Information Inaccuracy in Inventory Systems

by

Yun Kang

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2004

Author ....
Department of Mechanical Engineering
February 1, 2004

Certified by..
Stanley B. Gershwin
Senior Research Scientist
Thesis Co-Supervisor

Certified by..
Sanjay E. Sarma
Associate Professor of Mechanical Engineering
Thesis Co-Supervisor

Accepted by..
Ain A. Sonin
Chairman, Department Committee on Graduate Students

Information Inaccuracy in Inventory Systems

by

Yun Kang

Submitted to the Department of Mechanical Engineering
on January 19, 2004, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

It is critically important for inventory-carrying facilities to provide high availability of products at the minimal operating cost. To achieve this objective, many companies have automated their inventory operations and rely on the information system in critical decision makings. However, if the information is inaccurate, it may lead to high out-of-stocks and/or excess inventory. This thesis examines what the primary causes of the inaccuracy are, how and to what extent they degrade the inventory system performance, and what can be done to compensate for the inaccuracy.

Analytical and simulation modelling demonstrate that the inventory system performance is highly sensitive to the inaccuracy caused by stock loss, which is the disappearance of items (such as due to theft) not detected by the information system. That is, even a small level of stock loss accumulated over time can lead to inventory inaccuracy that disrupts the replenishment process and creates severe out-of-stocks. In fact, revenue losses due to out-of-stocks can far outweigh the property losses due to the disappearing items.

One way to deal with the inaccuracy problem is the use of RFID-based automatic product identification technology under development at the Auto-ID Center, which can provide the real-time and accurate information regarding the location and quantity of objects in supply chain. It is found that even when this technology provides imperfect measurement of the stock quantity, dramatic performance improvement can be achieved using an inventory control scheme based on dynamic programming.

Various other methods of compensating for the inventory inaccuracy are presented and evaluated. Analysis of each method reveals that the inventory inaccuracy problem can be effectively treated even without automatic identification technology in some situations. However, each method has weaknesses.

Thesis Supervisor: Stanley B. Gershwin
Title: Senior Research Scientist

Thesis Supervisor: Sanjay E. Sarma
Title: Associate Professor of Mechanical Engineering

Acknowledgments

A wise man said that history, big and small, is made during the process. I certainly did not have to wait till these last moments of my graduate years at MIT to appreciate all the individual experiences that have made and shaped me — both academically and personally. They are also the experiences that have placed me on a track in life that I had never imagined possible at the outset. I am indebted to many who have made these experiences invaluable.

This academic research would not have been possible without my thesis advisor Stan Gershwin. His love and capacity for research have always inspired and challenged me to produce a research work that is new, true, and important. Couple them with his commitment to guiding and nurturing the students, and you have a research advisor rare to find. I extend my gratitude to Sanjay Sarma my co-advisor. His leadership in the Auto-ID Center and guidance in the research provided me ample opportunities to inspire my academic work with real-life problems.

Away from the research labs and lecture rooms, I have been truly fortunate to have my wife Eunjung with me. She has never failed to give me the comfort, strength, and courage I needed to make the steps forward. She truly has been by my side throughout the years.

Many thanks to the Auto-ID Center members. Thanks to Robin Koh for having been my friend and colleague. I may have joked about his 'supply chain fairy tales' a number of times, but the wealth of his industry experience and insights on practical problems certainly have added spices to my academic work. Thanks to Toni Pommet who've always made the workplace warmer and more inviting. Thanks also to Tom Scharfeld, Dan Engels, and the rest of the Auto-ID Center members who participated in jumpstarting this transformational technology.

Last but not least, I express sincere thanks to all my spiritual mentors. You helped me through the most challenging times of graduate years at MIT. You allowed me to see and find what I had never imagined possible before, and to lay the foundation upon which my academic work and so many others were able to thrive. You taught me that I first needed to love to realize how much I am loved. You helped make this Ph.D. experience not just an meaningful end, but also a means to so much more to come.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

For many companies that operate inventory-carrying facilities, providing high product availability to customers at minimal operation costs is one of the key factors that determine the success of their businesses. Especially in industries where the competition is fierce and profit margins are thin, companies have automated the inventory management processes to better meet customer demand and reduce operational costs. For example, many retailers use an automatic replenishment system which tracks the number of products in the store and places an order to suppliers in a timely fashion with minimal human intervention.

By doing so, the companies depend on the accuracy of the computerized information system for critical decision making. Information regarding what products are where and in what quantity must be provided accurately to effectively coordinate the movement of the goods. However, if the information provided by the computer system is incorrect, the ability to provide the product to the consumers at the minimal operation cost is compromised. For example, if the computer's record of stock quantity in the facility does not agree with the actual physical stock, orders may not be placed to the supplier in time, or the facility could be carrying unnecessary inventory.

This research investigates the problems related to the information inaccuracy in inventory systems — what the inaccuracy is, what the causes are, and what impact it has on the performance of the inventory system. In addition to quantifying the costs of inaccuracy, this research also addresses various ways the inaccuracy can be mitigated to improve the system performance.

1.1 Research Motivation

The issues discussed here became apparent due to the work of the Auto-ID Center. The Auto-ID Center, founded in 1999 at the Massachusetts Institute of Technology, is sponsored by over 100 global companies, many of whom are leaders in their industries. Its aim is to create an automatic product identification system that can potentially replace bar-code technology. A radio frequency identification (RFID) tag, which is a microchip with an antenna, would be placed on physical objects in trade — a soda bottle, a pair of jeans, a car engine, etc. By placing the RFID readers that sense the presence of tagged objects throughout key locations in the supply chain, the objects can be tracked from the point of manufacture to and beyond the point of consumption. The Auto-ID Center is engaged in designing and deploying a global infrastructure that will make it possible for computers to provide accurate, real-time identification and location of objects.

