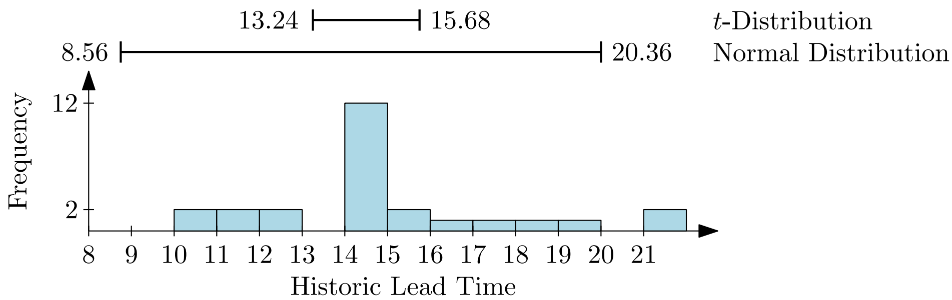

FIGURE 7.1: *Confidence intervals: t-Distribution vs Normal Distribution.*

The HLT & ND Inventory Model uses an iterative approach for predicting the unmet demand which may occur one lead time from now (L) until the next review's order arrives ($L + R$), which is used to determine the base order quantity (q). This is described in § 6.2.1. However, it has now been shown that real world lead times may fluctuate and more than one supplier may exist to choose from for each product. This means that the before-mentioned lead time (L) of a product should not simply be assumed as an exact value.

The lower confidence value ($LC_{p,s}$), calculated from the t-Distribution, is the probable shortest number of days that the order may take to arrive when ordered from supplier s for product p. By replacing the previously used lead time value (for the L measure) with the lower confidence value, the model is able to include additional predicted unmet demands for the base order quantity (q) in the event of an early delivery.

Similarly, where the $L + R$ measure is used to determine the final forecast time step, the upper confidence value ($UC_{p,s}$) can replace the lead time value to add any additional predicted unmet demands to the base order quantity in the event of the next delivery taking longer. Therefore, the inventory model will henceforth quantify orders between $LC_{p,s}$ and $UC_{p,s}+R$ of the forecast.

Cumulative Discrete Exponential Distribution (CDED)

The *Cumulative Discrete Exponential Distribution* (CDED) uses the frequency of an event to observe where the greatest rate of change occurs. If the frequency of a specific lead time occurring is denoted as $f_{p,s}(x_L)$ where x_L is the lead time value in days, then the CDED at X_L is denoted as $F_{p,s}(x_L)$:

$$F_{p,s}(x_L) = \sum_{i=\min(\mathcal{L}_{p,s})}^{x_L} f_{p,s}(i), \quad x_L = \{\min(\mathcal{L}_{p,s}), \min(\mathcal{L}_{p,s}) + 1, \ldots, \max(\mathcal{L}_{p,s})\}$$

Figure 7.2 is an example of the CDED created using the same historic data set as Figure 7.1. In this example, it can be observed that the greatest rate of change occurred where $x_L = 14$ days. Only a small percentage of historic lead times occurred before 14 days indicating a low probability of orders arriving in less than 2 weeks.

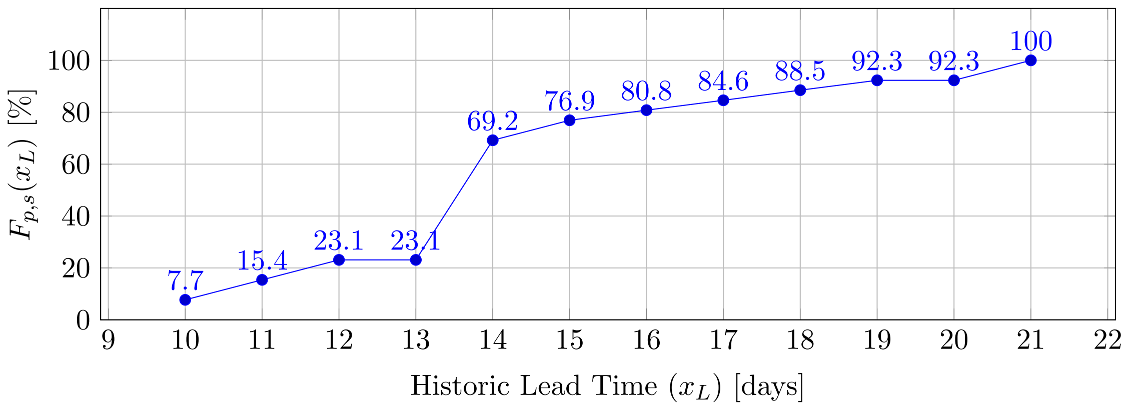

FIGURE 7.2: *Example of the Cumulative Discrete Exponential Distribution (CDED).*

Each supplier, with its own historic lead times, will produce some CDED values, $F_{p,s}(x_L)$. These values can be used to quantify each supplier's ability to meet the expected unmet demand values (UD_{tt}) estimated by the forecast $tt \in \{0, UC_{p,s} + R\}$ for $tt \geq LC_{p,s}$. This is achieved by first calculating the total expected value of the anticipated unmet demand to occur for product p, $E_{p,s}$. The expected value of each supplier is calculated between their own lower C.I. ($LC_{p,s}$) until the overall largest $UC_{p,s} + R$ value, U_p:

$$U_p = \max\{(UC_{p,s_1} + R), (UC_{p,s_2} + R), \ldots, (UC_{p,s_m} + R)\} \tag{7.1}$$

$$e_{p,s}(x_L) = UD_{x_L} \times F_{p,s}(x_L)$$

$$E_{p,s} = \sum_{i=LC_{p,s}}^{U_p} e_{p,s}(i)$$

For example, Table 7.1 shows the how the expected daily unmet demand (UD_{tt}) is determined by using the forecast demand (D_{tt}) for each time period tt. The inventory level (I_{tt}) is a count of the number of items expected to be in stock at the *start* of the day. In this example, there are no past orders scheduled to arrive. The total expected value, $E_{p,s}$, is the sum of the values highlighted in yellow. In this example the lower C.I. values for each supplier are $LC_{p,s_1} = 13$, $LC_{p,s_2} = 12$ and $LC_{p,s_3} = 11$. The upper C.I. values are $UC_{p,s_1} = 16$, $UC_{p,s_2} = 16$ and $UC_{p,s_3} = 15$. Given a review period of $R = 7$, the summation continues until $U_p = \max[(16+7), (16+7), (15+7)] = 23$.

TABLE 7.1: *Example of determining the supplier expected values, $E_{p,s}$.*

				s_1		s_2		s_3	
tt	I_{tt}	D_{tt}	UD_{tt}	$F_{p,s_1}(x_L)$	e_{p,s_1}	$F_{p,s_2}(x_L)$	e_{p,s_2}	$F_{p,s_3}(x_L)$	e_{p,s_3}
0	220	20	0	0	0	0	0	0	0
1	200	22	0	0	0	0	0	0	0
2	178	20	0	0	0	0	0	0	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
9	39	22	0	0	0	0.19	0	0.15	0
10	17	21	4	0.08	0.31	0.23	0.92	0.15	0,62
11	0	18	18	0.15	2.77	0.27	4.85	0.15	2.77
12	0	21	21	0.23	4.85	0.27	5.65	0.23	4.85
13	0	19	19	0.23	4.38	0,27	5.12	0,27	5.12
14	0	20	20	0.69	13.85	0,77	15.38	0,92	18.46
15	0	20	20	0.77	15.38	0,77	15.38	0,92	18.46
16	0	21	21	0.81	16.96	0,81	16.96	0,92	19.38
17	0	18	18	0.85	15.23	0,92	16.62	0,92	16.62
18	0	20	20	0.88	17.69	0,92	18.46	0,92	18.46
19	0	22	22	0.92	20.31	0,92	20.31	0,92	20.31
20	0	18	18	0.92	16.62	0,96	17.31	0,96	17.31
21	0	21	21	1	21	1	21	1	21
22	0	19	19	0	19	0	19	0	19
23	0	20	20	0	20	0	20	0	20
			$E_{p,s}$		180.42		191.19		201.73

7.1.2 Order cost

Based on the $E_{p,s}$ values in Table 7.1, Supplier 3 (s_3) is expected to be more likely to meet more of future demand. However, this alone does not make it the optimal supplier. Cost plays a major roll in choosing a supplier. For example, assume that the same three suppliers from before (s_1,s_2 and s_3) charge the prices shown in Table 7.2. Per item, Supplier 3 is actually the most expensive option. Perhaps one of the other suppliers are capable of producing near-equivalent results at a lower cost.

TABLE 7.2: *Example of supplier pricing.*

Supplier	Cost per batch [ZAR]	Batch size ($MOQ_{p,s}$)	Resulting cost per item [ZAR]
s_1	R20.00	10	R2.00
s_2	R28.00	15	R1.87
s_3	R15.00	5	R3.00

It has been previously stated that the base order quantity will be quantified between $LC_{p,s}$ and $UC_{p,s} + R$ of the forecast, meaning that suppliers which produce different lower and upper confidence intervals may also result in different order quantities ($q_{p,s}$). The supplier batch size ($MOQ_{p,s}$) of the product (p) is then used to determine the final order quantity that would be placed, $Q_{p,s}$. From this the cost of the order (Cp, s) can be calculated using the price set by the supplier per batch size ($c_{p,s}$):

$$Q_{p,s} = \left\lceil \frac{q_{p,s}}{MOQ_{p,s}} \right\rceil \times MOQ_{p,s}$$

$$Cp, s = \frac{Q_{p,s}}{MOQ_{p,s}} \times c_{p,s}$$

Consider the example in Table 7.1 and the pricing shown in Table 7.2. Table 7.3 demonstrates how the base order quantities are determined using the lower and upper confidence values, then scaled based on the defined batch size before determining the order cost.

TABLE 7.3: *Example of supplier order costs.*

Supplier	Base quantity $q_{p,s}$	Final quantity $Q_{p,s}$	Cost (Cp, s)
s_1	$q_{p,s_1} = \sum_{i=LC_{p,s_1}}^{UC_{p,s_1}+R} UD_i = \sum_{i=13}^{16+7} UD_i = 218$	220	R440.00
s_2	$q_{p,s_2} = \sum_{i=LC_{p,s_2}}^{UC_{p,s_2}+R} UD_i = \sum_{i=12}^{16+7} UD_i = 239$	240	R448.00
s_3	$q_{p,s_3} = \sum_{i=LC_{p,s_3}}^{UC_{p,s_3}+R} UD_i = \sum_{i=11}^{15+7} UD_i = 237$	240	R720.00

7.1.3 Choosing the optimal supplier

Next is to perform a cost-benefit comparison to determine which supplier has the best efficiency. Supplier 3 achieved the best total expected value ($E_{p,s_3} = 201.73$) due to a higher likelihood of an early delivery which might meet more demand. However, this is only 5.52% better than Supplier 2's expected value ($E_{p,s_2} = 191.19$) and costs R272.00 (60.71%) more. Logic dictates that Supplier 2, with a near-equivalent expected value and a much lower order cost, would be the best choice for placing an order.

In order to arrive at this result each supplier must be compared with the others, first in a benefit comparison and secondly in a cost comparison. The measure of benefit that supplier x has over supplier y ($B_{x,y}$) is calculated by dividing the total expected value of supplier x (E_{p,s_x}) by the total expected value of supplier y (E_{p,s_y}). Similarly, the measure of cost for supplier x over supplier y ($K_{x,y}$) is calculated by dividing the order cost with supplier x (C_{p,s_x}) by the order cost with supplier y (C_{p,s_y}).

$$B_{x,y} = \frac{E_{p,s_x}}{E_{p,s_y}}$$

$$K_{x,y} = \frac{C_{p,s_x}}{C_{p,s_y}}$$

A larger $B_{x,y}$ value is desired to indicate that supplier x has a much higher ability to meet demand. While a smaller $K_{x,y}$ value is desired, showing that supplier x is the cheaper option. The efficiency value ($\xi_{x,y}$) is therefore:

$$\xi_{x,y} = \frac{B_{x,y}}{K_{x,y}}$$

An efficiency value $\xi_{x,y} < 1.0$ indicates that supplier x is a worse choice than supplier y. While an efficiency value $\xi_{x,y} = 1.0$ indicates that suppliers x and y are equally good/poor choices. However, an efficiency value $\xi_{x,y} > 1.0$ indicates that supplier x is a better choice than supplier y and should be prioritised.

In the case of the three suppliers from Tables 7.1 & 7.3, the efficiency values have been calculated and shown in Table 7.4. Supplier 2 (s_2) achieved the largest overall efficiency value ($\xi_{s_2,s_3} = 1.5232$) making it the optimal supplier, as suggested earlier. This means that s_2 assumes the first position in the priority list, $\mathcal{P}_p[1] = s_2$.

TABLE 7.4: *Example of supplier efficiency values, part 1.*

	$B_{x,y}$			$K_{x,y}$			$\xi_{x,y}$		
x \ y	s_1	s_2	s_3	s_1	s_2	s_3	s_1	s_2	s_3
s_1	–	0.9437	0.8944	–	0.9821	0.6111	–	0.9608	1.4635
s_2	1.0597	–	0.9478	1.0182	–	0.6222	1.0408	–	1.5232
s_3	1.1181	1.0551	–	1.6364	1.6071	–	0.6833	0.6565	–

Now s_2 can be ignored, as shown in Table 7.5, in order to determine which supplier will have second priority. Of the two remain suppliers, s_1 and s_3, Supplier 1 had the largest efficiency value ($\xi_{s_1,s_3} = 1.4635$). Therefore, s_1 assumes the second position in the priority list, $\mathcal{P}_p[2] = s_1$. In this case, only one supplier remains and will thus assume the third position in the priority list, $\mathcal{P}_p[3] = s_3$. The final priority set is: $\mathcal{P}_p = \{s_2; s_1; s_3\}$.

TABLE 7.5: *Example of supplier efficiency values, part 2.*

	$B_{x,y}$			$K_{x,y}$			$\xi_{x,y}$		
x \ y	s_1	s_2	s_3	s_1	s_2	s_3	s_1	s_2	s_3
s_1	–	0.9437	0.8944	–	0.9821	0.6111	–	~~0.9608~~	1.4635
s_2	1.0597	–	0.9478	1.0182	–	0.6222	~~1.0408~~	~~–~~	~~1.5232~~
s_3	1.1181	1.0551	–	1.6364	1.6071	–	0.6833	~~0.6565~~	–

7.1.4 When to prioritise suppliers

Each healthcare facility has its own list of suppliers. For some products there will be a number of suppliers to choose from and for other products there may only be one available supplier. A contractual agreement must be established that spans some period of time, commonly 1a number of years. The specifics of these contracts vary and can define how-much or how-little the facility may order, and may even include deals that reduce prices in exchange for becoming the primary supplier. It is not in the scope of this paper to model every unique contract. Rather, the model will assume that suppliers get reviewed either monthly, biannually or annually. This review would be used to generate the priority set for use during the next period.

Suppliers, being aware of the review process, may be driven towards improving their performance in order to acquire a higher standing in the priority list. The number of entries from the historic lead time set used to calculate the priority list could have an affect on the suppliers' behaviour.

7.2 Order Irregularities

Suppliers do their best to meet the total order quantities issued to them, however there are always limits to what can be provided. Although suppliers often provide a minimum order quantity (MOQ), they seldom define a maximum/paramount order quantity (POQ). When suppliers are overwhelmed with orders they may not be able to meet every client's target order quantity. Each supplier may handle these situations differently. Suppliers could use prioritisation, fully supplying a few high priority clients while giving low priority clients nothing. The supplier could share the available product quantity across all clients providing everybody with the same number of items, or alternatively the same percentage of their order quantities, and back order the rest of each order for a later date. Whatever the supplier chooses to do, the healthcare facility may not be informed about the limitations and only witness the outcomes when receiving.

It can be assumed that the smaller the order quantity, the more likely the supplier is capable meeting the full demand. Ideally, the supplier will always deliver 100% of the order, as demonstrated in Figure 7.3a. Realistically suppliers are limited to some discrete level due to manufacturing rates and capacity constraints. Orders which exceed the supplier's ability are unlikely to receive the full order. This is described by Figure 7.3b.

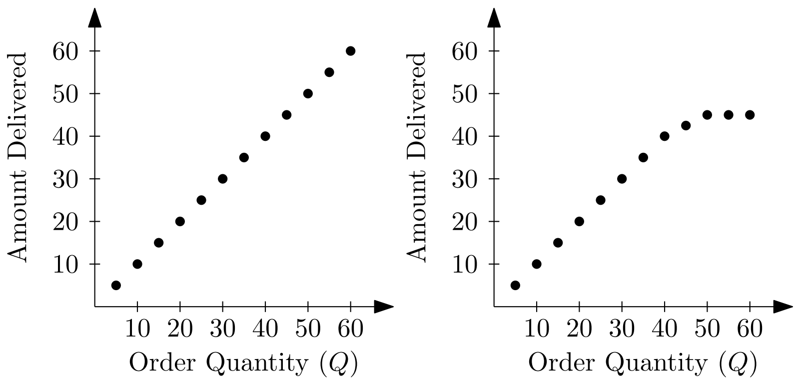

(a) *Ideal delivery behaviour* (b) *Realistic delivery behaviour*

FIGURE 7.3: *Example of ideal versus realistic delivery behaviours.*

Historic orders can be compared to the actual amount that arrived in order to determine the maximum allowable order quantity that may be placed with a supplier. Any additional items that need to be ordered will be issued to the next supplier on the priority list ($\mathcal{P}_p$). Figure 7.4 demonstrates how polynomial curves (degree= 2) have been fitted through four supplier's historic delivery amounts (AD) in order to estimate the trend and to determine the critical point associated with each supplier.

Although the critical point appears to identify the maximum likely quantity of product p to be delivered by supplier s, the linearity of the trend begins to falter long before reaching the

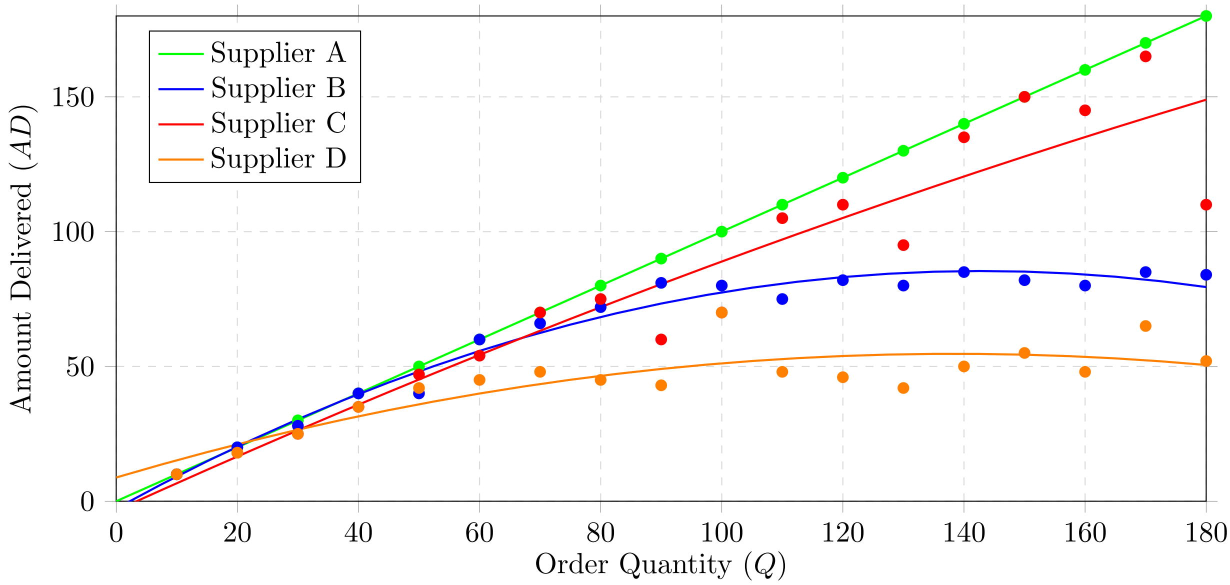

FIGURE 7.4: *Example of fitting a polynomial curve to delivery behaviours.*

critical point. It is the linearity between the order quantity issued and the amount received which identifies the capability of the supplier. Supplier A has a perfect delivery record resulting in an ideal linear relationship and will not curve nor have a critical point. In this scenario the paramount order quantity will be 50% greater than the largest historic order quantity to date.

$$POQ_{p,s} = Q_{p,s} \times 1.5, \quad \text{if perfect delivery record}$$

However, Suppliers C and D failed to keep up with the orders. In these non-ideal scenarios a different approach is necessary to determine the paramount order quantity. If $g_{p,s}(Q, AD)$ is a second degree polynomial equation fitted to the historic set of orders, then the $POQ_{p,s}$ is the last order quantity which can be found to result in an expected delivered amount larger than or equal to some significance level ($\alpha \in [0,1]$) of the desired order. For example, if $\alpha = 0.9$ then the paramount order quantity ($POQ_{p,s}$) is equal to the smallest order quantity ($Q_{p,s}$) which resulted in $\geq 90\%$ of the order being delivered according to the polynomial trend ($g_{p,s}(Q_{p,s}, AD_{p,s})$).

$$POQ_{p,s} = \min(Q_{p,s}), \quad \text{where} \quad \alpha Q_{p,s} \geq g_{p,s}(Q_{p,s}, AD_{p,s})$$

Figure 7.5 is a flow diagram depicting how an order for q_p items of product p gets distributed amongst the suppliers based on the supplier priority list ($\mathcal{P}_p$) and the respective $POQ_{p,s}$ values. If only one supplier is listed in the supplier priority list, then the only option is to place the full order with that supplier, scaled to the necessary batch size ($MOQ_{p,}$). If more than one supplier is available, then starting with the first supplier ($x = 1$) in the priority list, the desired amount (q_p) is scaled according to the supplier's batch size ($MOQ_{p,\mathcal{P}[\mathrm{x}]}$) to some value, $X_{p,\mathcal{P}[\mathrm{x}]}$. If this quantity does not exceed the $POQ_{p,\mathcal{P}[\mathrm{x}]}$ then the order is of a confident size and can be issued to the supplier. No further manipulation would be required.

However, if the paramount order quantity is exceeded, then the order with the current supplier (x) is reduced to the $POQ_{p,\mathcal{P}[\mathrm{x}]}$ value and the remaining quantity ($q_p = q_p - POQ_{p,\mathcal{P}[\mathrm{x}]}$) gets reviewed in the same manner for the next supplier in the priority list. If no further suppliers exist in the priority list, but some quantity still remains to be ordered ($q_p > 0$), it becomes necessary to exceed the suppliers' paramount order quantity levels. Starting again at the first supplier in the priority list ($x = 1$), the supplier's order quantity is increased by one batch size ($MOQ_{p,\mathcal{P}[\mathrm{x}]}$). If the amount that still needed to be distributed to suppliers (w) was less than the batch size, then the full quantity has been ordered and no further manipulation is required.

If the batch size was not enough to satisfy the amount that still needed to be distributed ($MOQ_{p,\mathcal{P}[\mathrm{x}]} < w$), then w is reduced by the batch size (which has already been appended to the order of supplier x). The next supplier in the priority list ($x + 1$) then has their order quantity increased by one batch size. The order quantities of each supplier continues to be incrementally increased by one batch size until the full desire order quantity has been achieved.

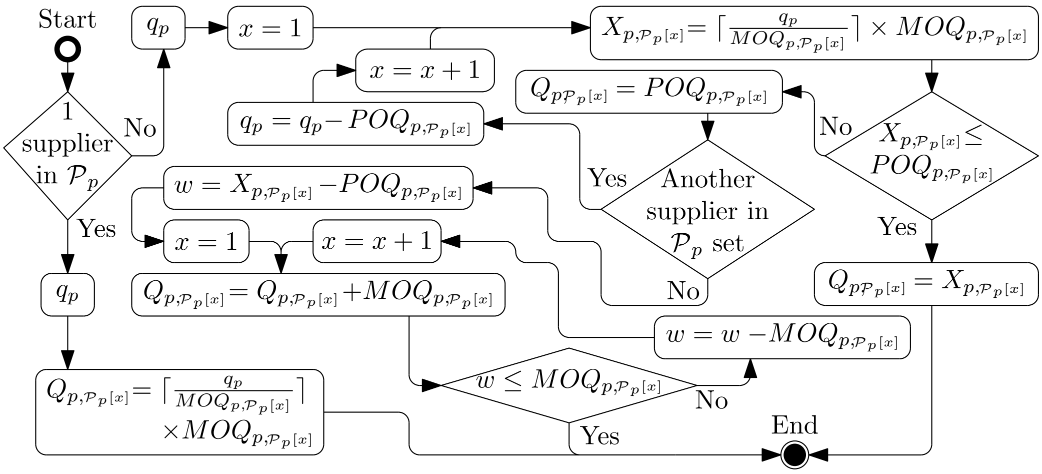

FIGURE 7.5: *Orders assignment based on supplier priority.*

7.3 Storage Limitations

Healthcare facilities vary in size and each have their own storage limitations in terms of the amount of inventory that can be kept on hand. It was found during the systematic literature review that some inventory policies account for this by using a Par value, used in order-up-to inventory models, to ensure that no over-ordering occurs. Such models were the Par inventory policy and Policies 5–9, described in § 4.6.4. It is important not to place orders which may, upon arrival, result in inventory levels exceeding the maximum available storage space.

7.3.1 Estimating the storage space used

The *The HLT & ND Inventory Model* performs an iterative daily forecast of the inventory's behaviour. New orders are expected to arrive after $LC_{p,s}$ days. For each time-step during the forecast of product p the remaining inventory at the end of the day ($i_{tt,p}$) can calculated by increasing the inventory level of the previous day ($i_{tt-1,p}$) by the expected orders for the current day (o_{tt}), minus the expected demand ($d_{tt,p}$).

$$i_{tt,p} = \begin{cases} i_{tt-1,p} + o_{tt} - d_{tt,p} & \text{,if } (i_{tt-1,p} + o_{tt}) > d_{tt,p} \\ 0 & \text{,else} \end{cases}$$

These inventory levels can be used to calculate the total expected storage space used at the end of each day (ES_{tt}). Provided that a decision maker has defined a measure of space (h_p) for one item of product p, ES_{tt} can be quantified. The value U_p is calculated using Equation 7.1 (page 65) and $P_{\max}$ is the number of products at the healthcare facility.

$$ES_{tt} = \sum_{p=1}^{P_{\max}} i_{tt,p} \times h_p, \quad \text{for each } tt \in [0, U_p]$$

7.3.2 Constraint for storage levels

The total holding space that is available in the facility (H) must be define in terms of the same measure as the h_p values were. If ES_{tt} for any day tt exceeds the total available holding space ($ES_{tt} > H$) then changes need to be made to the orders. Any expected overstocking which occurs before any of the new orders would arrive cannot be changed. However, any overstocking that may occur as a result of the new orders arriving can be reviewed. This is to say that space should be reviewed during the iterative forecast for $tt \in [\min\{LC_1, LC_2, \ldots\}, U_p]$.

7.3.3 Product priority categorization

It has now been established that orders may need to be changed in order to prevent overstocking, but choosing which orders to alter and the degree of change can be difficult to determine. In healthcare, it is inappropriate to prioritize products based on demand or cost. Drugs which are seldom required may still need to be readily available for irregular circumstances. Even tools such as syringes and needles come in different tip forms and sizes, each designed for a particular purpose [11].

To ensure that products get prioritised appropriately the priority matrix depicted in Table 7.6 will be used. A decision maker familiar with the products is tasked with assigning each product a *Product Category* code (A–K) which defines the level of importance of the product. A-category products are the most vital and will never have their orders altered. These products will be medication patients cannot go without. In cases such as tuberculosis or antiretrovirals, skipping even one treatment could result in the patient developing a resistance to the medication. K-category products are the least essential and orders can be completely reduced to zero. These items are often easy to replace with alternative products in the case of a stock-out.

TABLE 7.6: *Product priority matrix.*

Product	Priority Level										
Category	1	2	3	4	5	6	7	8	9	10	11
A	1	1	1	1	1	1	1	1	1	1	1
B	1	1	1	1	1	1	1	1	1	1	0.9
C	1	1	1	1	1	1	1	1	1	0.9	0.8
D	1	1	1	1	1	1	1	1	0.9	0.8	0.7
E	1	1	1	1	1	1	1	0.9	0.8	0.7	0.6
F	1	1	1	1	1	1	0.9	0.8	0.7	0.6	0.5
G	1	1	1	1	1	0.9	0.8	0.7	0.6	0.5	0.4
H	1	1	1	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3
I	1	1	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2
J	1	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1
K	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0

The model starts at *Priority Level 1* ($PL = 1$) and will incrementally increase until the space constraint has been met. Preference values ($Pr[PL, p]$) are extracted from Table 7.6 for each product (p) based on its category and the current PL. A preference value $Pr[PL, p] = 1.0$ indicates that 100% of the desired order quantity will be ordered. Similarly, a preference value of $Pr[PL, p] = 0.0$ means that none of the product will be ordered. Healthcare facilities that find themselves regularly using large priority levels $PL \geq 6$ should either acquire more storage space or shorten their review periods, placing more frequent orders.

7.3.4 Performing the space constraint

Figure 7.6 is a flow diagram depicting how the events described in § 7.3.1–7.3.3 are used together to enforce the storage space limitations. The model steps through each day of the forecast ($tt \in [\min\{LC_1, LC_2, \ldots\}, U_p]$), starting at the earliest lower confidence lead time value of the suppliers and ending at the U_p value calculated from Equation 7.1 on page 65.

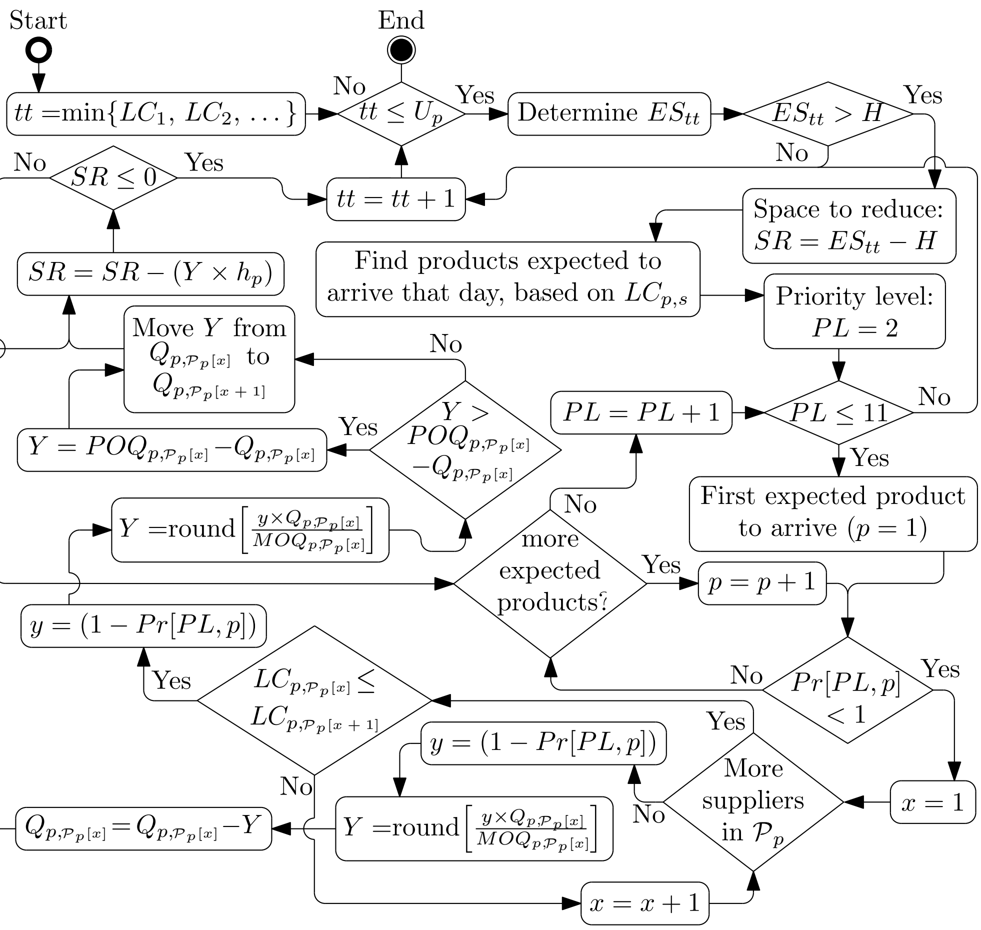

FIGURE 7.6: *Orders review and reassignment for storage space constraint.*

At the start of each time step in the forecast (tt), the estimated storage space in use at the end of the day (ES_{tt}) is inspected. If this expected value is *not* larger than the total available holding space at the facility ($ES_{tt} \leq H$), then the next time step in the forecast ($tt + 1$) can be reviewed. However, if $ES_{tt} > H$ then the storage space exceeding H is calculated. This is the amount of space that needs to be reduced (SR) to satisfy the storage constraint. The only event that would cause the used storage space to increase from $tt - 1$ to tt would be the arrival of product orders. Only products orders expected to arrive on that day (based on their respective $LC_{p,s}$ values) will require adjustments.

At this stage the priority level increases to $PL = 2$. Each product expected to arrive that day is reviewed. If the preference value of the product has changed ($Pr[PL, p] < 1$), then $(1 - Pr[PL, p]) \times 100\%$ of the product order, rounded to the nearest batch size (Y), attempts to get transferred to the next supplier in the suppliers priority list ($\mathcal{P}_p$). The next supplier must have an equal or larger $LC_{p,s}$ value than that of the current supplier. If there are no suppliers in the suppliers priority list that meet this condition, the desired amount cannot be moved. In this scenario the order quantity is simply reduced by the rounded preference amount (Y).

If transferring or removing Y items of the product has not reduced the storage space by the required amount (SR), then the next product expected to arrive that day is reviewed. If $SR > 0$ and no more products exist in the list, then the priority level is increased ($PL = PL + 1$) and the process repeats. If all eleven priority levels have been tested for that day and $SR > 0$, the model accepts that the storage space constraint may be exceeded due to the necessity of the products.

7.4 Project Progress

Objective VII was achieved during this chapter. This chapter completed the *Test phase* and the first target of the *Development phase* of the project methodology framework. The updated framework is shown in Figure 7.7.

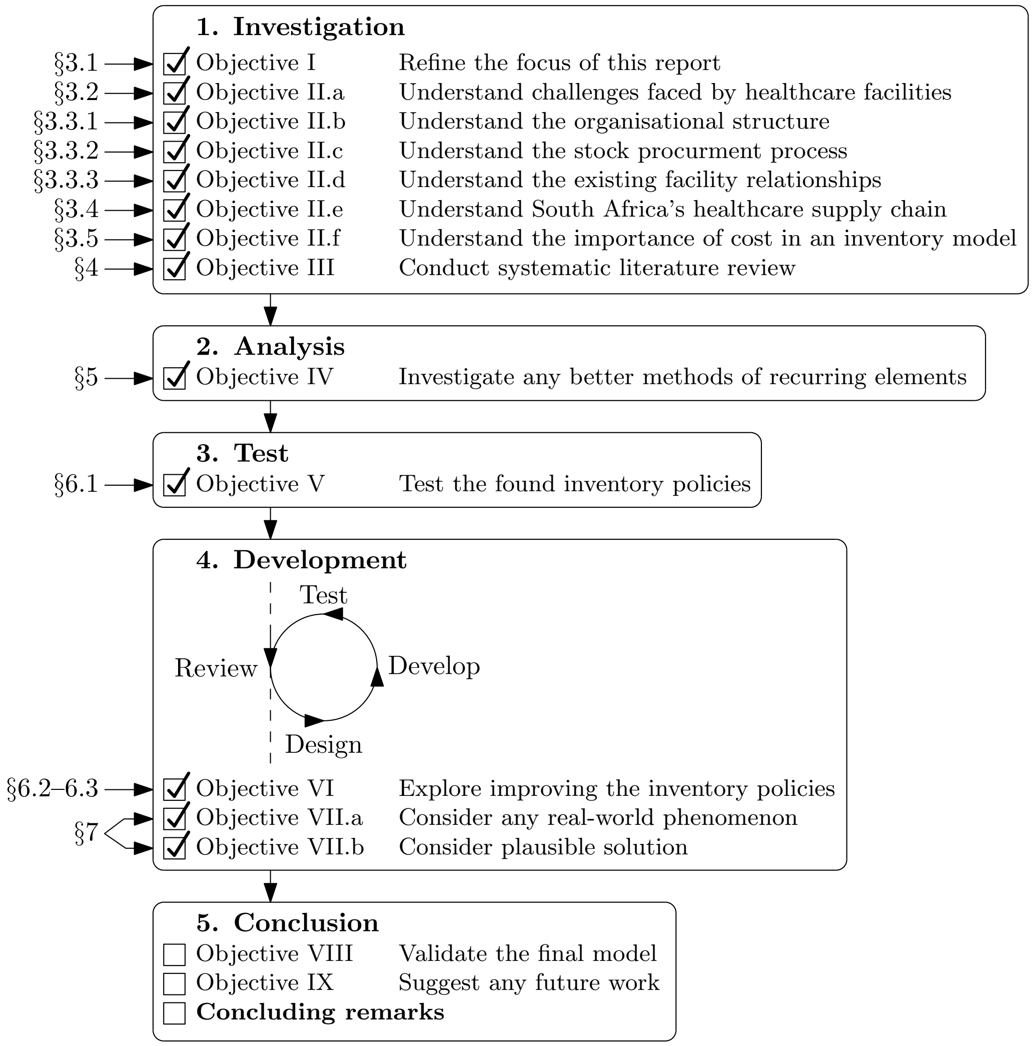

FIGURE 7.7: *Project methodology framework: Chapter 7.*

CHAPTER 8

Validation

Contents

This chapter will cover the first target of the *Conclusion phase* in the project methodology framework, § 2.2, which aims to achieve Objective VIII from the project objectives list, § 1.4.