Figure 1-1 shows the components of the Auto-ID Center technology and how they interact with one another. In the microchip of each tag is stored a number called Electronic Product Code (EPC) [Jos00]. The EPC is similar to the barcode numbering scheme in identifying the item's manufacturer and product category, but by having more digits than the barcode, it is able to uniquely identify every single item in trade. The RFID reader transmits electromagnetic waves that power the RFID tag, and if the tagged object is within the reading range of the reader (which can vary from a few inches to a few meters), the tag sends back the EPC stored on the microchip.

The local controller, once it retrieves the EPCs from the reader, can find the location in the Internet where the information about the products are stored (such as date of manufacture, ingredients, expiration date, and etc.). It does so by transmitting the EPC to a server that provides a look-up service called Object Naming Service (ONS) [Jos00]. ONS serves as a directory of the EPCs — it takes as an input an EPC and returns the address in the Internet where the information about the object resides. Once the location is determined, the local controller then contacts the server which stores the information about the object. This server uses a structured language called Product Markup Language (PML) to describe the product in a way that can be understood by other computers [Bro01]. PML is based on the widely used extensible markup language (XML), and makes it possible to share the information necessary for common business tasks.

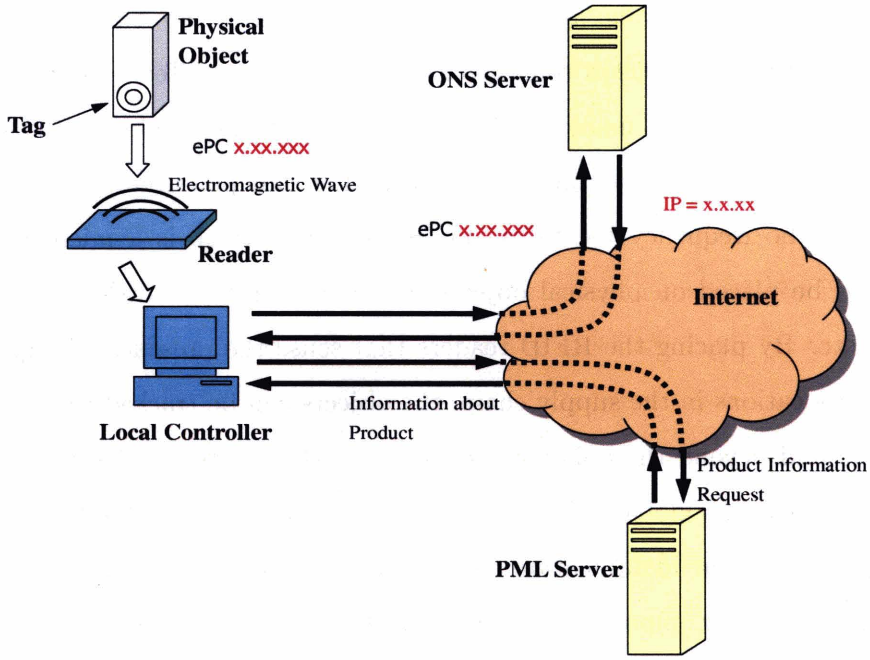

Figure 1-1: Auto-ID Center Technology

Therefore, Auto-ID Center technology attempts to extend the network of computers to interact with the physical objects. Although the Internet enables sharing of information between one computer and another, computers still remain unable to interact with the physical objects. The Auto-ID Center, by using the RFID technology and by developing a universal and open standard for identifying products, allows the computers to gather and share the information about the physical objects in the world.

By creating a network of physical objects, the Auto-ID Center technology can provide a wide range of applications throughout the supply chain. Figure 1-2 lists some of the potential opportunities in a generic, simplified supply chain in which the physical goods typically flow in the direction of the arrows. For example, the manufacturers are able to track their manufactured goods uniquely, and in case of a recall, can trace only the defective goods instead of the entire batches of items shipped. At warehouses and distribution centers, material transactions (such as shipment verification) at both the inbound and outbound sides can be facilitated through the use of automatic identification. Retailers can process the checkouts faster by not having to scan individual products with the barcode scanner. By constantly monitoring the items on the shelves, the store inventory system can also predict

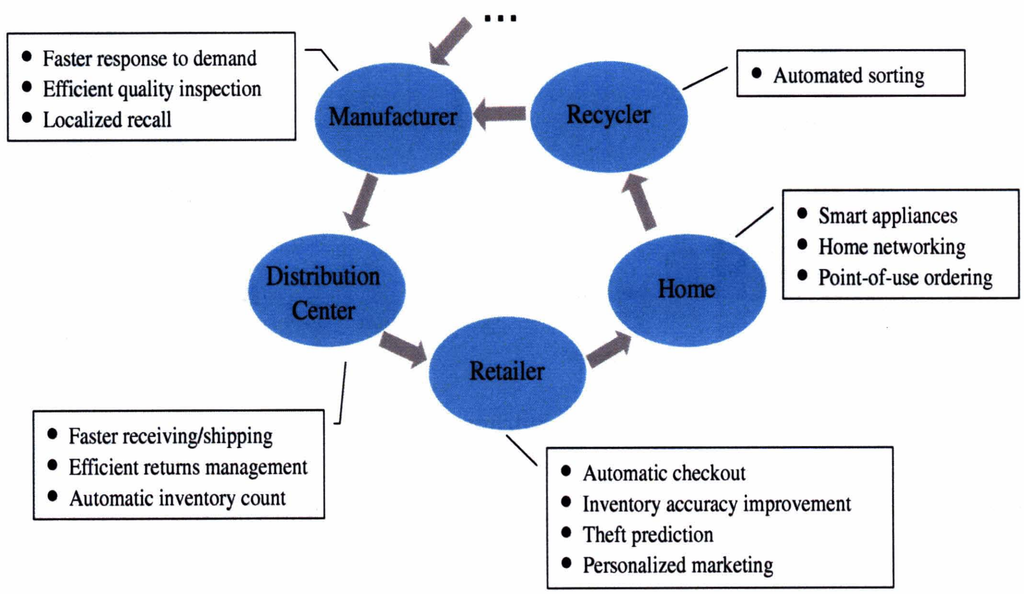

Figure 1-2: Potential applications of the Auto-ID Center technology in supply chain

a bulk pilferage when an unusually high number of items are taken off the shelf at one time. The applications extend beyond the stores. At consumers' homes, the microwave ovens can be made 'smart' by reading the tag ID on the food package with its built-in RFID reader and downloading the cooking instructions from the food manufacturer.

1.2 Inventory Inaccuracy

While working with a number of select sponsors in the area of consumer goods to understand the potential applications of the Auto-ID Center technology, we learned something that is contrary to a popular belief. That is, retailers are not very good at knowing how many products they have in the stores.

1.2.1 Real-life Example

Consider a global retailer who will be referred to as Company A for confidentiality. Each store carries thousands of product lines (also known as SKUs — stock keeping units), and as a common practice for any inventory-carrying facility, it conducts a physical count of all the items in all the stores at least once a year for financial reporting purposes. After the manual inventory verification, the stores are able to compare the stock quantity in the inventory

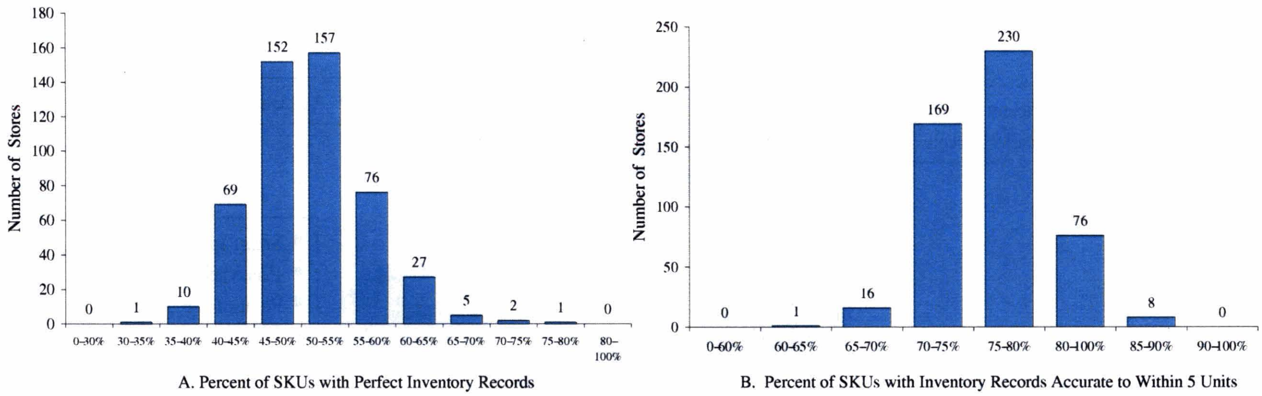

Figure 1-3: Inventory accuracy in Company A stores

record (which is stored in the computer information system) and the actual stock quantity. For each store, the percentage of SKUs whose inventory record matches the actual stock perfectly is calculated. Define this as the *perfect inventory accuracy* of a store. Figure 1-3A summarizes the inventory accuracy of a large subset of Company A's stores.

According to the histogram, the best performing store is the one in which only 75%-80% of its inventory records match the actual inventory. In one store, two thirds of its inventory records are inaccurate. On average, the inventory accuracy of Company A stores is only 51%. In other words, only about a half the SKUs have perfectly accurate inventory records.

Another measure of the inventory accuracy can be obtained by relaxing the requirement and allowing the inventory record of a SKU be considered accurate if it agrees with the actual stock within ± 5 items. A histogram for this definition is shown in Figure 1-3B. Under this definition, the average accuracy of Company A stores rises to 76%. What this means is that on average, the inventory record for one out of four SKUs in the store deviates from the actual stock by six or more items.

The impact of inaccurate inventory records on the performance of retailers like Company A can be severe because the stores rely on the inventory record to make important operations decisions. Since Company A stores carry thousands of SKUs, tracking the inventory record of every SKU manually is very time-consuming. Instead, the stores use an automatic replenishment system in which the inventory record of each SKU is monitored and the computer system determines the order quantity based on the inventory record readings. If there is an error in the inventory record, items may not be ordered in a timely fashion, resulting in out-of-stocks or excess inventory.

Raman et al. reports similar findings from a study done with a leading retailer. Out of close to 370,000 SKUs investigated, more than 65% of the inventory records did not match the physical inventory at the store-SKU level. Moreover, 20% of the inventory records differed from the physical stock by six or more items. The retailer in the report also used information technology extensively to automate the replenishment processes [RDT01].

1.2.2 Causes of the Inventory Inaccuracy

These findings indicate that perfect inventory records are difficult to maintain and inventory records are very likely to be incorrect. The causes of discrepancies in the records are many, and some of the commonly observed ones are discussed here: stock loss, transaction error, inaccessible inventory, and incorrect product identification.

Stock loss, also known as *shrinkage* in industry, includes all forms of loss of the products available for sale. One common example is theft, which can be committed by both shoppers (external theft) and employees (internal theft). It also includes collusion between customers and staff and the unauthorized consumption (such as eating) of the stock by both shoppers and employees. In addition, the vendors can also steal merchandise while in the store performing replenishment duties for their merchandise. Stock loss can also occur when products are rendered unavailable for sale by becoming out of date, damaged, or spoiled.

Stock loss can be categorized into known and unknown stock loss. The former refers to all losses that are identified by the store personnel and reflected in the computer inventory record (such as out-of-date products that are taken off the shelf and written off the books). The latter refers to the rest of the losses not detected and thus not updated into the record. Undetected theft, for example, would fall under this category. It is the unknown stock loss that creates inventory record inaccuracy.