8.1 Inventory Model Validation

Chapter 6 tested and developed various inventory policy models in the attempt to find a suitable solution capable of quantifying orders for South Africa's public healthcare facilities. A scenario was developed to test the inventory policies. This depicted a healthcare facility keeping 100 types of products in stock. The product specifications were extracted from the real-world Master Procurement Catalogue [80] provided by the NDoH and were selected in a way that represented the lead time distribution of the full contract set (recall Table 3.3 from § 3.4.1).

The HLT & Normal Distribution Inventory Model (§ 6.3) achieved impressive results with safety stock ranging from thirty percent ($ss \geq 30\%$) and service levels ranging from as low as ten percent ($\alpha \geq 10\%$). Recall that the service level is used to find the minimum forecast value based on a normal distribution of the historic demand set. A low service level is desired. Therefore, provided that procurement data and daily inventory levels are available, appropriate order quantities can be determined and issued to suppliers. Subject matter experts (SMEs) working on the Stock Visibility System (SVS), described in § 1.1, have shown great interest in the HLT & Normal Distribution Inventory Model. However, the next phase of their system, the *SVS 2.0*, will encourage healthcare facilities to update inventory levels only once a week. The HLT & Normal Distribution Inventory Model will need to be tested under these conditions, using weekly- as opposed to daily-inventory levels.

8.1.1 The Revised HLT & ND Inventory Model

The larger the gap between each recorded inventory level, the greater the uncertainty of the demand during that period of time. Taking into consideration that all prior inventory models tested in this research paper were provided with ideal records of uninterrupted daily inventory levels, a similar assumption of diligent weekly stock counts will be applied to this model. This will allow for a fair side-by-side comparison of the results after testing.

The proposed HLT & ND Inventory Model (§ 6.3) has used historic daily demand levels to predict future demand and inventory behaviour. At each review period, the most recent 30 days historic demand would be used to create a forecast of the expected future demand. These daily demand levels were calculated using the known daily inventory levels and the order arrivals. Capturing inventory levels only once a week create scenarios that vary the demand visibility over each week. Seven explicit scenarios (*Scenarios A–G*) will be described, as well as the solution for calculating the weekly demand in each case. Implementing these solutions into the existing inventory model will create the *The Revised HLT & ND Inventory Model* in an attempt to satisfy the SVS 2.0 design.

Resolving demand prediction challenges

The most ideal scenario occurs when both the current inventory level this week (I_T) and the previous inventory level last week (I_{T-1}) are positive. This can be subdivided into two unique scenarios; **Scenario A**, which received some positive order quantity during the week ($O_{T-1|T}$) from suppliers and **Scenario B**, which did not receive any orders. Scenarios A and B are depicted by Figures 8.1 and 8.2 respectively. Notice that in both scenarios $I_T > O_{T-1|T}$. This indicates that some of last week's inventory still exists and no stock-outs occurred. For this reason the weekly demand ($D_{T-1|T}$) can confidently be calculated as $D_{T-1|T} = I_{T-1} + O_{T-1|T} - I_T$.

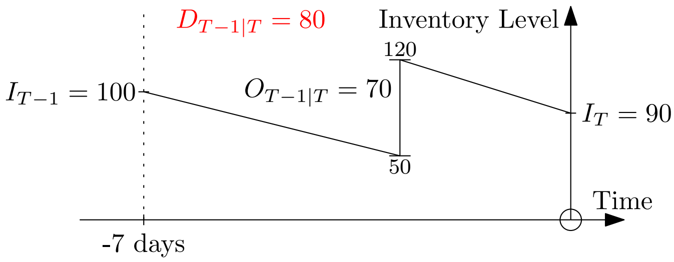

FIGURE 8.1: *Weekly inventory capture: Scenario A.*

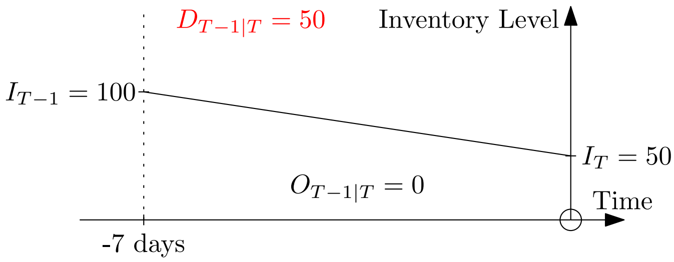

FIGURE 8.2: *Weekly inventory capture: Scenario B.*

Scenario C, depicted by Figure 8.3, represents what may happen when $I_T \leq O_{T-1|T}$. Both week's inventory levels are positive, but a stock-out may have occurred at some point during the week. Figure 8.4 shows that it is also possible that no stock-out occurred. Because the actual demand may be larger than calculated, the historic weekly demands will be reviewed. To prevent under-assuming the demand and producing low forecast values, the set of historic weekly demands ($\mathcal{D}$) will be reviewed.

$$\mathcal{D} = \{D_{T-n}, \ldots, D_{T-1}\}, \quad n = \text{number of historic weekly demands in the set}$$

The maximum historic value in the set may have been a rare occurrence. Therefore, to prevent over-assuming the weekly demand the inverse of the cumulative normally distributed value ($F^{-1}(\alpha, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$) at an 80% service level ($\alpha = 0.8$) will be used. The mean ($\mu_{\mathcal{D}}$) and standard deviation ($\sigma_{\mathcal{D}}$) are calculated from the historic demand set ($\mathcal{D}$) which must have at least two entries. The weekly demand will be set to the maximum value between the calculated weekly demand and $F^{-1}(\alpha, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$.

$$D_{T-1|T} = \max(I_{T-1} + O_{T-1|T} - I_T, F^{-1}(0.8, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}}))$$

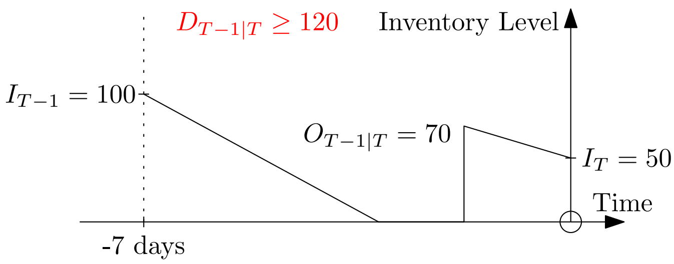

FIGURE 8.3: *Weekly inventory capture: Scenario C — stock-out.*

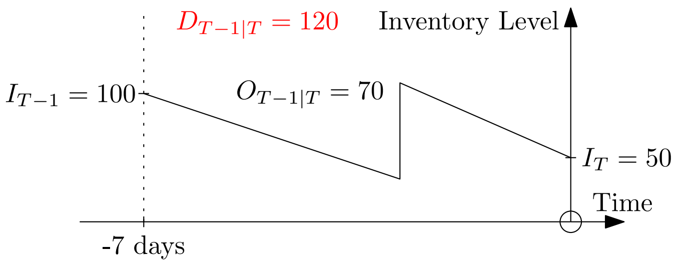

FIGURE 8.4: *Weekly inventory capture: Scenario C — no stock-out.*

Scenario D is the event of a current a stock-out ($I_T = 0$), provided there was stock last week ($I_{T-1} > 0$), but no order arrivals. As seen in Figure 8.5, the demand for Scenario D will be larger than or equal to I_{T-1}. It is not possible to know when during the week the inventory became depleted. Similarly to Scenario C, to prevent under-assuming the demand, the weekly demand ($D_{T-1|T}$) will be set to the maximum value between I_{T-1} and $F^{-1}(\alpha, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$.

$$D_{T-1|T} = \max(I_{T-1}, F^{-1}(0.8, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}}))$$

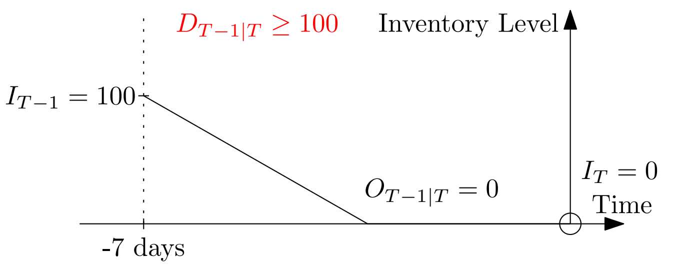

FIGURE 8.5: *Weekly inventory capture: Scenario D.*

Scenario E provides no information on the week's demand. Both inventory levels are stocked-out and no orders were delivered, as shown in Figure 8.6. The maximum historic weekly demand will be assigned the inverse of the cumulative normally distributed value.

$$D_{T-1|T} = F^{-1}(0.8, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$$

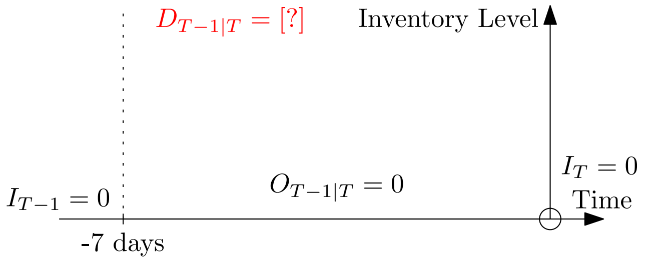

FIGURE 8.6: *Weekly inventory capture: Scenario E.*

Scenario F, depicted by Figure 8.7, occurs when the previous week was stocked-out ($I_{T-1} = 0$), an order arrived during the week and the current inventory level indicates remaining stock ($I_T > 0$). Provided that the date of the order was logged, it is possible to determine the number of days which the inventory satisfied demand ($\triangle t$). Dividing the known demand ($O_{T-1|T} - I_T$) by $\triangle t$ will give an average daily demand, then multiplied by seven will estimate the weekly demand.

$$D_{T-1|T} = \left(\frac{O_{T-1|T} - I_T}{\triangle t} \right) \times 7$$

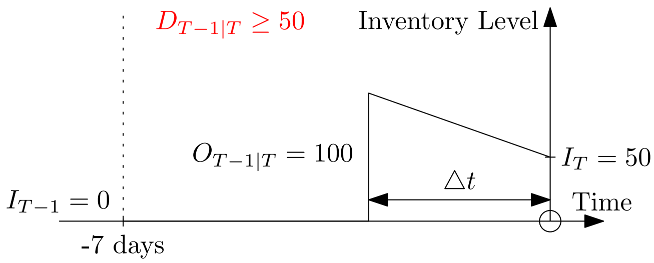

FIGURE 8.7: *Weekly inventory capture: Scenario F.*

Scenario G is when both weeks show stocked-out inventory levels, but an order did arrive during the week. Similarly to Scenario F, an estimate of the weekly demand can be calculated. However, there is no way of knowing how quickly $O_{T-1|T}$ was consumed during $\triangle t$. For this reason the maximum between the estimated weekly demand and the $F^{-1}(0.8, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$ value will be selected.

$$D_{T-1|T} = \max \left(\frac{O_{T-1|T}}{\triangle t} \times 7, F^{-1}(0.8, \mu_{\mathcal{D}}, \sigma_{\mathcal{D}}) \right)$$

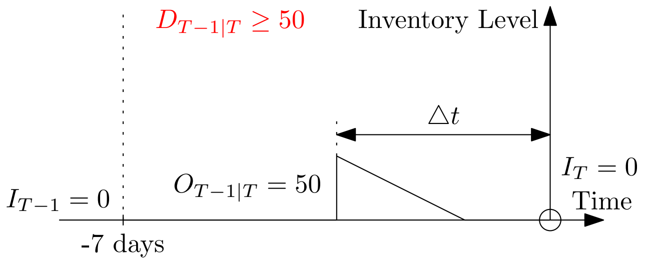

FIGURE 8.8: *Weekly inventory capture: Scenario G.*

Transform the weekly demand for daily forecasting

The weekly demands have been assumed using the weekly inventory levels and order information. Because the original HLT & Normal Distribution Inventory Model (§ 6.3) iteratively observes the expected behaviour of the future inventory levels, daily forecast demand is required. To perform a forecast returning daily demand, the historic data must also be daily. The most recent entries in the of weekly demand historic set ($\mathcal{D}$) will be used for the forecast. These weekly demand levels will simply be divided by seven to estimate the daily historic demand (d_t) which occurred during that week, as shown in Figure 8.9.

$$d_t = \left\lceil \frac{D_{T-1|T}}{7} \right\rceil, \quad \text{for all days } t \text{ between weekly inventory levels } T-1 \text{ and } T$$

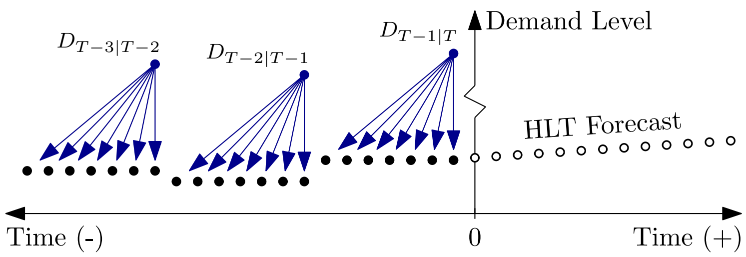

FIGURE 8.9: *Transforming weekly demand for daily forecasting.*

8.1.2 Testing the Revised HLT & Normal Distribution Inventory Model

The Revised HLT & Normal Distribution Inventory Model was tested using the same simulation used to test the other inventory policies in Chapter 7. The model was tested for both weekly and monthly review periods $R = \{7, 30\}$. This will imitate the behaviour of South African public healthcare facilities. In each case, the model was tested using 4, 8 and 12 weeks of historic demand values to observe the difference in performance.

Table 8.1 provides a comparative view of the total unmet demand (number of items for all 100 products in the test) which resulted from weekly reviews ($R = 7$), using the original HLT & Normal Distribution Inventory Model with daily inventory levels, as found in § 6.3.2, versus using the Revised HLT & Normal Distribution Inventory Model. All three cases of the revised model out-performed the original daily inventory level model. Results appear to improve as the number of used historic demand values increases. This may have resulted from using the inverse of the cumulative normally distributed value at $\alpha = 0.8$ as a buffer for uncertain scenarios, which could have assumed a higher demand than what truly occurred and placed larger orders.

Larger orders may cause overstocking and result in expired items. However, Table 8.2 shows that none of the models resulted in expired items. This signifies that all items were consumed before their shelf life ended. Table 8.3 shows the maximum inventory level to occur from each model. The maximum inventory levels actually appear to be very similar between 20% and 30% safety stock (ss).

TABLE 8.1: *Revised HLT & ND Inventory model results, $R = 7$: Total unmet demand.*

Total unmet demand: $R = 7$										
	Daily inventory levels — 30 days					Weekly inventory levels — 4 weeks				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	35 296	1 756	15	0	0	321	17	0	0	0
20%	10 790	240	15	0	0	78	17	0	0	0
30%	2 386	225	15	0	0	13	17	0	0	0
40%	577	185	15	0	0	0	0	0	0	0
50%	186	88	15	0	0	0	0	0	0	0
α	Weekly inventory levels — 8 weeks					Weekly inventory levels — 12 weeks				
10%	35	0	0	0	0	35	0	0	0	0
20%	76	0	0	0	0	16	0	0	0	0
30%	16	0	0	0	0	16	0	0	0	0
40%	0	0	0	0	0	0	0	0	0	0
50%	0	0	0	0	0	0	0	0	0	0

TABLE 8.2: *Revised HLT & ND Inventory model results, $R = 7$: Total expired items.*

Total expired items: $R = 7$										
	Daily inventory levels — 30 days					Weekly inventory levels — 4 weeks				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	0	0	0	0	0	0	0	0	0	0
20%	0	0	0	0	0	0	0	0	0	0
30%	0	0	0	0	0	0	0	0	0	0
40%	0	0	0	0	0	0	0	0	0	0
50%	0	0	0	0	0	0	0	0	0	0
α	Weekly inventory levels — 8 weeks					Weekly inventory levels — 12 weeks				
10%	0	0	0	0	0	0	0	0	0	0
20%	0	0	0	0	0	0	0	0	0	0
30%	0	0	0	0	0	0	0	0	0	0
40%	0	0	0	0	0	0	0	0	0	0
50%	0	0	0	0	0	0	0	0	0	0

The unmet demand resulting from monthly review periods is shown in Table 8.4. A very high safety stock is required for the Revised HLT & Normal Distribution Inventory Model to satisfy the demand. The best results were obtained from using 8 weeks of historic data at 40–50% safety stock, where $\alpha \geq 40\%$. This may sound like a lot of stock, but that is actually not the case. By comparing the number of expired items resulting from these conditions, shown in Table 8.5, it becomes evident that the Revised HLT & Normal Distribution Inventory Model performed better at 40% and 50% safety stock than the original model performed at 30% and 40%.