Transaction error occurs typically at the inbound and outbound sides of the facility. At the inbound side, shipments that arrive from the suppliers have to be registered into the store information system and this registration involves a manual, error-prone process. There may be discrepancy between the shipment record and the actual shipment, and if it goes unnoticed by the receiving clerk, the inventory record will not reflect the actual stock accurately. On the outbound side, the checkout registers are not exempt from contributing to the inventory record errors. Typically, the cashiers are rewarded based on the speed of checkouts, and when a shopper brings similar products with identical price, they may choose to scan only one of the products and process them as identical SKUs. The result is that the inventory record of the scanned product decreases more than it should, while that of other products is left unchanged.

Inaccessible inventory refers to products that are somewhere in the facility but are not available because they cannot be found. This can happen when a consumer takes a product from the shelf and places it at another location. It can also happen in the back room or any other storage area in the store. The inaccessible inventory will eventually be found and made ready for sale. However, a long time may pass until this happens, and until then, the inaccessible products are no different from being nonexistent as far as revenue is concerned.

Incorrect product identification can occur in several different ways. Wrong labels can be placed on the products by both the suppliers and the stores. When the barcode on these labels are scanned during receiving or checkout, the inventory record for wrong items will change. Incorrect identification can also happen during manual inventory counts.

What makes inventory inaccuracy seem like an insurmountable problem is the sheer volume of the products handled in the stores. Typical retail stores, being at the far end of the supply chain, are the merge points of thousands of products that come in all different categories, shapes, and sizes, and tens of thousands of items may come into and go out of the store in a single day. For this reason, keeping track of the location of every item and making sure the inventory record agrees with the actual stock quantity is a daunting task.

1.2.3 The Stock Loss Problem

Determining which causes contribute to inventory record error and in what proportion is no less difficult than maintaining the accuracy of the inventory record itself. While the stores admit the gravity of inventory inaccuracy problems and consider it to be one of the major obstacles to the successful execution of their operations, they often do not know when and where it occurs and in what magnitude. However, of all the inventory error causes discussed, industry findings suggest that the unknown stock loss can be a dominant factor for many SKUs.

What makes the unknown stock loss differ from the other causes discussed here is the direction of the inventory record error. Since the loss of the physical items are not reported in the record, the inventory record overestimates the stock. On the other hand, the other causes — transaction errors, inaccessible inventory, and incorrect product identification — can make the error either positive or negative. While it would be almost impossible to break down the inventory error into individual causes, the results of manual inventory counts can reveal some truth about the extent to which unknown stock loss contributes to the inventory inaccuracy. If the inventory record overestimates the actual stock persistently, it is likely that unknown stock loss is the dominant cause of the inaccuracy.

Consider again Company A whose stores carry brands from Company B, who is a global consumer goods manufacturer. To understand the extent of the inventory inaccuracy problem, the two companies decided to pick the topmost selling product from Company B and monitor how the inventory record and the actual inventory change over the period of eight weeks. Dozens of Company A's stores were selected in several regions of North America, and field observers visited the stores once a week and manually counted the stock quantity of the product. At the outset of this testing, the inventory record was set to exactly match the actual inventory. At the end of the testing, however, the actual inventory was less than the inventory record, and the total adjustment was 5% of sales quantity on average over the stores tested. In a thin margin retail industry, this figure is a substantial loss in the bottom line profit.

Company C is a leading supermarket chain who also uses automatic replenishment system for its stores, and in a recent year reported combined known and unknown stock loss of 1.14% of sales in monetary value. Among the product categories that have the highest rates of stock loss were batteries and razor blades, whose stock loss equaled 8% and 5% of sales, respectively. Both of these are products characterized by high value and small size, and thus it was believed that theft accounted for most of the losses.

There are also few industry reports that shed light on the magnitude of the stock loss at the macroscopic level. An extensive study on the magnitude of stock loss was conducted by ECR Europe. Based on a sampling of 200 companies with dominant share of the consumer goods industry in Europe, the study reports that stock loss amounts to 1.75% of sales annually for the retailers. This figure translates to 13.4 billion euros annually. Of this, 59% (or, 1% of total sales) was unknown to the retailers — meaning that the stores did not know where or how the products were lost [Eur01].

Every year, the University of Florida publishes a similar industry-wide empirical research on retail inventory shrinkage in the US [Hol03]. In the most recent report, 118 retailers from 22 different retail markets reported an average stock loss equaling 1.7% of total annual sales, a figure very close to the result from the ECR Europe. It further reports that the retailers perceive theft by the shoppers, employees, and vendors account for 80% of the total stock loss.

Since the stock loss figures are typically obtained by comparing the manual count of all inventories and the store inventory records, these findings suggest that overall in the retail industry, the inventory record error tends to have nonzero mean. The magnitude of this error, however, can vary significantly from one product to another, and the stores are able to estimate this figure for all of its SKUs at the end of yearly audit.

For these reasons, and to simplify our models, we focus on stock loss as a primary cause of inventory record error throughout the thesis.

1.3 Literature Review

The literature in the field of inventory management is vast. However, survey of the literature indicates that almost all of the research that address inventory policies assume perfect knowledge of the inventory is available. There is a scarcity of works that address the causes and consequences of inventory error. Here we summarize published documents that are most closely related to this work.

Iglehart et al. (1972) considers a reorder-point stocking policy subjected to random demand and inventory record error. Assuming the stocking policy is designed to protect only against variations in the demand and lead time, the optimal combination of additional safety stock and frequency of cycle count is obtained. This optimal combination minimizes the sum of the inventory holding cost and the counting cost, subject to the service level meeting a desired target.

Morey (1985) also investigates reorder-point based policies and develops a closed-form expression that relates the service level and three factors that affect it: frequency of cycle count, safety stock level, and the magnitude of inventory record error. This formulation is intended to serve as a very conservative, 'back-of-the-envelope' calculation tool for inventory managers to estimate the service level improvement due to combination of one or more of the three options.