Additionally, the maximum inventory levels of the Revised HLT & Normal Distribution Inventory Model are substantially lower than the original model for both review periods ($R = \{7, 30\}$), observed from Tables 8.3 and 8.6. Comparatively, the Revised HLT & Normal Distribution Inventory Model requires at least 10% more safety stock than the original model to reach the same inventory levels.

TABLE 8.3: *Revised HLT & ND Inventory model results, $R = 7$: Maximum inventory level.*

Maximum inventory level: $R = 7$					
$\alpha \backslash ss$	Daily inventory levels — 30 days 10%	20%	30%	40%	50%
10%	85 177	106 377	128 331	147 325	168 915
20%	87 841	109 083	130 547	149 583	170 859
30%	90 632	111 834	132 478	151 618	173 292
40%	93 276	113 735	134 667	153 357	175 743
50%	95 251	115 258	137 091	155 353	177 420
α	Weekly inventory levels — 4 weeks				
10%	102 264	117 094	127 583	140 532	158 756
20%	102 867	117 613	126 793	141 179	159 467
30%	102 941	118 928	126 557	140 917	160 105
40%	103 503	119 516	127 198	141 095	160 540
50%	103 573	118 835	128 087	141 811	161 109
α	Weekly inventory levels — 8 weeks				
10%	101 056	111 673	123 740	138 650	156 175
20%	101 199	112 706	123 952	138 454	156 746
30%	102 233	114 224	124 448	138 792	157 644
40%	101 715	114 885	125 285	139 214	158 393
50%	102 170	113 871	124 984	139 837	158 182
α	Weekly inventory levels — 12 weeks				
10%	98 487	110 871	123 744	138 020	155 971
20%	99 317	111 238	123 713	138 213	156 436
30%	99 413	111 114	124 106	138 513	156 959
40%	99 872	112 513	124 099	139 359	157 973
50%	100 512	112 034	124 672	139 850	159 197

TABLE 8.4: *Revised HLT & ND Inventory model results, $R = 30$: Total unmet demand.*

Total unmet demand: $R = 30$										
	Daily inventory levels — 30 days					Weekly inventory levels — 4 weeks				
$\alpha \backslash ss$	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	31 552	1 412	80	0	0	40 950	27 909	10 197	909	409
20%	7 412	0	0	0	0	40 950	27 909	9 873	909	0
30%	1 054	0	0	0	0	40 842	27 369	8 145	693	0
40%	0	0	0	0	0	43 819	22 401	8 793	0	0
50%	0	0	0	0	0	38 287	21 969	1 341	0	0
α	Weekly inventory levels — 8 weeks					Weekly inventory levels — 12 weeks				
10%	42 276	30 177	11 925	1 125	0	39 802	30 099	13 143	1 665	0
20%	42 276	30 177	11 925	1 125	0	39 802	30 099	13 143	1 665	0
30%	42 276	29 745	11 817	1 125	0	39 802	30 099	13 143	1 665	0
40%	41 952	28 665	8 145	0	0	39 694	29 343	12 819	585	0
50%	44 115	24 237	4 473	0	0	42 937	27 507	8 607	0	0

TABLE 8.5: *Revised HLT & ND Inventory model results, $R = 30$: Total expired items.*

Total expired items: $R = 30$										
	Daily inventory levels — 30 days					Weekly inventory levels — 4 weeks				
α \ ss	10%	20%	30%	40%	50%	10%	20%	30%	40%	50%
10%	484	1 533	2 578	4 467	8 069	0	9	714	1952	13633
20%	539	1 573	2 758	4 467	8 269	0	9	714	1952	14504
30%	573	1 613	2 758	4 467	8 369	0	9	714	1952	16006
40%	773	1 653	2 758	4 467	8 519	0	9	714	1952	15174
50%	773	1 653	2 758	4 467	8 569	0	9	714	2239	15934
α	Weekly inventory levels — 8 weeks					Weekly inventory levels — 12 weeks				
10%	0	0	0	180	3 924	167	48 290	79 368	119 575	166 982
20%	0	0	0	180	2 914	167	1 198	79 368	119 575	166 982
30%	0	0	0	180	4 414	167	1 198	79 258	119 575	166 982
40%	0	0	0	180	4 874	167	1 198	4 418	127 985	178 022
50%	0	0	0	180	5 684	167	1 198	4 498	69 467	94 814

TABLE 8.6: *Revised HLT & ND Inventory model results, $R = 30$: Maximum inventory level.*

Maximum inventory level: $R = 30$					
	Daily inventory levels — 30 days				
α \ ss	10%	20%	30%	40%	50%
10%	345 601	391 695	442 007	491 883	537 928
20%	349 852	399 402	448 402	496 206	541 876
30%	352 998	402 204	451 636	499 560	545 724
40%	358 620	405 661	455 427	503 474	550 517
50%	363 181	409 099	459 139	506 954	554 212
α	Weekly inventory levels — 4 weeks				
10%	319 584	355 449	397 567	439 399	514 066
20%	319 734	355 409	397 999	439 609	511 156
30%	319 784	355 657	399 079	440 892	509 904
40%	316 382	355 281	397 813	440 797	512 408
50%	317 700	355 277	399 980	442 425	513 882
α	Weekly inventory levels — 8 weeks				
10%	317 838	352 247	392 085	435 586	474 400
20%	317 988	352 357	392 245	435 736	474 652
30%	318 088	352 357	392 345	435 886	475 725
40%	319 592	352 475	393 417	436 235	476 484
50%	314 226	353 755	393 989	437 523	477 685
α	Weekly inventory levels — 12 weeks				
10%	319 929	402 494	474 208	555 197	647 983
20%	320 279	354 868	474 382	555 371	648 157
30%	320 539	355 118	468 338	555 487	648 273
40%	320 739	355 990	396 224	538 805	620 037
50%	315 564	357 164	396 814	466 544	514 167

8.1.3 Concluding remarks

The Revised HLT & Normal Distribution Inventory Model is fully capable of producing acceptable order quantities for public healthcare facilities which place stock orders once a week ($R = 7$). It is recommended to use 8–12 weeks of historic data, a safety stock of 20-30% and service level of $\alpha = 10\%$. When orders are only issued once a month ($R = 30$) it is recommended to use 8 weeks of historic data, with a safety stock of 40–50% and service level of $\alpha = 40\%$.

Considering that the original HLT & Normal Distribution Inventory Model was recommended a 30% safety stock in § 6.3.2, a 40% safety stock for the revised model is acceptable based on the results seen from the maximum inventory levels and total expired items. In fact, the Revised HLT & Normal Distribution Inventory Model has outperformed the original model in terms of both inventory levels and expired items, for both weekly and monthly review periods.

A summary of the recommended conditions for both the original and revised HLT & Normal Distribution Inventory Models is summarised in Table 8.7. This specifies what safety stock, service level and number of historic inventory levels to use based on the healthcare facility's review period (R) and available historic inventory levels.

TABLE 8.7: *Updated summary of the recommended inventory model conditions.*

Inventory Model	Recommended Conditions			
	Review Period (R)	Safety Stock (ss)	Service Level (α)	Historic Inventory Levels
HLT & ND	7	30%	20%	30 days, daily
	30	30%	20%	
Revised HLT & ND	7	20-30%	10%	8–12 weeks
	30	40–50%	40%	8 weeks

8.2 Subject Matter Expert Validation

All validation described in this section is qualitative feedback which has been acquired from subject matter experts (SMEs). Four SMEs were contacted to review the work done in this project. Each SME was presented with the work, then given the opportunity to ask questions and discuss future work. The same validation sheet was provided to each SME. Each SME had to answer four questions per section of work between "Strongly disagree" and "Strongly agree". The completed forms have been recreated in Appendix D.

SMEs 1–3

A visit was conducted to the Cape Medical Depot (CMD) where a a two-and-a-half hour long meeting was held with three of the depot's supply chain experts. One of these SMEs manage the acquisition of prescription drugs, and another one manages the acquisition of tools and equipment for the depot. All three SMEs have to make predictions to ensure that the storage capacity is not exceeded at any given time. Chapters 6 and 7 were presented in detail. Additionally, the means of adapting the inventory model to allow weekly stock counts (§ 8.1) was discussed. The feedback obtained from these three subject matter experts are shown in Tables D.1–D.3, Appendix D.

SME 4

The fourth SME is a supply chain specialist with more than ten years experience in stock acquisition, working at a distribution company in Cape Town. This SME was introduced through an mutual friend working at the same company. This company has been struggling with their many suppliers for some time. Orders seldom get delivered in full and often arrive late. This incurs hefty penalty fees on the company for being understocked when they are expected to sufficiently supply their customers. This mimics the supplier behaviour described in Chapter 7. Although this distribution company is not a healthcare facility, the same need exists for lead time predictions (§ 7.1.1), supplier prioritisation (§ 7.1.3) and allocation of order quantities (§ 7.2). The validation feedback is shown in Table D.4, Appendix D.

8.2.1 The validation feedback

The three SMEs from the CMD provided feedback on the inventory policy models and SME 4 from the distribution company solely provided feedback on the storage limitations section. In all cases the validation feedback indicated that the work had been effectively explained and made easy to understand. The SMEs found all of the work to be realistic and accurately depict the behaviour of South Africa's public healthcare facilities, especially hospitals.

SME 4 stated that the problems addressed in Chapter 7 perfectly captures what occurs at their distribution centre. The work done in § 7.1 and 7.2 will be applied at their distribution centre in the near future to predict their suppliers' behaviour and increase the confidence of their orders. Additionally, members aiding in the development of the Stock Visibility System (SVS) have reach out and would like to see the Revised HLT & Normal Distribution Inventory Model (§ 8.1.1) tested in a pilot study[1], which may occur during 2020.

The validation for the original HLT & Normal Distribution Inventory Model (§ 6.3) is summarised in Figure 8.10. Although there was no consensus on whether this model requires improvement for determining order quantities, all three SMEs believe that the model has the potential to help with future supply chain operations at healthcare facilities, when the infrastructure to capture daily inventory levels is available.

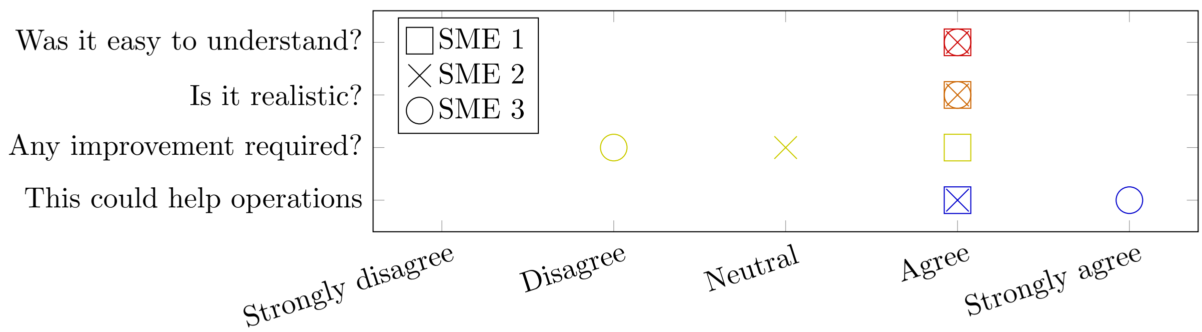

FIGURE 8.10: *Scatter plot showing the feedback spread from three subject matter experts on the HLT & ND Inventory Model (assuming* ***daily*** *historic demand values) (§ 6.3).*

The Revised HLT & Normal Distribution Inventory Model (§ 8.1.1) was also validated. However, at the time of validation the model had only been tested for 4 weeks of historic data, which was not optimal. It is now clear, as concluded in § 8.1.3, that the model performs better

[1] A small scale application of the work used to evaluate the outcomes and improve upon the design prior to a full-scale launch.

with 8–12 weeks of historic data. The validation, using only 4 weeks of historic inventory levels, is summarised in Figure 8.11. There was no consensus on whether this model requires improvement, but all three SMEs believe that this model can assist in supply chain operations at healthcare facilities, provided that weekly inventory levels are readily available.

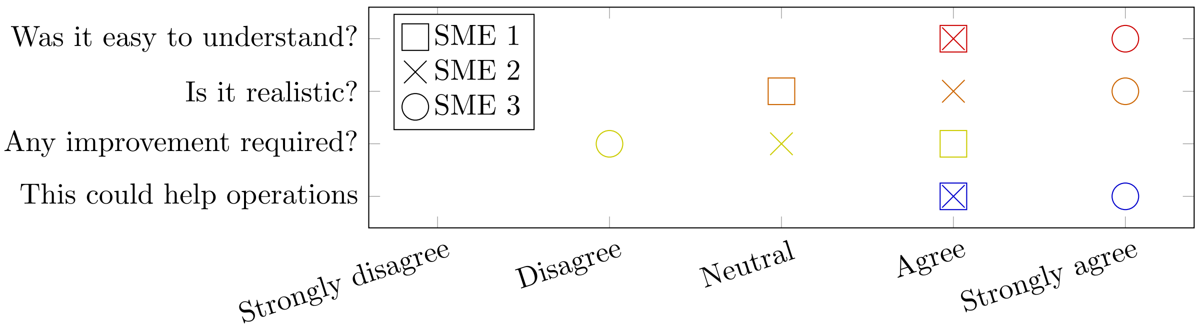

FIGURE 8.11: *Scatter plot showing the feedback spread from three subject matter experts on the Revised HLT & ND Inventory Model (assuming* ***weekly*** *historic demand values) (§ 8.1.1).*

The section of work which prioritises suppliers (§ 7.1) was validated as a whole. In other words, the manner in which the three specific subsections ('Varying Lead times', 'Order cost' and 'Choosing the optimal supplier') came together in order to construct the supplier priority list was validated. The results of this validation are summarised in Figure 8.12. All four SMEs deal with suppliers and thus provided feedback.

The SMEs agree that prioritising suppliers before issuing orders can help improve the confidence of receiving the desired order quantities during supply chain operations. One SME from the CMD wanted to see improvement to the model. However, they later mentioned that this was due to many healthcare facilities not logging their orders correctly, causing problems. It is not the model itself which requires improvement, but the infrastructure necessary to utilise it.

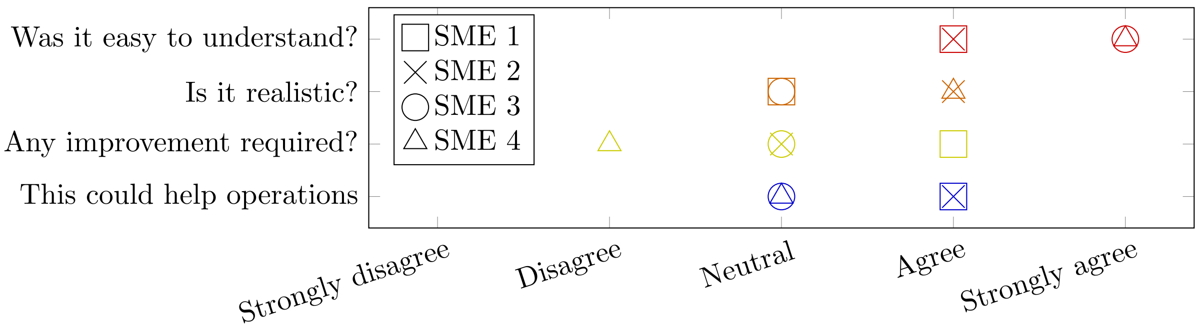

FIGURE 8.12: *Validation from SMEs: Supplier prioritisation — As a whole (§ 7.1).*

The validation results for the varying lead times section (§ 7.1.1) is summarised in Figure 8.13. Two SMEs stated that some improvement could be done, but this was curiosity directed towards further investigation for using historic lead time information. The sections on order cost (§ 7.1.2) and choosing optimal suppliers (§ 7.1.3) both received positive feedback, not requiring any improvement to the model and indicating that the work done could help with supply chain operations. The feedback is summarised in Figures 8.14 and 8.15 respectively.

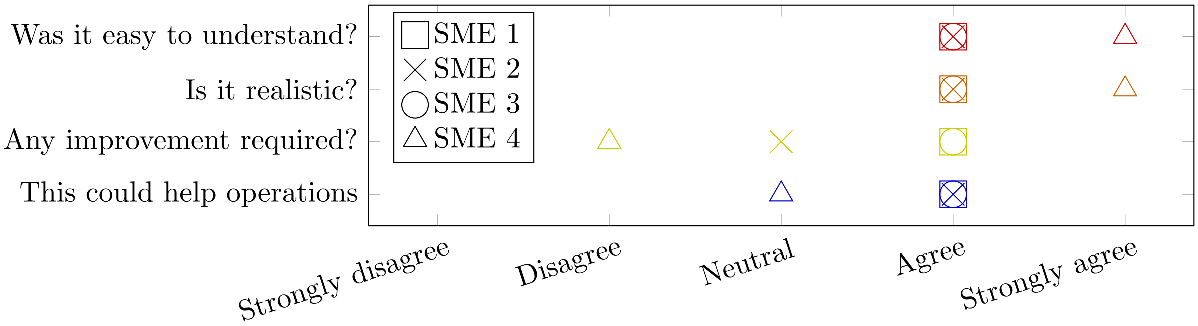

FIGURE 8.13: *Validation from SMEs: Supplier prioritisation — Varying lead times (§ 7.1.1).*

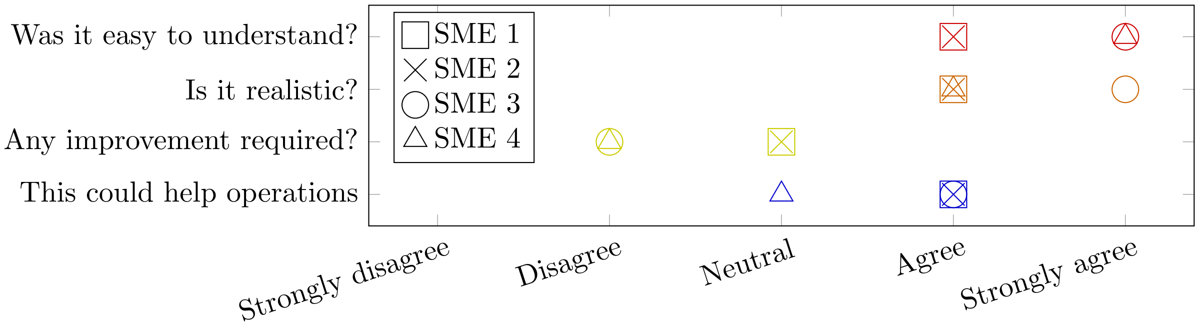

FIGURE 8.14: *Validation from SMEs: Supplier prioritisation — Order cost (§ 7.1.2).*

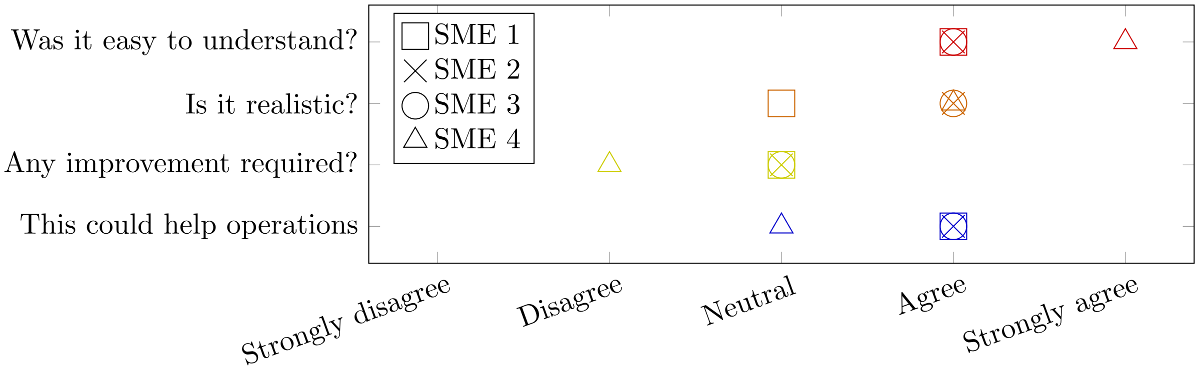

FIGURE 8.15: *Validation from SMEs: Supplier prioritisation — Choosing optimal supplier (§ 7.1.3).*

The section addressing order irregularities (§ 7.2), which determines how to divide an order quantity between the suppliers to increase the confidence of acquiring the full amount, received good reviews. The feedback is summarised in Figure 8.16. SME 1 thought some improvement to the work could be done, but did not describe what nor why. However, SME 4 was so confident that they wish to implement this work at their distribution company.

Only SME 4 provided feedback on the limitations section (§ 7.3), having taken the time to read the work. This work attempts to ensure that new orders will not cause a warehouse to exceed its storage capacity. Order quantities are adjusted by incrementally changing the priority level, effecting each product's order quantity based on the categorization, which specifies how crucial the product is at the facility. The feedback is summarised in Figure 8.17.

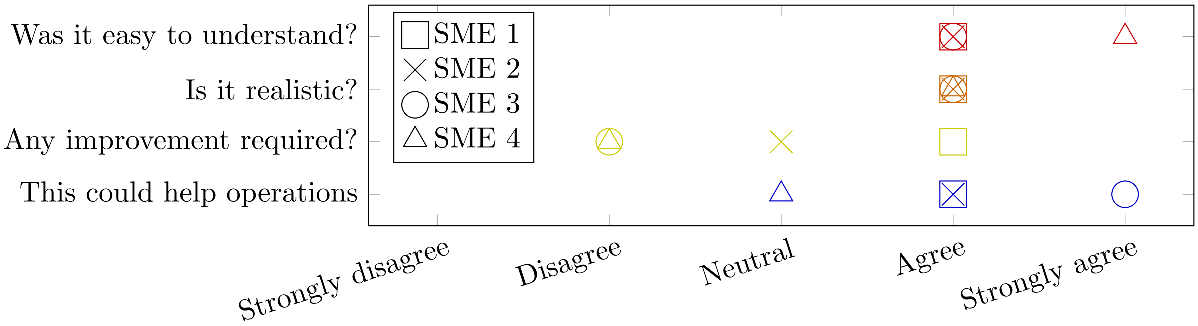

FIGURE 8.16: *Validation from SMEs: Order irregularities (§ 7.2).*

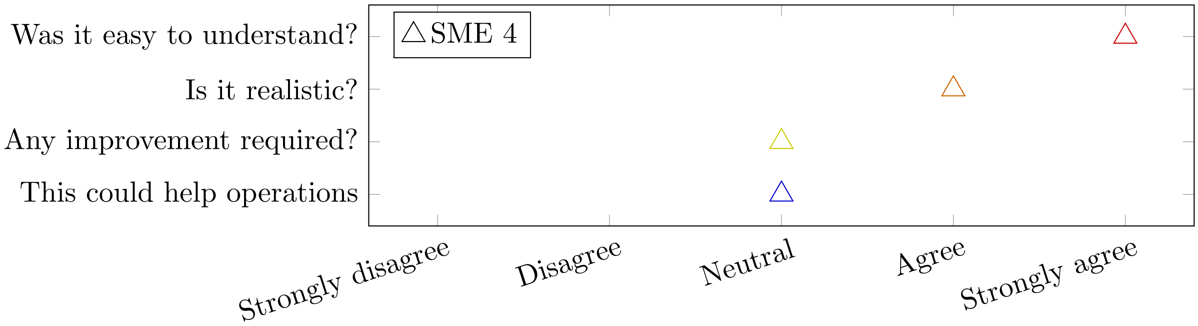

FIGURE 8.17: *Validation from SMEs: Storage limitations (§ 7.3).*

8.2.2 The validation summary

The SME validation results described in Figures 8.10–8.15 have been tallied by question and plotted in Table 8.8. A box and whisker diagram of these results has been plotted in Figure 8.18 to describe the distribution of the total SME validation. This describes the total results for all work validated during § 8.2.1. It is conclusive that the models were understood by the SMEs, improving the confidence of their reviews. The SMEs found all the models to be realistic potentially helpful towards current supply chain operations. Although most of the feedback did not indicate a need for improvement, there does exist a desire to see the models developed further. This would likely be achieved during one of the pilot tests, which would have these new theoretical models refined for real-world implementation.

TABLE 8.8: *Summary of the SME validation feedback.*

Question	Strongly disagree	Disagree	Neutral	Agree	Strongly agree	Total
Was it easy to understand?	0	0	0	18	9	27
Is it realistic?	0	0	4	20	3	27
Any improvement required?	0	9	12	6	0	27
This could help operations	0	0	7	17	3	27

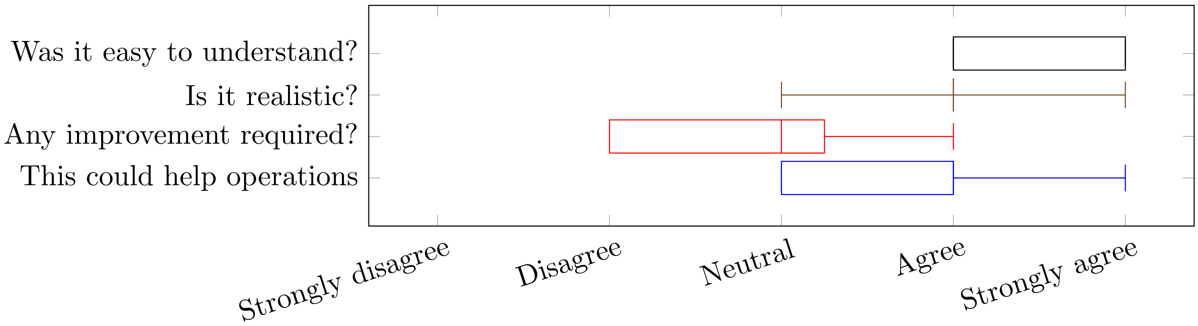

FIGURE 8.18: *Box and whisker diagram of the total SME validation.*

The box and whisker diagram helps to visualise the distribution of the SME feedback along the validation scale ("Strongly disagree"–"Strongly agree"), but does not necessarily describe the magnitude of the results. Pie charts use size to quickly describe the relationship between percentages. The pie charts in Figure 8.19 describe the extent of answers obtained for each question from Table 8.8. Figures 8.19a, 8.19b and 8.19d show a majority consensus to support the models and a zero percentage disagreement thereof.

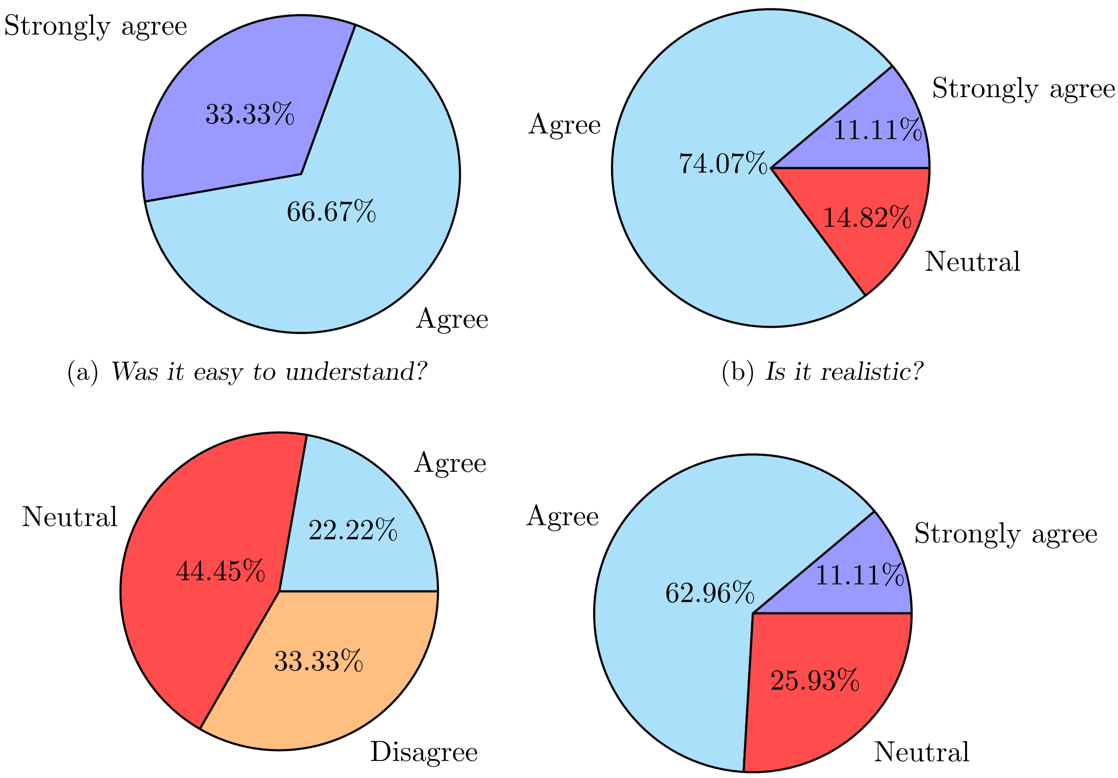

(a) *Was it easy to understand?* (b) *Is it realistic?* (c) *Any improvement required?* (d) *This could help operations*

FIGURE 8.19: *Pie charts summarizing the SME validation feedback.*

8.3 Project Progress

This chapter has completed Objective VIII, the validation of the final model developed in this thesis, and the first target in the *Conclusion phase* of the project methodology framework (§ 2.2). The validation was not limited to the inventory ordering policy, but included the work which may help improve the confidence of receiving orders. The updated framework is shown in Figure 8.20.

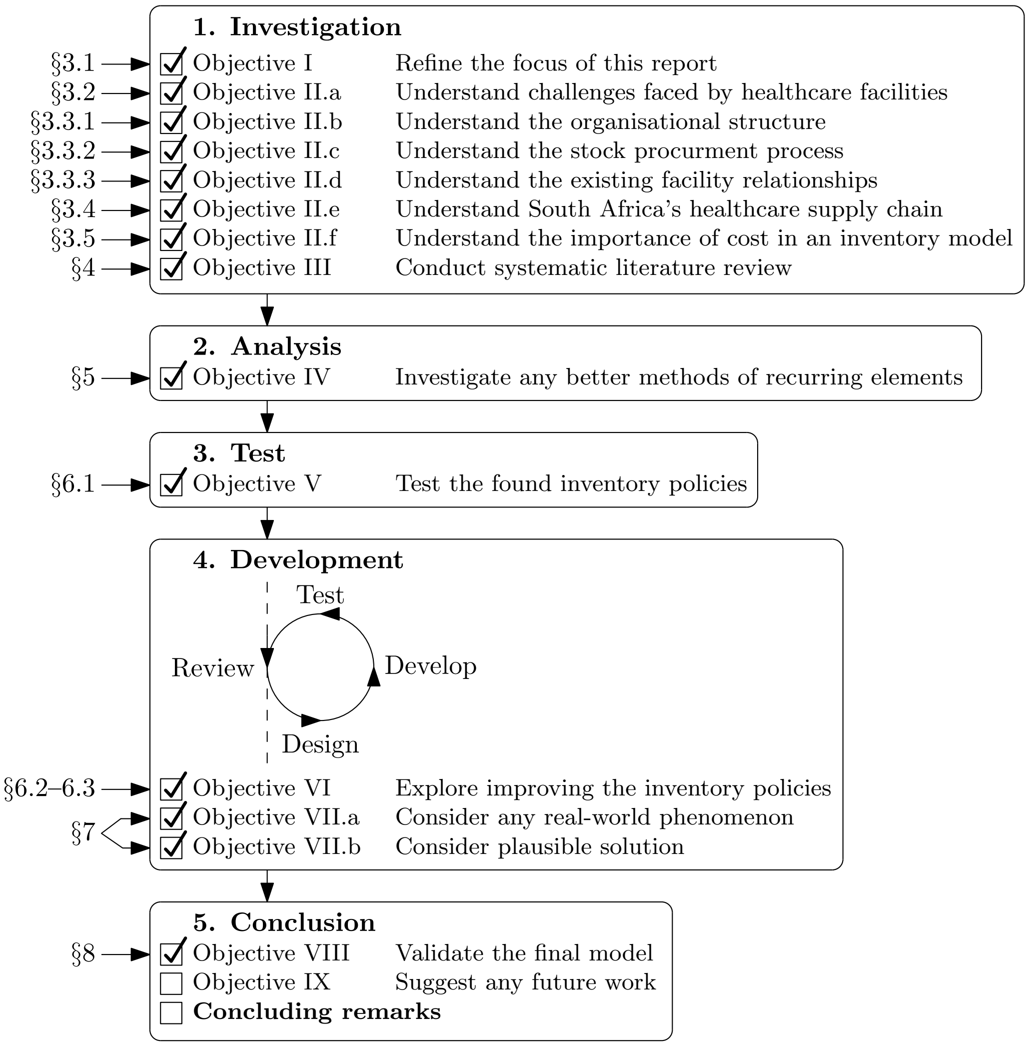

FIGURE 8.20: *Project methodology framework: Chapter 8.*

CHAPTER 9

Conclusion

Contents

This chapter will provide a summary of the work covered in each chapter, an appraisal of the work done, describe some contributions which the work has achieved and suggest some possible future work. The last two targets in the *Development phase* of the project methodology framework (§ 2.2) will be addressed. These targets aim to achieve Objective IX from the project objectives list (§ 1.4) and provide some final concluding remarks on the project.

9.1 Thesis Summary

Chapter 1 provided a short background describing the lack of stock visibility that physicians have at public hospitals and clinics across South Africa, despite their important role in issuing the orders required to sustain the healthcare facility. Once the problem was defined, a list of objectives were created which would strive towards creating an inventory policy model suitable for South Africa's public healthcare facilities and improving the confidence of placing stock orders. The project scope was focused at public facilities, but included the private sector in the investigation.

Popular methodology structures were discussed in Chapter 2 and a final project methodology framework was created to direct the flow of research when conducting the study. This framework was briefly revised and updated at the end of each chapter to track progress.

Chapter 3 began with considering the level of research that will be conducted and the importance of understanding how this work may be used to improve the situation of possible end users. This helped to refine the focus of the project, necessary for Objective I. A field study was conducted at five site visits and with the aid of discussions with three final year medical students to understand the conditions of South Africa's healthcare facilities. The healthcare organisational structure, procurement process and relationships seen at healthcare facilities was described. The conditions of South Africa's healthcare supply chain was outlined using information from the National Department of Health and findings from the site visits. The impact of trying to minimize cost was investigated and it was concluded that the primary objective should always be to meet the full demand. After meeting demand, money can be saved by reducing the number of expired items due to their high disposal cost. This concluded the six parts of Objective II.

A systematic literature review (SLR) was conducted in Chapter 4. This study attempted to find the best literature on solutions in healthcare facilities, such as inventory policies, decision support systems, improved organisational structures and stock management. The most significant effect on order quantities were the inventory policies. These policies were described mathematically, concluding Objective III and ending the Investigation phase of the project methodology framework. All literature on healthcare inventory policies were studies conducted overseas (outside of Africa) and it was observed that only one method was used to estimate future demand, the *Moving Average* forecast method. This was the most promising area to investigate better methods of application, for Objective IV.

Chapter 5 conducted a study on popular alternative forecast methods, addressing Objective IV. The study concluded that Holt's Linear Trend (HLT) method was most suited due to its impressive accuracy scores and low computational times. This ended the Analysis phase of the project methodology framework.

Chapter 6 began by testing the inventory policies found during the SLR. None of the models were able to produce good orders under the conditions of South Africa's public healthcare supply chain. This completed Objective V and the Test phase of the project methodology framework. These results also supported the need to improve on the inventory policies, as planned for Objective VI. Using the better forecast technique, the HLT method, the Iterative Forecast Inventory Model was created. This model stepped through the predicted forecast and captured how the inventory behaved, considering information on expected order arrivals and each product's shelf life. This method performed much better than the SLR inventory policies, but allowed decreasing forecast trends to under-assume demand levels. The HLT & ND Inventory Model was created to correct this flaw by using a normal distribution fit over the historic demand set to define a minimum demand level for the forecast. This model performed very well at both fully meeting demand and eliminating expiring items, achieving Objective VI.