Morey (1986) calculates the minimum required frequency between audits that maintains an inventory record accuracy (not the service level) in pre-specified limits. The optimal audit frequency is determined for two types of audits: perfect audits which eliminate all discrepancies between the book and actual inventory, and imperfect audits which leave errors in records.

A number of works have appeared that address the effective timing of cycle counts in multiple SKU environments. Cantwell (1985), Edelman (1984), and Reddock (1984) discuss the ABC analysis which assigns differing tolerances in inventory accuracy depending on the proportion of the total sales of the products. The size of tolerance would be directly related to the frequency of cycle counts. Neely (1987) proposes a few more methods of determining when to count, including increasing the cycle count frequency for high-activity SKUs.

In environments where there are many SKUs and the cost of manually counting the entire inventory becomes prohibitive, the inventory managers have the option of choosing and counting only a portion of the SKUs. Various sampling techniques exist to perform this task, and are explored in Buck and Sadowski (1983), Dalenius and Hodges (1959), Cochran (1977), Arens and Loebecke (1981), and Martin and Goodrich (1987).

Bernard (1985) and Graff (1987) discuss managerial steps that can be taken to make the cycle counts more effective and to improve the inventory record accuracy in multi-item production environment. Graff (1987) also emphasizes that cycle count merely provides a measurement of the inventory, and it alone is inadequate to control or improve the accuracy.

Various definitions and measures of inventory accuracy are presented in Ernst et al. (1984), Buker (1984), Chopra (1986) and Young (1986). Ernst et al. (1984) also proposes using a control chart to monitor the changes in the inventory accuracy. It serves as a tool for the inventory manager to identify when to look for non-random variability in the inventory accuracy. Hart (1998) provides a case study of a company that used a control chart.

A few works address inventory inaccuracy in MRP (Manufacturing Resource Planning) systems. French (1980) identifies numerous sources of work-in-process inventory inaccuracy. Krajewski (1987) uses a large-scale simulation to assess which factors in a MRP-based production environment (inventory inaccuracy being one of them) have the biggest impact on performance. Brown (2001) also uses simulation to investigate the impact of not only the frequency of error, but the magnitude of error and the location in the bill of material structure where the error takes place.

1.4 Thesis Overview

In this thesis, the research work in inventory inaccuracy is largely divided into two parts. The first part (Chapters 2 and 3) investigates what happens when the inventory record error created by unknown stock loss is left uncorrected and how much the system performance is degraded as a result. Chapter 2 treats inventory systems operating under the (Q,R) policy and Chapter 3 is focused on the fixed review period, base stock policy. At the end of Chapter 3, we present a comparison of the impact that inventory inaccuracy has on the performance of the two policies.

The second part (Chapters 4 through 7) primarily addresses the question of what can be done to deal with the inventory record error and thereby improve the performance of the system. In Chapter 4, various compensation methods are discussed and modelled, including the Auto-ID Center technology which has motivated this research. Chapter 5 and 6 probes deeper into the inventory control problems using the Auto-ID Center technology. The optimal ordering policy is developed for two types of cases: the perfect information case in which the technology provides an accurate measurement of the stock quantity, and the imperfect information case in which the measurement is erroneous. Chapter 7 investigates a near-optimal inventory control scheme for the current inventory systems in which no Auto-ID Center technology is used.

Chapter 2

Inventory Inaccuracy in the (Q,R) Policy

The questions addressed in this chapter and the next are: what impact does inventory record error have on the performance of inventory systems if the error is not corrected, what are the mechanisms by which the inaccuracy degrades the performance, and in what circumstances does it affect the performance most? Specifically, we answer these questions for the (Q,R) policy in this chapter and for the fixed review, base stock policy in Chapter 3.

The rationale behind choosing the (Q,R) policy is its wide presence in the retail industry, including the mass merchandisers who carry a large number of product lines available for consumers at their retail stores. Company A mentioned in the previous chapter, for example, uses policies based on the (Q,R) for many of its fast-moving products.

2.1 Mechanisms of the (Q,R) Policy

In the (Q,R) policy, the inventory is monitored continuously[1] and an order of fixed quantity Q is placed to the supplier if the sum of inventory on-hand (quantity in the facility available for sale) and on-order (quantity ordered to the supplier and waiting to be received and made ready for sale) is less than or equal to the reorder point R. The time between placing an order and its arrival is called *lead time*. Figure 2.1 illustrates the mechanism of this policy.

The reorder point is set so that when an order is placed, enough inventory exists in the

[1] In practice, monitoring is often done daily.

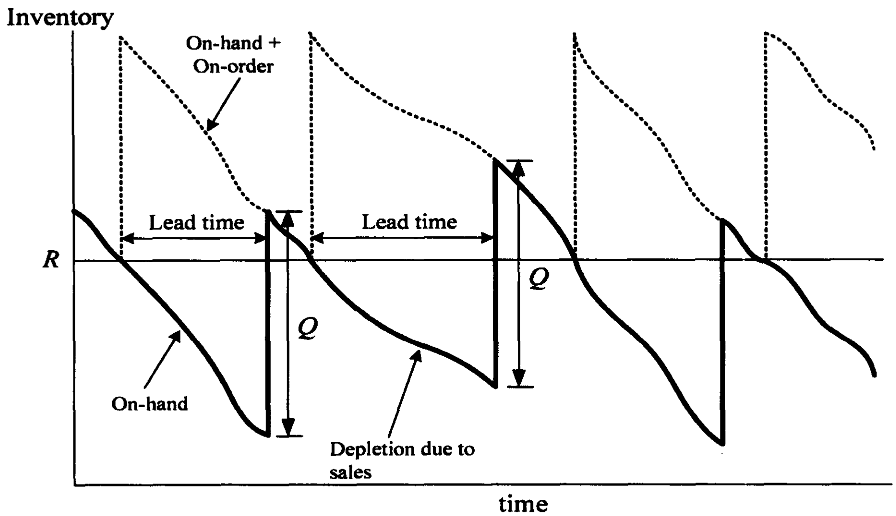

Figure 2-1: (Q,R) Policy

facility to meet the demand until the order arrives. Thus the reorder point has a critical bearing on the performance of this policy. If it is set too low, inventory will be depleted frequently and out-of-stocks will occur. If it is set too high, then the facility will be carrying unnecessary inventory.

The reorder point is often explained as consisting of two parts: the expected value of total demand during lead time and safety stock:

$$R = (\text{expected demand during lead time}) + (\text{safety stock}). \tag{2.1}$$

If the demand is known and constant, then setting the reorder point equal to the total expected demand during lead time would ensure that all demand would be satisfied. However, if there is randomness in the system — such as in the demand or supplier lead time — then the reorder point will have to be higher to cover the uncertainties. This extra inventory is safety stock.

The (Q,R) policy is effective in its timing of orders and thus provides high availability of the stock at the minimal inventory level, provided that the on-hand inventory information used during the review is accurate. In reality, however, the exact value of on-hand inventory is often unknown, and many stores estimate the on-hand inventory based primarily on two measurements that they have access to: the incoming shipments and outgoing sales.

The data for the former is obtained either through order transaction records or shipment verifications, and the latter through a technology commonly used that keeps track of bar-code-scanned sales at the checkout registers (called POS — *Point of Sales*). By updating the computerized inventory record whenever these two events are observed (the industry terminology for this inventory record is *perpetual inventory*), the stores are able to automate the inventory review and order placement processes with minimal human intervention.

According to this method, the inventory record at the beginning of period $k+1$, denoted $\tilde{x}_{k+1}$, is determined from the inventory record at the beginning of the previous period, $\tilde{x}_k$, the quantity received in the previous period, h_k, and the quantity sold in the previous period, a_k, through the relationship

$$\tilde{x}_{k+1} = \tilde{x}_k + h_k - a_k. \tag{2.2}$$

In reality, the inventory record suffers from accuracy problems even if the incoming shipment and sales are known exactly. As discussed earlier, the unknown stock loss is an example of the causes of the error.

Throughout the research, we make a fundamental assumption that differs from those of traditional inventory models: stores do not know the exact value of on-hand inventory at the time of ordering. Therefore, our models distinguish between the inventory record and the actual inventory, and recognize the discrepancy between these two caused by stock loss.

2.2 Stochastic Simulation Model

To see how the stock loss, by creating a discrepancy between the actual inventory and the inventory record, can affect the performance of the (Q,R) policy, consider a single-item inventory model with the following assumptions:

- The demand for purchase during each period k, w_k, is assumed to be independent and normally distributed with mean μ_w and standard deviation σ_w. (That is, a normal distribution with these parameters is used and the negative demands are discarded.)
- The demand for stock loss in period k, v_k, is also independent and identically distributed, and is generated from a Poisson distribution with mean λ.
- The lead time is known and fixed at L.

- The demand occurring when there is zero actual on-hand inventory is lost (no backlog).

A Poisson distribution is chosen for stock loss to prevent assigning negative values when the mean of the distribution is small.

The sequence of events in each period is assumed to be as follows:

1. The inventory record is reviewed and an order is placed to the supplier.
2. The incoming order is received.
3. Sales and stock loss take place.

Denote by x_k the actual inventory at the beginning of period k. According to this sequence, there is $x_k + h_k$ available for meeting the demand for purchase and stock loss in period k. When the sum of demand for purchase and stock loss exceeds the available inventory, the available inventory is divided proportionately to meet the the two demands. The sales in period k is then

$$a_k = \begin{cases} w_k & \text{if } w_k + v_k \le x_k + h_k, \\ (x_k + h_k)\frac{w_k}{w_k+v_k} & \text{otherwise.} \end{cases} \tag{2.3}$$

Since sales can only take on integer values, the quantity in the second line is rounded to the nearest integer. The changes in actual inventory and the inventory record are then

$$\tilde{x}_{k+1} = \tilde{x}_k + h_k - a_k \tag{2.4}$$

$$x_{k+1} = x_k + h_k - a_k - \min(v_k, x_k + h_k - a_k). \tag{2.5}$$

The *min* term represents the *actual* stock loss in period k: it is the smaller of the stock loss demand v_k (sufficient on-hand inventory) and the difference between the available inventory $x_k + h_k$ and sales a_k (insufficient on-hand). Since stock loss is not seen by the inventory record, it is not included in Equation (2.4).

Figure 2-2 shows the evolution of the actual inventory and the inventory record in a sample simulation run. The average demand μ_w is 10 and the standard deviation σ_w is 2. The average daily stock loss λ is 0.2, which is 2% of the average demand. Lead time L is 3 periods and the fixed order quantity Q is 50. The operation end time t_f is chosen to be 365 periods since the standard procedure in the industry is to conduct a physical count of the stock at least once a year and reconcile the inventory record. The initial inventory is