Theoretically, the product order quantities have been solved. However, the effects of real-world supplier behaviour threatened acquiring these desired quantities. Chapter 7 investigated the reality of supplier performance and how these occurrences can be utilised to improve the confidence of receiving the order quantities which get placed. This was achieved by first prioritising which supplier should be ordered to first based on delivery time and cost. Thereafter, a paramount order quantity is calculated for each supplier which defines the maximum amount expected to be delivered, for some level of confidence. Additionally, a solution was suggested for keeping inventory levels from exceeding storage capacity using a priority matrix. More essential products will receive a higher categorization. The priority level is gradually changed, scaling the size of product orders based on their categorization until the estimated capacity size after delivery is within the storage limitations. This achieved both parts of Objective VII, completing the Development phase of the project methodology framework.

Chapter 8 acquired validation on the work done, as required for Objective VIII. Current public healthcare facilities are not yet ready to capture daily inventory levels, but effort is being made to ensure weekly inventory levels will be consistently recorded in the near future. The Revised HLT & ND Inventory Model was designed, taking into account all the scenarios which may occur when only capturing weekly inventory levels. The results appeared promising. The models developed and improved in Chapters 6–8 were presented to four subject matter experts (SMEs), who provided qualitative validation on the work. All the feedback was positive. One SME working as a supply chain specialist at a distribution company in Cape Town has requested the work done in Chapter 7 to be applied at their facility. Additionally, members aiding in the development of the Stock Visibility System would like to see the Revised HLT & Normal Distribution Inventory Model tested in a pilot study in the near future.

This chapter, Chapter 9, will achieve the final two parts required for the Conclusion phase in the project methodology framework, bringing this research project to an end.

9.2 Thesis Appraisal

The work conducted in this thesis has attempted to solve a range of problems commonly experienced during the process of acquiring stock at public hospitals and clinics. This section will describe both the strengths and weaknesses of this work.

9.2.1 Strengths

Before proceeding with any work on the models, a thorough understanding of the conditions inconveniencing public healthcare facilities, and the physicians that work in them, was conducted. This inspiration came from the Three Worlds Framework (§ 3.1) which highlighted the importance of knowing whom the research is for. The site visits conducted during this project, the input from medical students and the feedback obtained from subject matter experts all supported the need for a new inventory policy model capable of confidently producing and acquiring order quantities.

In the pursuit of a feasible solution, several individual models were created. Each model is designed to resolve a unique problem contributing to the confident acquisition of stock. The original intent was to produce a single model solely for use at healthcare faculties, downstream of the supply chain. However, a few models were individually created. Each model addressed a different topic, but contributed towards placing more confident orders. A representation of how these models link together is shown in Figure 9.1.

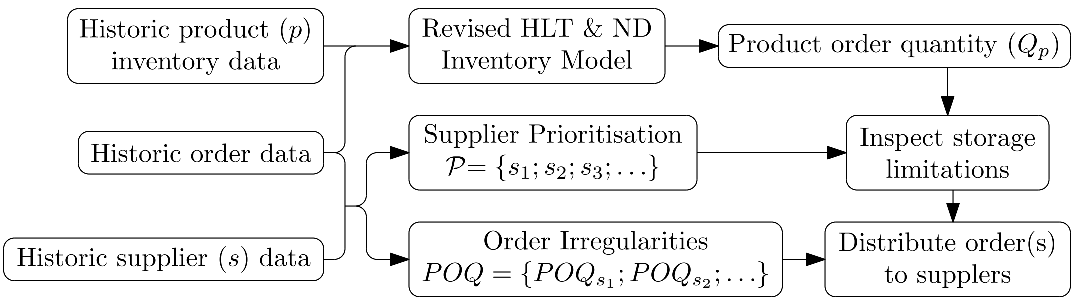

FIGURE 9.1: *Using thesis models to place more confident orders.*

It was never the intent to influence the behaviour or performance of upstream entities in the supply chain, as defined by the scope in § 1.5. However, some of the work done can be applied further upstream as well to assist in supply chain operations.

Figure 9.2 describes where the four primary models of this report can be applied throughout the supply chain to improve operations. The Revised HLT & ND Inventory Model was designed for facilities under downstream conditions. Without any testing, there is no certainty that this model will perform well for provincial depots. Medical depots are attempting to reduce the number of direct deliveries from manufacturers to healthcare facilities. However, many contracts will remain which allow a healthcare facility to choose between suppliers.

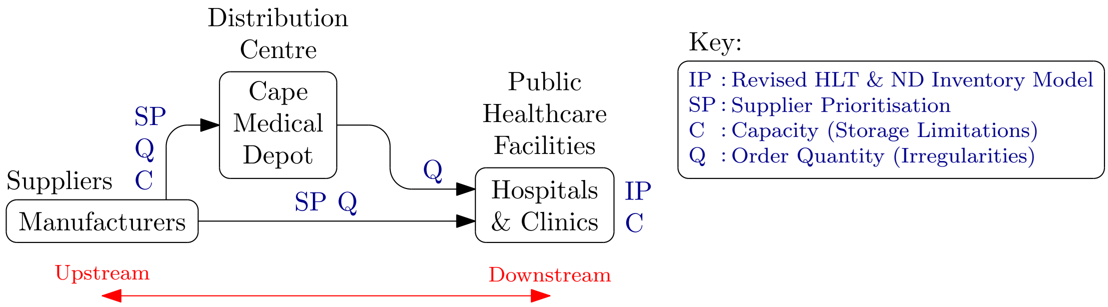

FIGURE 9.2: *Application of thesis work in the healthcare supply chain.*

9.2.2 Weaknesses

The inventory policy models in this paper have been tested for an environment which uses ideal First-In, First-Out. The model also assumed the expiry date of batches to be one shelf life from the day of a delivery's arrival. This was an acceptable assumption due to the large review periods, which separate orders by 7–30 days, reducing the likelihood of a later-batch being manufactured earlier. However, supplier warehouses and distribution centres may not always apply First-In, First-Out perfectly when satisfying the demand of healthcare facilities. Additionally, some batches of stock may have been waiting longer in storage before delivery than others. The real expiry date is in fact determined by the date of manufacture, which can only be inspected by the healthcare facility upon acquisition.

The Revised HLT & ND Inventory Model was created to accommodate the capture of weekly inventory levels as apposed to the original daily inventory levels design. Scenarios were considered that may result from capturing data in this fashion. However, the revised model did not consider inventory counts being performed earlier, or later, than the planned date. This would cause the number of days between inventory counts to vary. The historic demand set, $\mathcal{D}$, would then include a variety of demands not all resulting from seven days. This would effect all *Scenarios C–G* (pages 78–80) which use the inverse of the cumulative normally distributed value of the set and assume consistent seven day intervals.

9.3 Thesis Contributions

By taking the time to communicate with physicians and medical students at the start of the project, the causes for over- and understocking were identified. The importance of meeting all demand without causing an excessive amount of expired items was recognized. The investigation chapter identified that foreign inventory policies found it difficult to accommodate the long lead times and review periods experienced in South Africa's healthcare supply chain. An original inventory model was created to perform under these conditions. This model performed well assuming the availability of daily inventory levels. This same model was then adapted to use weekly inventory levels and showed promising results. This research is unique and provides an alternative to existing software solutions owned by companies, whom have not yet proven to be a successful solution in South Africa's healthcare supply chain.

Subject matter experts have shown interest in furthering the work of the inventory model and conducting a pilot test at one, or more, healthcare facilities in the near future. These models address the first two points of the three key focus efforts described in § 3.6. The inventory model is capable of producing ideal order quantities for all the listed products at the healthcare facility in only a few seconds, which saves time for physicians to tend to their patients and rest.

Models were created to improve the confidence of orders by evaluating the historic performance of suppliers. These original models contribute a new way of reviewing supplier behaviour and has caught the attention of specialists whom wish to implement and test the work in a real-world distribution company.

9.4 Future Work

A lot has been achieved during the course of this thesis. However, there still exists many possibilities to take this research further. This section will suggest ideas for future work, by briefly providing an explanation of each.

9.4.1 Implementation into healthcare facilities

I The Revised HLT & ND Inventory Model should be improved to allow for varying review dates, rather than relying on a fixed one of exactly seven days. This was explained when describing the thesis weaknesses, § 9.2.2.

II Not many public healthcare facilities have a well-designed functioning local database system. A design must be created which will allow healthcare facilities to capture all the necessary information required to run the desired models. Additionally, the design may consider allowing existing local database systems to be converted with little effort.

III To achieve the best results, the inventory models require First-in, First-out (FIFO) to be enforced at each healthcare facility. An investigation into implementing FIFO at facilities should be done, taking into account that the best solution should allow physicians to acquire stock effortlessly. If physicians have to waist time requesting or finding stock, then it is likely that hoarding will occur throughout the facility, as seen during the first public hospital site visit (§ 3.2.1), reducing visibility and resulting in more expired items. Possible solutions may include:

   (a) Shelving design organisation,

   (b) Kanban-like methodologies, or

   (c) Decision support systems (DSS).

IV A pilot test of the final inventory model should be conducted at one, or more, real-world healthcare facilities. The model should be tested against existing electronic stock management software, such as the RxSolution [85].

V Capturing inventory levels can be achieved by manually counting the stock on hand at the end of the day, or it can be monitored by updating the database whenever an item gets used or sold. There are many healthcare facilities which cannot afford the luxury of scanners or computers. This is where the Stock Visibility System (SVS) can make a difference. However, an internet connection is required to use the SVS. There will always be some form of infrastructure required. A detailed investigation can be done on of the levels of infrastructure available, considering the costs and benefits of each.

VI The supplier prioritisation model uses sets of historic lead times, achieved by each supplier, to analyse the expected performance. The number of historic entries in these sets may change the outcome of the test, affecting the priority list. A study can be conducted, testing the model for different set sizes, to find the optimal number of historic entries. This was suggested at the start of § 7.1.1 and at the end of § 7.1.4 to drive supplier behaviour.

9.4.2 Additional investigations

I During § 3.5 the costs involved with acquiring and holding stock at healthcare facilities was investigated. This is solely from the perspective of the healthcare facility. The purchase cost of each contract from the Master Procurement Catalogue, which includes the cost of delivery, have been calculated to accommodate the common ordering schedule of public healthcare facilities; once a week, or once a month. This also allows the suppliers to better plan their production, holding and distribution. An in-depth study can be conducted on optimising the total cost imposed on the supply chain, both in terms of monetary value and time.

II If work is being done to increase stock visibility throughout the supply chain, then research can be done on utilising data from further upstream in the supply chain. If a healthcare facility is able to monitor the stock availability at distribution centres, and possibly even suppliers, then it may be possible to make even better decisions. A test can be carried out to determine this change in performance and weigh the advantages with the disadvantages of using information sharing.

III Public healthcare facilities have very little stock visibility. This makes it difficult to find the desired item, or to know whether it is still in stock. Stock visibility requires regularly updated systems to deliver the best results. It is difficult to create an inexpensive solution to providing stock visibility, but even simplistic dashboards showing the most recent inventory data could make a difference.

IV Logging stock information by pen and paper will allow healthcare facilities to manage the admin side of inventory. However, this information becomes lost in books and folders with no access to readily use it for better decision making. However, manually entering stock data into a computer can sometimes be even slower. An investigation into inexpensive, fast and accurate ways of capturing stock behaviour could improve the performance of inventory systems. It would also assist the stock visibility problem described in the Future Work VII.

9.5 Project Progress

Objective IX has been achieved in this chapter, § 9.4, completing the second target in the *Conclusion phase* of the project methodology framework. The final objective, concluding remarks on the report, were covered in § 9.1–9.3. All thesis objectives have successfully been achieved and the project framework has been executed seamlessly. The completed project framework is presented in Figure 9.3.

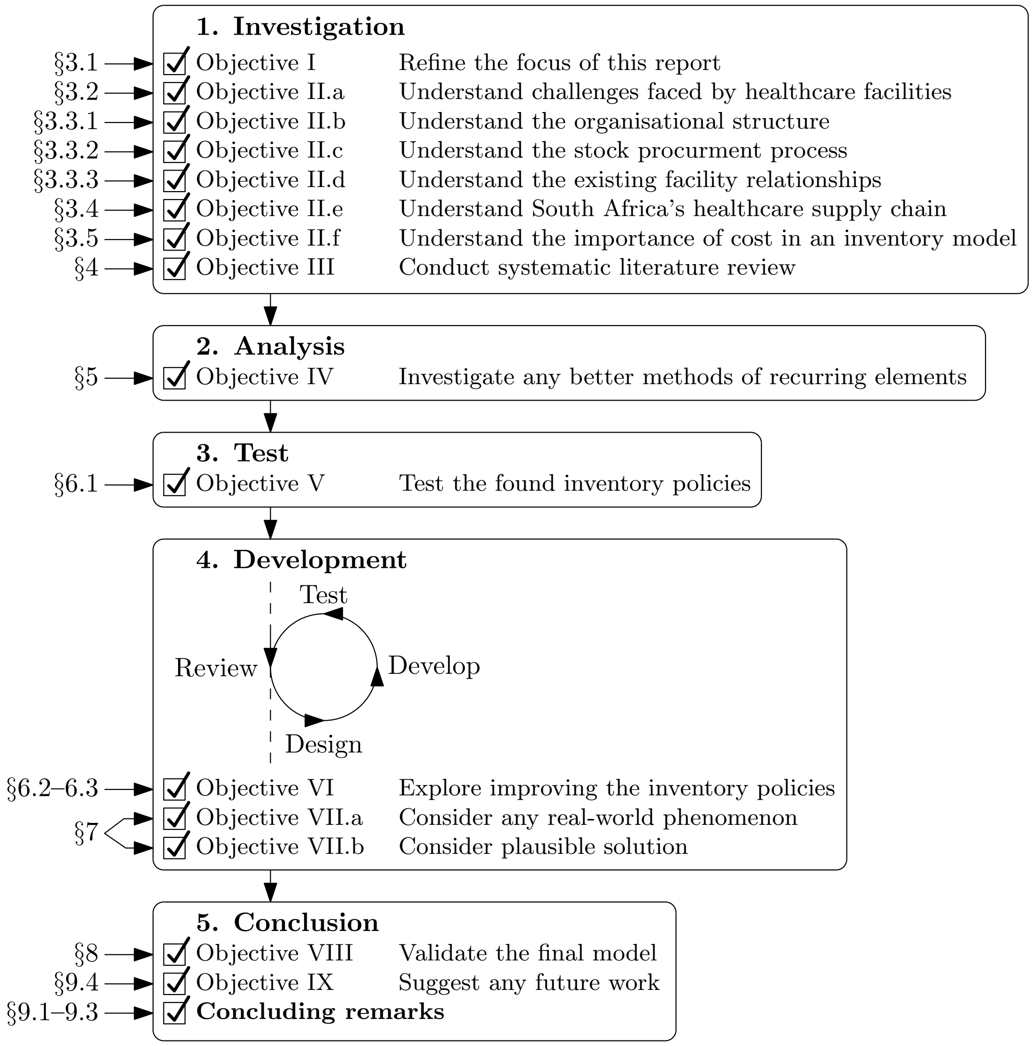

Figure 9.3: *Project methodology framework: Chapter 9.*

References

[1] Abramowitz PW, 1984, *Controlling financial variables - purchasing, inventory control, and waste reduction*, **41(17)**, pp. 309–317.

[2] AccountingExplained, 2013, *Days' Inventory on Hand Ratio*, Accessed: 2018-07-06.

[3] Altman DG & Gardner MJ, 1988, *Statistics in Medicine: Calculating confidence intervals for regression and correlation*, British medical journal (Clinical research ed.), **296(6631)**, pp. 1238.

[4] Asllani A, Culler E & Ettkin L, 2014, *A simulation-based apheresis platelet inventory management model*, **54(1)**, pp. 2730–2735.

[5] Avert, 2018, *HIV and AIDS in South Africa*, Accessed: 2018-06-16, URL: `https://www.avert.org/professionals/hiv-around-world/sub-saharan-africa/south-africa`.

[6] Awaya T, Ohtaki K.-i, Yamada T, Yamamoto K, Miyoshi T, Itagaki Y.-i, Tasaki Y, Hayase N & Matsubara K, 2005, *Automation in Drug Inventory Management Saves Personnel Time and Budget*, **125(5)**, pp. 427–432.

[7] Axsäter S, "Forecasting", in: *Inventory Control*, 8th ed., Sweden: Springer, 2015, chap. 2, pp. 7–16.

[8] Ballentine R, Ravin RL & Gilbert JR, 1976, *ABC inventory analysis and economic order quantity concept in hospital pharmacy purchasing*, **33**, pp. 552–555.

[9] Barron P, Padarath A, Ranchod S, Adams C, Burger R, Carvounesi A, Stewart J, Dreyer K, Smith A & van Biljon C, 2017, *Twenty years of the South African Health Review; South Africa's hospital sector: old divisions and new developments*, Health Systems Trust, Durban.

[10] Budgen D & Brereton P, 2006, *Performing systematic literature reviews in software engineering*, Proceedings of the Proceedings of the 28th international conference on Software engineering: ACM, pp. 1051–1052.

[11] Cancaster B, 2019, *SELECTING SYRINGES AND NEEDLES*, Accessed: 2019-07-23, URL: `https://www.vitalitymedical.com/blog/selecting-syringes-and-needles.html`.

[12] Chopra S & Meindl P, "Managing Economies of Sale in a Supply Chain: Cycle Inventory & Managing Uncertainty in a Supply Chain: Safety Inventory", in: *Supply Chain Management: Strategy, Planning, and Operation*, 6th ed., Boston Pearson, 2016, chap. 11 & 12, pp. 312–357.

[13] Chowles T, 2016, *NDoH Deploys SVS mHealth App to Prevent Stock-outs - eHealth News ZA*, Accessed: 2018-03-16, URL: `http://ehealthnews.co.za/ndoh-svs-mhealth-stock-outs/`.

[14] Chowles T, 2015, *Vodacom Business Tackling ARV Stockouts - eHealth News ZA*, Accessed: 2018-03-16, URL: http://ehealthnews.co.za/vodacom-business-arv-stockouts/.

[15] Clarivate, 2019, *Web of Science*, Accessed: 29/04/2019, URL: http://apps.webofknowledge.com/WOS%7B%5C_%7DGeneralSearch%7B%5C_%7Dinput.do?product=WOS%7B%5C&%7Dsearch%7B%5C_%7Dmode=GeneralSearch%7B%5C&%7DSID=F6Zw7cyjxUeuQtjkAVg%7B%5C&%7DpreferencesSaved=.