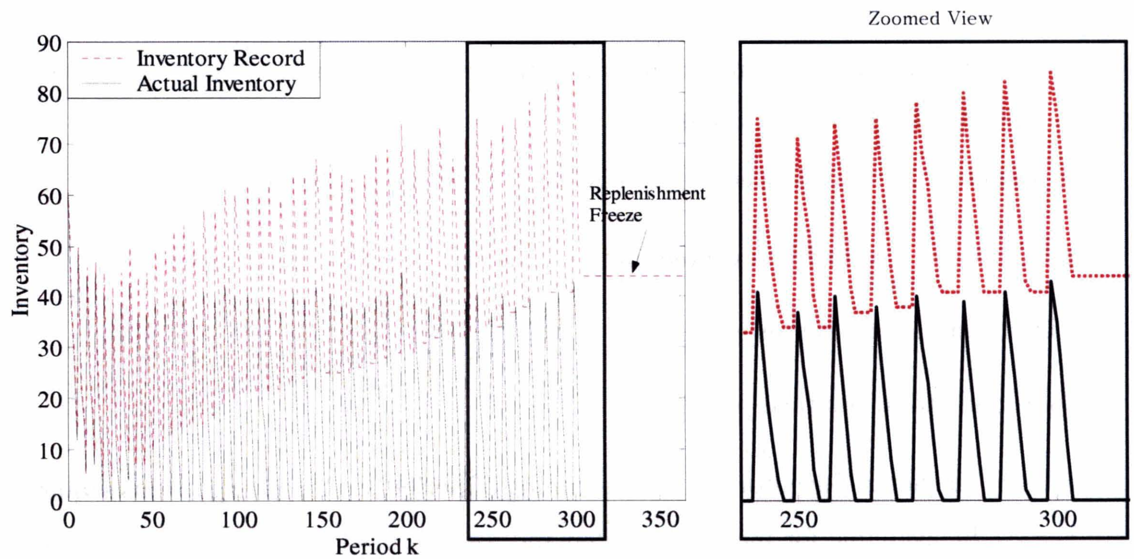

Figure 2-2: A sample simulation run showing inventory evolution of (Q,R) system when subjected to stock loss

$R+Q-\mu_w L$ (chosen to be consistent with the deterministic model to follow this section). It was found through many simulation repetitions that in the absence of stock loss, a reorder point R of 41 produced stockout rate — defined as the total lost sales as a percentage of total demand over the operating period — of 0.5%. We assume that this is the desired target stockout rate and use the corresponding reorder point $R = 41$ when simulating cases in which the stock loss occurs.

Initially, the inventory record and actual inventory are equal. However, the inventory record is not aware of the stock loss and starts to diverge from the actual inventory. As the gap between the two curves widens, the actual inventory curve hits zero more frequently, creating lost sales. Initially, there would be partial out-of-stocks (i.e., a portion of the daily demand during a period would be lost). However, as the inventory error grows further, out-of-stock worsens and the periods in which the entire demand is lost start to appear. This is seen by the flat portions of the actual inventory curve lying on the x-axis. On average, the duration of this complete out-of-stock gets longer with time.

We also observe a continual rise in the inventory record cycles. The two curves jump by the same quantity at the beginning of each cycle since the replenishment quantity is fixed at Q. However, until the next replenishment, the inventory record curve drops slower than the actual inventory curve since it does not recognize the stock loss, thereby making the inventory record cycles rise over time. In fact, when the inventory record cycles continue

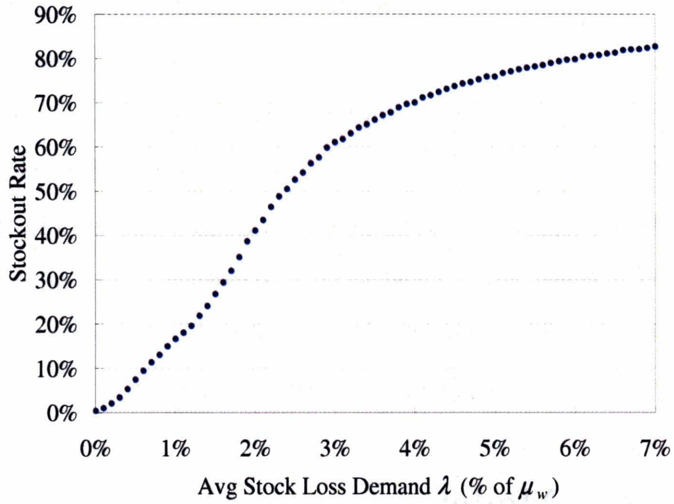

Figure 2-3: Simulation data points for stockout rate vs. average stock loss demand

to rise even further, the system eventually reaches a point where the inventory record stays above the reorder point and no order is placed. Such 'freezing' of replenishment is undesirable since it leads to extremely high lost sales.

Figure 2-3 shows how stockout rate is impacted by stock loss as the average stock loss demand λ is varied from 0% to 7% of average demand. Each data point is the average of 500 independent simulation runs.

When there is no stock loss ($\lambda = 0$), the inventory system achieves the target stockout rate of 0.5%. As the stock loss increases, the error in the inventory record grows and cumulative lost sales rises. Shortly after the stock loss of 1%, the inventory system experiences freezing of replenishment and the stockout rate curve becomes steeper. What is of interest is how fast the stockout rate rises with the stock loss in the system. Even when the stock loss is as small as 1% of average demand (i.e., 1 item disappears for every 100 items in demand by shoppers on average), the error accumulating in the inventory record is large enough to disrupt the replenishment process and make 17% of the total demand from the shoppers lost due to out-of-stocks. When the average stock loss is at 2.4%, more than half of the demand ends up as lost sales. It does not require a high level of stock loss to degrade the in-stock performance of the system. Therefore, when nothing is done to correct the inventory error, stockout rate is highly sensitive to the inventory inaccuracy created by stock loss.

The results also convey a compelling managerial insight. Items lost to shoplifters, for example, are direct loss to the retailer, but the chain reaction created by shoplifting — error in the inventory record, untimely replenishment, and out-of-stocks — creates lost sales substantially greater than the items stolen. Results show that the lost sales quantity can be ten to twenty times higher. Even if the comparison is made in bottom line monetary values[2], the unrealized revenue due to lost sales is substantial in highly competitive retail environments. This means that to effectively control the stock loss problem, management needs to pay close attention to maintaining inventory accuracy.

2.3 Deterministic Model

Whereas in the previous section the demand for purchase and stock loss were assumed to be stochastic and discrete, in this section they are treated as constant and continuous. With this simplification, we look for a closed-form solution for the system performance given the parameters of the (Q,R) policy. Moreover, by developing a model with deterministic demand and stock loss, the role that randomness plays in the inventory inaccuracy problem can also be examined.

Assume demand for purchase and stock loss occur at a rate of w and v units per time, respectively. The lead time L is again fixed and known, and the assumption regarding excess demand (lost sales and no backlog) remains unchanged from the previous model. Also, the ordering decision is made in accordance with the (Q,R) policy. Figure 2-4 shows how the on-hand inventory record and actual inventory evolve over time, along with the inventory position, defined as the sum of on-hand and on-order quantity.