[16] Cooke MW, Higgins J & Kidd P, 2003, *Use of emergency observation and assessment wards: a systematic literature review.*

[17] Coovadia H, Jewkes R, Barron P, Sanders D & McIntyre D, 2009, *The health and health system of South Africa: historical roots of current public health challenges*, Elsevier Ltd.

[18] Cordenier C, 2011, *To what extent is the Panopticon a suitable analogy for power?*, Accessed: 2018-03-15, URL: http://www.e-ir.info/2011/03/02/to-what-extent-is-the-panopticon-a-suitable-analogy-for-power/.

[19] Council for Medical Schemes, 2016, *Annual Report 2015–2016*, (Unpublished) Technical Report, Council for Medical Schemes, Pretoria.

[20] Davis AM, Bersoff EH & Comer ER, 1988, *A strategy for comparing alternative software development life cycle models*, **14(10)**, pp. 1453–1461.

[21] Day C & Gray A, 2016, *Health and Related Indicators.*

[22] Dotoli M, Fanti MP, Iacobellis G, Martino L, Moretti A & Ukovich W, 2010, *Modeling and management of a hospital department via Petri nets*, Proceedings of the, pp. 1–6.

[23] Econex, 2015, *Private hospital expenditure and relation to utilisation – observations from the data*, (Unpublished) Technical Report, Econex, Stellenbosch.

[24] Elsevier, 2019, *Scopus*, Accessed: 2019-04-29, URL: https://www-scopus-com.ez.sun.ac.za/search/form.uri?display=basic.

[25] Esmaili N, Norman B & Rajgopal J, 2015, *A Heuristic for Selecting Multi-item Inventory Review Policies*, Proceedings of the, pp. 1050–1060.

[26] Fernandes R, Gouveia JB & Pinho C, 2010, *Overstock – A Real Option Approach*, **3(2)**, pp. 98–107.

[27] Finance Train, 2018, *Inventory Turnover and Days of Inventory on Hand (DOH)*, Accessed: 2018-07-06, URL: https://financetrain.com/inventory-turnover-and-days-of-inventory-on-hand-doh/.

[28] Foucault M, 1995, *Discipline & Punish The Birth of the Prison*, 2nd Edition, Foucault, Michel, New York, USA.

[29] Francher DO & Beasley DB, 1981, *Cost cutting through the EOQ.*

[30] Gardner ES & McKenzie E, 1985, *Forecasting Trends in Time Series*, **31(10)**, pp. 1237–1246.

[31] Gardner ES & McKenzie E, 2009, *Why the damped trend works*, **62(6)**, pp. 1177–1180.

[32] Gebicki M, Mooney E, Chen SJ(& Mazur LM, 2014, *Evaluation of hospital medication inventory policies*, **17(3)**, pp. 215–229.

[33] Gibson D, 2004, *The gaps in the gaze in South African hospitals*, **59(10 SPEC. ISS.)**, pp. 2013–2024.

[34] Hart G, 1996, *The five W's: an old tool for the new task of audience analysis*, Technical Communication, **43(2)**, pp. 139–145.

[35] Hyndman R, 2018, *Forecasting Using R*, (Unpublished) Technical Report, DataCamp.

[36] Iannone R, Lambiase A, Miranda S, Riemma S & Sarno D, 2015, *Cost savings in hospital materials management: look-back versus look-ahead inventory policies*, **22(1)**, pp. 60.

[37] Ibanez L, 2017, *A quick overview to Agile development – Medium*, Accessed: 2018-05-07, URL: `https://medium.com/@LazaroIbanez/a-quick-overview-to-agile-5c87ffc9e0f2`.

[38] Jensen PA & Bard JF, "Inventory Theory - The (s, Q) Inventory Policy", in: *Operations Research Models and Methods*, 1st ed., Australia: John Wiley and Sons, 2003, chap. 5, pp. 1–4.

[39] Jonrinaldi, Primadi MY & Hadiguna RA, 2017, *A joint inventory policy under permissible delay in payment and stochastic demand (Case study: Pharmacy Department of Pariaman Hospital)*, (Unpublished) Technical Report.

[40] Kachwee M & Hartmann D, 2013, *Hospital Supply Chain Management and Optimisation*, Proceedings of the 1th SAIIE25 Proceedings, Stellenbosch, Cape Town, South Africa, pp. 1–12.

[41] Kelle P, Woosley J & Schneider H, 2012, *Pharmaceutical supply chain specifics and inventory solutions for a hospital case*, **1(2-3)**, pp. 54–63.

[42] Khan K, Kunz R, Kleijnen J & Antes G, 2003, *Systematic Reviews to Support Evidence-based Medicine: how to review and apply findings of healthcare research*, 1st Edition, Royal Society of Medicine Press, London.

[43] Kingston RL, Pfeifle CE, Cipolle RJ, Zaske DE, Eells CB & Johnson HL, 1982, *Comprehensive Pharmaceutical Services for a State Correctional Facility*, American Journal of hospital Pharmacy, **39(9)**, pp. 86–89.

[44] Kinyon K, 2001, *The Panopticon Gaze in Kundera's The Unbearable Lightness of Being*, **42(3)**, pp. 243–251.

[45] Kitchenham B, 2004, *Procedures for performing systematic reviews*, **33(TR/SE-0401)**, pp. 1–28.

[46] Knowles JD & Corne DW, 2000, *Approximating the nondominated front using the Pareto Archived Evolution Strategy*, Evolutionary Computation, **8(2)**, pp. 149–172.

[47] Leon A & Koch AS, 2004, *Agile Software Development Evaluating the Methods for Your Organization*, Artech House, Inc., Norwood, MA, USA.

[48] Levit A, 2014, *5 Reasons Hierarchy Is Killing Your Culture - OPEN Forum*, Accessed: 2018-03-28, URL: `https://www.americanexpress.com/us/small-business/openforum/articles/5-reasons-hierarchy-is-killing-your-culture/`.

[49] Llopis Roca F, Juan Pastor A, Ferré Losa C, Martín Sánchez F, Llorens Soriano P, Sempere Montes G, Navarro Bustos C & Martínez Ortiz De Zárate M, 2014, *Register of short-stay units in Spain, the REGICE project, study 1: Location, organizational features, and staffing*, **26(1)**, pp. 57–60.

[50] Lu He JL, Cuevas JLT, Kokash MZ, Banimustafa EA & Khasawneh MT, 2017, *Optimization in Pharmaceutical Supply Chain Inventory management for Disaster Planning*, Proceedings of the 67th Annual Conference and Expo of the Institute of Industrial Engineers 2017, pp. 1583–1588.

[51] Lyon T, 2012, *Two Bin Supply System at St Clair Hospital - YouTube*, Accessed: 2018-04-26, URL: https://www.youtube.com/watch?v=yjSwwPF5BUU.

[52] Maestre JM, Fernández MI & Jurado I, 2017, *An application of economic model predictive control to inventory management in hospitals*, **71**, pp. 120–128.

[53] Mahalakshmi M & Sundararajan M, 2013, *Traditional SDLC Vs Scrum Methodology – A Comparative Study*, **3(6)**, pp. 192–196.

[54] Makombo T, 2016, *Fast Facts: Public health sector in need of an antidote*, **6(298)**, pp. 1–6.

[55] Mettler T & Eurich M, 2012, *A "design-pattern"-based approach for analyzing e-health business models*, **1(2)**, pp. 77–85.

[56] Mouton J, 2001, *How to Succeed in Your Master's and Doctoral Studies: A South African Guide and Resource Book*, 1st Edition, van Schaik, Pretoria.

[57] Myny D, Van Goubergen D, Limère V, Gobert M, Verhaeghe S & Defloor T, 2009, *Determination of standard times of nursing activities based on a nursing minimum dataset*, **66(1)**, pp. 92–102.

[58] Nau R, 2014, *Forecasting with moving averages*, (Unpublished) Technical Report, Durham, North Carolina.

[59] Noel M, 1978, *Inventory control: a review for hospital pharmacists*, **13**, pp. 617–631.

[60] Otexts.com, 2018, *Trend methods: Holt's linear trend method*, Accessed: 2019/05/13, URL: https://otexts.com/fpp2/holt.html.

[61] Parlar M, Perry D & Stadje W, 2011, *FIFO Versus LIFO Issuing Policies for Stochastic Perishable Inventory Systems*, Methodology and Computing in Applied Probability, **13(2)**, pp. 405–417.

[62] Parliament of the Republic of South Africa, 1996, *Constitution of the Republic of South Africa No. 108 of 1996*, **108(38)**, pp. 1241–1331(56).

[63] Pasandideh SHR, Niaki STA & Yeganeh JA, 2010, *A parameter-tuned genetic algorithm for multi-product economic production quantity model with space constraint, discrete delivery orders and shortages*, Advances in Engineering Software, **41(2)**, pp. 306–314.

[64] Paulussen T, Jennings N, Decker K & Heinzl A, 2003, *Distributed Patient Scheduling in Hospitals*, 18th International Joint Conference on Artificial Intelligence, May, pp. 1224–1229.

[65] Pegram AM, Grainger M, Jones K & While AE, 2015, *An exploration of the working life and role of the ward manager within an acute care hospital organisation*, **20(4)**, pp. 312–328.

[66] Principles of Health Science, 2013, *Hospital Organizational Structure*, Accessed: 2018-01-29, URL: https://www.youtube.com/watch?v=FpQEwbAV3Qw.

[67] Quia.com, 2013, *Organizational Structure of a Hospital*, (Unpublished) Technical Report.

[68] Riley C, 2018, *Jeremy Bentham and Equity: The Court of Chancery, Lord Eldon, and the Dispatch Court Plan*, **39(1)**, January, pp. 29–57.

[69] Rosales CR, Magazine M & Rao U, 2015, *The 2Bin system for controlling medical supplies at point-of-use*, **243(1)**, pp. 271–280.

[70] Royce WW, 1970, *Managing the development of large software systems*, **26(August)**, pp. 1–9.

[71] Saelim N, Porouhan P & Premchaiswadi W, 2016, *Improving organizational process of a hospital through Petri-net based repair models*, (Unpublished) Technical Report.

[72] SAPA, 2012, *Aids drug shortage in 6 provinces*, Accessed: 2018-06-14, URL: https://www.sowetanlive.co.za/news/2012-05-30-aids-drug-shortage-in-6-provinces/.

[73] Schneider H, 1978, *Methods for Determining the Re-order Point of an (s , S) Ordering Policy when a Service Level is Specified*, **29(12)**, pp. 1181–1193.

[74] Schooler JW, 2014, *Turning the Lens of Science on Itself: Verbal Overshadowing, Replication, and Metascience*, Perspectives on Psychological Science, **9(5)**, pp. 579–584.

[75] Schwaber K, 1997, *SCRUM Development Process*, Springer, London.

[76] Siddaway A, 2014, *What is a systematic literature review and how do I do one?*, (Unpublished) Technical Report, University of Stirling, Stirling, UK.

[77] Singh G, 2018, *7 Methods to Perform Time Series Forecasting (with Python codes)*, Accessed: 2018-08-01, URL: https://www.analyticsvidhya.com/blog/2018/02/time-series-forecasting-methods/.

[78] Smartsheet, 2016, *What's the Difference? Agile vs Scrum vs Waterfall vs Kanban*, Accessed: 2018-05-04, URL: https://www.smartsheet.com/agile-vs-scrum-vs-waterfall-vs-kanban.

[79] Smith KI, Everson RM, Fieldsend JE, Murphy C & Misra R, 2008, *Dominance-based Multiobjective Simulated annealing*, IEEE Transactions on Evolutionary Computation, **12(3)**, pp. 323–342.

[80] South African Department of Health, 2018, *Medicine: Medicine List*, Accessed: 2018-07-19, URL: http://www.health.gov.za/index.php/medicine.

[81] Southwest Solutions Group, 2018, *Par vs Kanban Hospital Inventory — Two Bin Supply System for Nursing Supplies*, Accessed: 2018-04-26, URL: http://www.southwestsolutions.com/healthcare/par-vs-kanban-hospital-inventory-two-bin-supply-system-for-nursing-supplies.

[82] Statistics South Africa, 2016, *Mid-year population estimates 2016*, (Unpublished) Technical Report, Statistics South Africa, Pretoria.

[83] Staudhammer C & LeMay V, 2000, *Height prediction equations using diameter and stand density measures*, The Forestry Chronicle, **76(2)**, pp. 303–309.

[84] Takeuchi H & Nonaka I, 1986, *New New Product Development Game*, **64(1)**, pp. 1–10.

[85] USAID — SIAPS Program, 2016, *RxSolution*, Accessed: 2019-11-04, URL: http://siapsprogram.org/tools-and-guidance/rxsolution/.

[86] VanDerLinde LP, 1983, *System to Maximise Inventory Performance in a Small Hospital*, **40(3)**, pp. 70–73.

[87] Verma EP, 2015, *Definition of software development life cycle—Phases of SDLC - Yuvayana*, Accessed: 2018-03-20, URL: http://er.yuvayana.org/definition-of-software-development-life-cyclephases-of-sdlc/.

[88] Walther BA & Moore JL, 2015, *The concepts of bias, precision and accuracy, and their use in testing the performance of species richness estimators, with a literature review of estimator performance*, **28(6)**, pp. 818.

[89] West Cape News, *Health department throws away R3,4 million worth of expired medicine*, Accessed: 2018-06-14, URL: http://westcapenews.com/?p=4747.

[90] Wilson KJ, Hodge R, Bivens D, Wilson KJ, Hodge R & Bivens D, 2015, *Reducing Stockouts in a Cancer Center's Ambulatory Care Clinics*, **27(3)**, pp. 99–108.

[91] Woosley JM, Wiley-Patton S, Kelle P & Schneider H, 2009, *Decision support in healthcare supply chain management and pharmaceutical inventory control*, Proceedings of the 15th Americas Conference on Information Systems, AMCIS Proceedings, California, pp. 1–10.

[92] Xu E, Wermus M & Bauman DB, 2011, *Development of an integrated medical supply information system*, **5(3)**, pp. 385–399.

APPENDIX A

Systematic Literature Review Data

This appendix contains two tables discussed in Chapter 4, the *Systematic Literature Review*. Table A.2 is referenced during the selection process of the *Screening* section (§ 4.4) on page 31. It provides the source titles for all the literature chosen after reading every abstract. Additionally, the table also indicates which of these chosen records could be acquired for reading. Table A.1 indicates where the research for each of the acquired literature was conducted. This literature is dominantly studies performed in developed countries, with exception to two records obtained from Indonesia and Thailand. This was referenced in § 4.4.4.

TABLE A.1: *Systematic Literature Review: Geographic locations.*

Continent	Country	*Classification*	City	Number	Total
Asia	Indonesia	*Developing*	Padang	1	3
	Japan	*Developed*	Otaru	1	
	Thailand	*Developing*	Bangkok	1	
Europe	Belgium	*Developed*	Ghent	2	10
	Italy	*Developed*	Bari	1	
			Fisciano	1	
	Spain	*Developed*	Sevilla	2	
	Sweden	*Developed*	Stockholm	1	
	Switzerland	*Developed*	Bern	1	
	United Kingdom	*Developed*	Coventry	1	
			London	1	
North America	United States	*Developed*	California-Berkeley	1	13
			Chicago	1	
			Louisiana	2	
			Michigan	1	
			Minnesota	1	
			New York	4	
			Philadelphia	1	
			Tennessee	1	
			Texas	1	
Total				**26**	**26**

TABLE A.2: *SLR: List of records' source titles.*

Source type	Chosen	Acquired
Article	30	17
Abstracts of Health Care Management Studies	1	
Abstracts of Hospital Management Studies	1	
American Journal of Hospital Pharmacy	5	4
Canadian Journal of Hospital Pharmacy	1	
Control Engineering Practice	1	1
Deutsche Apotheker Zeitung	1	
Emergencias	1	1
Enterprise Information Systems	1	1
Farmacia Hospitalaria	1	
Gesundheitswesen	1	
Health Care Management Science	1	1
Hospital and Community Psychiatry	1	
Hospital Financial Management	1	1
Hospital Pharmacy	1	1
Indian Journal of Public Health Research and Development	1	
International Journal of Healthcare Management	1	1
International Journal of Services and Operations Management	1	1
Journal of Advanced Nursing	1	1
Journal of Research in Nursing	1	1
Medizinische Klinik - Intensivmedizin und Notfallmedizin	1	
Operations Research for Health Care	1	1
Revue Hospitaliere de France	1	
The International Journal of Health Planning and Management	1	
Transfusion	1	1
Yakugaku Zasshi	1	1
(Undefined)	1	
Book Chapter	1	0
Low-Wage Work in Denmark	1	
Conference Proceeding	7	6
15th Americas Conference on Information Systems 2009, AMCIS 2009	1	1
2010 IEEE Workshop on Health Care Management, WHCM 2010	1	1
67th Annual Conference and Expo of the 2017 Industrial and Systems Engineering Research Conference	1	1
AIP Conference Proceedings	2	1
IJCAI International Joint Conference on Artificial Intelligence	1	1
2016 14th International Conference on ICT and Knowledge Engineering	1	1
Review	3	3
Best Practice and Research: Clinical Anaesthesiology	1	1
Emergency Medicine Journal	1	1
EMJ - Engineering Management Journal	1	1
Total	**41**	**26**

APPENDIX B

Programming script

This appendix contains all the referenced code mentioned during the report. This does not include the code for the simulations conducted in Chapters 6–9, which consist of roughly 450–700 lines of code each. The code shown here is simply to share the packages which were used to perform some of the methods discussed, tested and implemented during this thesis. These packages are free and readily available, making them the ideal choice to allow future work or real-world implementation.

B.1 Forecasting Code

This section shows how the forecasting methods described in Chapter 5 were tested using existing Python packages. In every test, the same three data sets were read in from CSV files as DataFrames. These data sets are the "flat", "increasing" and "decreasing" demand sets described in § 5.2.2 and represented in Figure 5.1. Each set contains 365 days (1 year) of historic demand values ($t = \{-1, \ldots, -365\}$) and 45 days known future values ($t = \{0, \ldots, 44\}$).

Historic DataFrames were created from the demand data sets called "historic_Flat", "historic_Increasing" and "historic_Decreasing" respectively, which represent the 1 year historic demand at that moment in time. Table B.1 briefly describes the construct of these three DataFrames. In the case of forecasts which use a window size (p), the most recent 30 days of the historic demand was selected ($t = \{-1, \ldots, -30\}$). Results from each forecast was individually saved to an excel sheet for inspection and accuracy testing, using the remaining 45 days future values.

Moving Average Method (MA):

```
import numpy as np

# Window size
p = 30

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand[(-1)*p:]))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand[(-1)*p:]))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand[(-1)*p:]))
```

TABLE B.1: *DataFrames containing the historic demand values used for forecast testing.*

	historic_Flat		historic_Increasing		historic_Decreasing	
Index	t	Demand	t	Demand	t	Demand
-365	-365	100	-365	100	-365	100
-364	-364	101	-364	105	-364	79
-363	-363	100	-363	102	-363	94
-362	-362	108	-362	90	-362	99
-361	-361	79	-361	83	-361	119
⋮	⋮	⋮	⋮	⋮	⋮	⋮
-5	-5	104	-5	158	-5	15
-4	-4	99	-4	182	-4	16
-3	-3	119	-3	195	-3	34
-2	-2	99	-2	170	-2	27
-1	-1	101	-1	183	-1	28

Weighted Moving Average Method (WMA):

```
import numpy as np
import math

# Window size
p = 30

# Weight
w = [0] * p
denominator = 0
for k in range(0, p):
   denominator = denominator + (p-k)
for j in range(0, p):
   w[j] = ((p-j)/denominator)

# Calculate forecast values
F_Flat = 0
F_Increasing = 0
F_Decreasing = 0
for z in range(0, p):
   F_Flat = F_Flat + (w[z] * historic_Flat.Demand[(-1)*(z+1)])
   F_Increasing = F_Increasing + (w[z] * historic_Increasing.Demand[(-1)*(z+1)])
   F_Decreasing = F_Decreasing + (w[z] * historic_Decreasing.Demand[(-1)*(z+1)])

# Final forecast values
d_F['Flat'] = math.ceil(F_Flat)
d_F['Increasing'] = math.ceil(F_Increasing)
d_F['Decreasing'] = math.ceil(F_Decreasing)
```

Moving Simple Exponential Smoothing Method (MSES):

```
import numpy as np
from statsmodels.tsa.api import SimpleExpSmoothing

# Window size
p = 30

# Calculate forecast values
fit1 = SimpleExpSmoothing(np.asarray(historic_Flat['Demand'])).fit()
d_F['Flat'] = np.ceil(fit1.forecast(len(future)))
fit2 = SimpleExpSmoothing(np.asarray(historic_Increasing['Demand'])).fit()
d_F['Increasing'] = np.ceil(fit2.forecast(len(future)))
fit3 = SimpleExpSmoothing(np.asarray(historic_Decreasing['Demand'])).fit()
d_F['Decreasing'] = np.ceil(fit3.forecast(len(future)))

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand))
```

Naive Approach Method (NA):

```
import numpy as np

# Final forecast values
d_F['Flat'] = historic_Flat.Demand[-1]
d_F['Increasing'] = historic_Increasing.Demand[-1]
d_F['Decreasing'] = historic_Decreasing.Demand[-1]
```

Holt's Linear Trend Method (HLT):

```
import numpy as np
from statsmodels.tsa.api import Holt

# Calculate forecast values
fit1 = Holt(np.asarray(historic_Flat['Demand'])).fit()
d_F['Flat'] = np.ceil(fit1.forecast(len(future)))
fit2 = Holt(np.asarray(historic_Increasing['Demand'])).fit()
d_F['Increasing'] = np.ceil(fit2.forecast(len(future)))
fit3 = Holt(np.asarray(historic_Decreasing['Demand'])).fit()
d_F['Decreasing'] = np.ceil(fit3.forecast(len(future)))

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand))
```

Holt-Winter Multiplicative Method (HWA):

```
import numpy as np
from statsmodels.tsa.api import ExponentialSmoothing

# Calculate forecast values
fit1 = ExponentialSmoothing(np.asarray(historic_Flat['Demand']),
seasonal_periods=7, trend='mul', seasonal='mul').fit()
d_F['Flat'] = np.ceil(fit1.forecast(len(future)))
fit2 = ExponentialSmoothing(np.asarray(historic_Increasing['Demand']),
seasonal_periods=5, trend='mul', seasonal='mul').fit()
d_F['Increasing'] = np.ceil(fit2.forecast(len(future)))
fit3 = ExponentialSmoothing(np.asarray(historic_Decreasing['Demand']),
seasonal_periods=4, trend='mul', seasonal='mul').fit()
d_F['Decreasing'] = np.ceil(fit3.forecast(len(future)))

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand))
```

Simple Average Method (SA):

```
import numpy as np

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand))
```

Holt-Winter Additive Method (HWA):

```
import numpy as np
from statsmodels.tsa.api import ExponentialSmoothing

# Calculate forecast values
fit1 = ExponentialSmoothing(np.asarray(historic_Flat['Demand']),
seasonal_periods=7, trend='add', seasonal='add').fit()
d_F['Flat'] = np.ceil(fit1.forecast(len(future)))
fit2 = ExponentialSmoothing(np.asarray(historic_Increasing['Demand']),
seasonal_periods=4, trend='add', seasonal='add').fit()
d_F['Increasing'] = np.ceil(fit2.forecast(len(future)))
fit3 = ExponentialSmoothing(np.asarray(historic_Decreasing['Demand']),
seasonal_periods=3, trend='add', seasonal='add').fit()
d_F['Decreasing'] = np.ceil(fit3.forecast(len(future)))

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand))
```

Damped Trend Method (DT):

```
import numpy as np
from statsmodels.tsa.api import Holt

# Calculate forecast values
fit1 = Holt(np.asarray(historic_Flat['Demand']), damped=True).fit()
d_F['Flat'] = np.ceil(fit1.forecast(len(future)))
fit2 = Holt(np.asarray(historic_Increasing['Demand']), damped=True).fit()
d_F['Increasing'] = np.ceil(fit2.forecast(len(future)))
fit3 = Holt(np.asarray(historic_Decreasing['Demand']), damped=True).fit()
d_F['Decreasing'] = np.ceil(fit3.forecast(len(future)))

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand))
```

Simple Exponential Smoothing Method (SES):

```
import numpy as np
from statsmodels.tsa.api import SimpleExpSmoothing

# Calculate forecast values
fit1 = SimpleExpSmoothing(np.asarray(historic_Flat['Demand'])).fit()
d_F['Flat'] = np.ceil(fit1.forecast(len(future)))
fit2 = SimpleExpSmoothing(np.asarray(historic_Increasing['Demand'])).fit()
d_F['Increasing'] = np.ceil(fit2.forecast(len(future)))
fit3 = SimpleExpSmoothing(np.asarray(historic_Decreasing['Demand'])).fit()
d_F['Decreasing'] = np.ceil(fit3.forecast(len(future)))

# Final forecast values
d_F['Flat'] = math.ceil(np.average(historic_Flat.Demand))
d_F['Increasing'] = math.ceil(np.average(historic_Increasing.Demand))
d_F['Decreasing'] = math.ceil(np.average(historic_Decreasing.Demand))
```

B.2 Inverse of the Cumulative Normal Distribution Code

This work was discussed and referenced in § 4.6.6 and § 6.3.1. The variables describe; *service_level* (α) is the confidence percentage, *mu* (μ) is the average and *sigma* (σ) is the standard deviation.

The inverse of the cumulative normal distribution function, $F^{-1}(\alpha, \mu, \sigma)$:

Excel Code:

=NORM.INV(α, μ, σ)

Python Code:

```
import scipy.stats
F_1 = scipy.stats.norm.ppf(service_level, mu, sigma)
```

APPENDIX C

Testing the Inventory Policies

This appendix contains the 100 products used for testing purposes in this thesis. All products began with no stock on hand (SOH) and were allocated the same storage space (1 unit size). The minimum order quantity (MOQ), lead time (L) and price of the products were randomly extracted from the Master Procurement Catalogue, following the distribution of lead times specified in Table 3.3, as described in § 3.5.2. The shelf life was randomly generated between 2 and 36 months in intervals of 30 days, as specified during § 3.4.5. The product categorization was created to test Policy 8 in § 6.1.1 which uses different z-scores based on the product category, according to Table 6.1.

TABLE C.1: *Product specifications for inventory policy testing.*

Product	SOH	MOQ	Lead time	Shelf life	Space size	Category	Price
Product 1	0	108	7	480	1	C	R129.67
Product 2	0	144	14	780	1	C	R65.19
Product 3	0	108	7	270	1	C	R168.40
Product 4	0	50	14	480	1	C	R41.11
Product 5	0	20	14	510	1	C	R17.65
Product 6	0	100	14	870	1	C	R43.96
Product 7	0	10	14	600	1	B	R93.81
Product 8	0	5	14	1080	1	A	R483.00
Product 9	0	6	14	180	1	B	R96.66
Product 10	0	400	14	480	1	B	R3.47
Product 11	0	140	14	720	1	C	R10.35
Product 12	0	20	14	1080	1	C	R19.42
Product 13	0	1	14	990	1	A	R161.58
Product 14	0	1	14	1080	1	B	R508.30
Product 15	0	5	14	1080	1	C	R260.36
Product 16	0	10	14	360	1	C	R6.26
Product 17	0	20	14	180	1	C	R25.83
Product 18	0	6	14	420	1	C	R46.21
Product 19	0	40	14	90	1	C	R25.47
Product 20	0	40	14	60	1	C	R24.73

TABLE C.2: *Product specifications for inventory policy testing.*

Product	*SOH*	*MOQ*	Lead time	Shelf life	Space size	Category	Price
Product 21	0	300	14	960	1	C	R5.96
Product 22	0	80	14	750	1	C	R8.63
Product 23	0	75	14	510	1	C	R28.58
Product 24	0	360	14	540	1	C	R5.62
Product 25	0	500	14	930	1	B	R4.79
Product 26	0	600	14	180	1	B	R4.75
Product 27	0	300	14	1020	1	C	R8.22
Product 28	0	60	14	930	1	A	R10.64
Product 29	0	75	14	90	1	C	R48.02
Product 30	0	360	14	690	1	C	R8.91
Product 31	0	500	14	390	1	C	R6.90
Product 32	0	600	14	360	1	B	R6.65
Product 33	0	50	14	120	1	C	R20.23
Product 34	0	100	14	270	1	C	R20.26
Product 35	0	100	14	510	1	C	R21.28
Product 36	0	144	14	780	1	C	R28.54
Product 37	0	100	14	390	1	C	R23.72
Product 38	0	100	14	90	1	C	R25.97
Product 39	0	80	14	540	1	C	R15.53
Product 40	0	40	14	480	1	C	R24.70
Product 41	0	120	14	690	1	C	R20.70
Product 42	0	100	14	630	1	C	R7.19
Product 43	0	100	14	120	1	A	R8.94
Product 44	0	10	14	480	1	C	R28.18
Product 45	0	100	14	810	1	A	R2.13
Product 46	0	22	14	870	1	C	R6.44
Product 47	0	10	14	300	1	C	R806.00
Product 48	0	1	14	690	1	B	R60.53
Product 49	0	30	14	1020	1	C	R11.67
Product 50	0	500	14	600	1	A	R5.20
Product 51	0	50	14	720	1	C	R35.41
Product 52	0	48	14	300	1	C	R247.60
Product 53	0	4	14	180	1	B	R316.41
Product 54	0	480	14	660	1	C	R3.73
Product 55	0	10	14	360	1	C	R22.36
Product 56	0	600	14	1080	1	C	R2.36
Product 57	0	160	14	90	1	C	R2.49
Product 58	0	10	14	330	1	C	R9.80
Product 59	0	10	14	870	1	C	R15.89
Product 60	0	10	14	750	1	C	R26.35
Product 61	0	20	14	780	1	C	R23.58
Product 62	0	12	14	990	1	C	R94.23
Product 63	0	1	14	960	1	B	R77.38
Product 64	0	1	14	840	1	B	R38.70
Product 65	0	90	14	510	1	C	R9.18

TABLE C.3: *Product specifications for inventory policy testing.*

Product	*SOH*	*MOQ*	Lead time	Shelf life	Space size	Category	Price
Product 66	0	60	14	630	1	C	R47.41
Product 67	0	640	14	630	1	C	R8.71
Product 68	0	640	14	420	1	C	R8.82
Product 69	0	360	14	750	1	C	R11.30
Product 70	0	24	14	1020	1	B	R235.00
Product 71	0	24	14	840	1	C	R212.00
Product 72	0	24	14	630	1	C	R195
Product 73	0	1	28	570	1	A	R103.84
Product 74	0	1	14	60	1	A	R136.28
Product 75	0	10	21	60	1	C	R56.20
Product 76	0	10	14	180	1	B	R190.39
Product 77	0	10	14	1050	1	C	R39.79
Product 78	0	10	14	330	1	C	R25.32
Product 79	0	10	14	960	1	C	R58.67
Product 80	0	10	14	780	1	C	R36.4
Product 81	0	10	14	60	1	B	R92.21
Product 82	0	10	14	300	1	C	R57.21
Product 83	0	48	14	450	1	C	R124.81
Product 84	0	5	14	180	1	B	R88.15
Product 85	0	60	14	840	1	C	R60.22
Product 86	0	30	14	360	1	C	R34.44
Product 87	0	60	14	540	1	C	R34.44
Product 88	0	100	14	60	1	C	R51.22
Product 89	0	30	14	480	1	C	R32.48
Product 90	0	1	14	570	1	A	R79.69
Product 91	0	1	14	270	1	A	R169.05
Product 92	0	20	14	810	1	C	R33.53
Product 93	0	20	14	210	1	C	R49.34
Product 94	0	20	14	1080	1	C	R26.05
Product 95	0	80	14	900	1	C	R12.88
Product 96	0	1600	14	780	1	C	R1.54
Product 97	0	1000	14	810	1	C	R1.67
Product 98	0	192	14	90	1	C	R6.49
Product 99	0	1000	14	960	1	C	R7.19
Product 100	0	1000	14	570	1	C	R18.04

APPENDIX D

Validation Forms

This appendix contains the results of the four completed validation forms as filled out by the subject matter experts (SMEs). SME feedback 1–3 represent the forms completed by SMEs 1–3 described in § 8.2. Similarly, SME feedback 4 represent the form completed by SME 4.

TABLE D.1: *Validation results: SME feedback 1 – 10 October 2019.*

§	Section	Question/Remark	strongly disagree	disagree	neutral	agree	strongly agree
6.3	Inventory Policy: Daily historic demand	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?				X	
		This model could help operations.				X	
8.1.1	Inventory Policy: Weekly historic demand	Was it easy to understand?				X	
		Is it realistic?			X		
		Does the model require improvement?				X	
		This model could help operations.				X	
7.1	Supplier Prioritisation: As a whole	Was it easy to understand?			X		
		Is it realistic?			X		
		Does the model require improvement?				X	
		This model could help operations.			X		
7.1.1	Supplier Prioritisation: Varying lead times	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?				X	
		This model could help operations.				X	
7.1.2	Supplier Prioritisation: Order cost	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
7.1.3	Supplier Prioritisation: Choosing optimal supplier	Was it easy to understand?				X	
		Is it realistic?			X		
		Does the model require improvement?			X		
		This model could help operations.				X	
7.2	Order Irregularities	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?				X	
		This model could help operations.				X	

Table D.2: *Validation results: SME feedback 2 – 10 October 2019.*

§	Section	Question/Remark	strongly disagree	disagree	neutral	agree	strongly agree
6.3	Inventory Policy: Daily historic demand	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
8.1.1	Inventory Policy: Weekly historic demand	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
7.1	Supplier Prioritisation: As a whole	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
7.1.1	Supplier Prioritisation: Varying lead times	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
7.1.2	Supplier Prioritisation: Order cost	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
7.1.3	Supplier Prioritisation: Choosing optimal supplier	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
7.2	Order Irregularities	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	

TABLE D.3: *Validation results: SME feedback 3 – 10 October 2019.*

§	Section	Question/Remark	strongly disagree	disagree	neutral	agree	strongly agree
6.3	Inventory Policy: Daily historic demand	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?		X			
		This model could help operations.					X
8.1.1	Inventory Policy: Weekly historic demand	Was it easy to understand?					X
		Is it realistic?					X
		Does the model require improvement?		X			
		This model could help operations.					X
7.1	Supplier Prioritisation: As a whole	Was it easy to understand?					X
		Is it realistic?			X		
		Does the model require improvement?			X		
		This model could help operations.			X		
7.1.1	Supplier Prioritisation: Varying lead times	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?				X	
		This model could help operations.				X	
7.1.2	Supplier Prioritisation: Order cost	Was it easy to understand?					X
		Is it realistic?					X
		Does the model require improvement?		X			
		This model could help operations.				X	
7.1.3	Supplier Prioritisation: Choosing optimal supplier	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?			X		
		This model could help operations.				X	
7.2	Order Irregularities	Was it easy to understand?				X	
		Is it realistic?				X	
		Does the model require improvement?		X			
		This model could help operations.					X

TABLE D.4: *Validation results: SME feedback 4 – 22 August 2019.*

§	Section	Question/Remark	strongly disagree	disagree	neutral	agree	strongly agree	Comments
7.1	Supplier Prioritisation: As a whole	Was it easy to understand?					X	
		Is it realistic?				X		
		Does the model require improvement?		X				
		This model could help operations.			X			Yes: Procurement
7.1.1	Supplier Prioritisation: Varying lead times	Was it easy to understand?					X	
		Is it realistic?					X	
		Does the model require improvement?		X				
		This model could help operations.			X			Yes: Procurement
7.1.2	Supplier Prioritisation: Order cost	Was it easy to understand?					X	
		Is it realistic?				X		
		Does the model require improvement?		X				
		This model could help operations.			X			Yes: Procurement
7.1.3	Supplier Prioritisation: Choosing optimal supplier	Was it easy to understand?					X	
		Is it realistic?				X		
		Does the model require improvement?		X				
		This model could help operations.			X			Yes: Procurement
7.2	Order Irregularities	Was it easy to understand?					X	
		Is it realistic?				X		
		Does the model require improvement?		X				
		This model could help operations.			X			Yes: Procurement
7.3	Storage Limitations	Was it easy to understand?					X	
		Is it realistic?				X		
		Does the model require improvement?			X			
		This model could help operations.			X			Yes: Procurement