The deterministic model exhibits all the essential features seen in the stochastic simulation model — the growing gap between the recorded and actual inventories, the continual rise of the inventory record cycles, and the eventual freezing of replenishment as the inventory record stays above the reorder point.

For convenience, we break the inventory evolution into two time intervals. Let Region A consist of the group of cycles in which no out-of-stock occurs. The time t_A marks the end of this region. Region B consists of the group of complete cycles that fall between t_A and

[2]Since stock loss is lost property and the lost sales is lost revenue, comparison of the impact on profit would require the profit margin of the product. The margin, however, varies widely based on the pricing strategies of the retailers, and could range from a small percentage to multiple times the product cost.

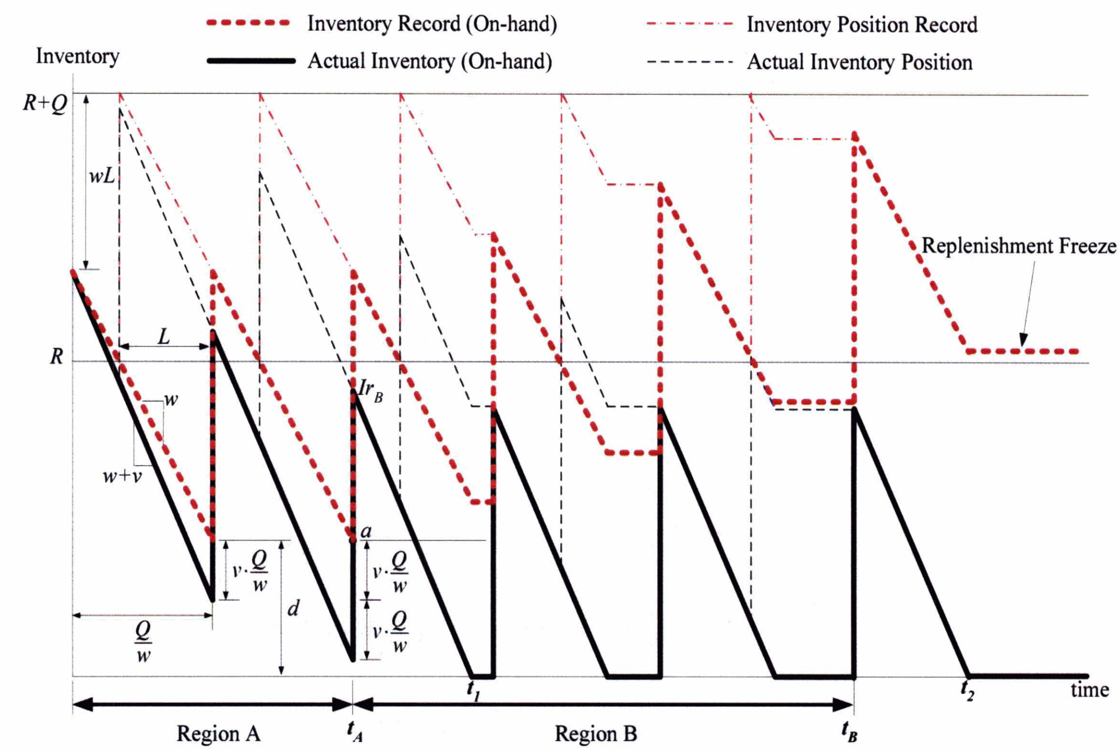

Figure 2-4: (Q,R) Policy subjected to stock loss under deterministic demand

the time of last order arrival prior to system freezing. Let t_B denote the end of this region.

We can compute various performance measures of the system — the time of first out-of-stock t_1, the time of replenishment freeze t_2, and the stockout rate S_{out}. The exact calculations for these quantities can be obtained. However, by making the plausible assumptions that the stock loss rate v is small and there are many replenishment cycles before the end of operation t_f, we arrive at expressions that are much simpler and yet able to approximate the exact calculation very closely. Appendix A describes in detail how the expressions for both exact and approximate values are found. Approximations for t_1 and t_2 are

$$t_1 \approx \frac{R - wL}{v} + \frac{Q}{w + v} \tag{2.6}$$

$$t_2 \approx t_1 + \frac{L(w + v)}{v} + \frac{L}{2}\Big(\frac{wL(W + v)}{vQ} - 1\Big) + \frac{Q}{w + v}. \tag{2.7}$$

The stockout rate S_{out} is determined by adding all the horizontal, flat portions of on-hand actual inventory curve in Figure 2-4 and dividing it by the operating time t_f. It is approximated by

$$S_{out} \approx \begin{cases} 0 & \text{if } t_f < t_1, \\ \frac{1}{t_f}\left[\frac{m(m+1)}{2}\frac{vQ}{w(w+v)}\right] & \text{if } t_1 \leq t_f < t_2, \\ \frac{1}{t_f}\left[\frac{L}{2}\Big(\frac{wL(w+v)}{vQ} + 1\Big) + t_f - t_2\right] & \text{if } t_f \geq t_2, \end{cases} \tag{2.8}$$

where m appearing in the second expression is the number of complete cycles between t_1 and t_f if t_f is located in Region B, approximated by

$$m \approx \frac{2w + v}{2v} + \sqrt{\Big(\frac{2w + v}{2v}\Big)^2 + 2\frac{w(w + v)}{vQ}(t_f - t_1)}. \tag{2.9}$$

Figure 2-5 shows the exact and approximate calculation for stockout rate when the purchase demand rate w is 10 units per time and the stock loss demand rate v is varied from 0 to 0.7 units per time (0% to 7% of purchase demand rate). The lead time L is 3, order quantity Q is 50, and the operation is allowed to continue until $t_f = 365$. The reorder point R is set to equal to the lead time demand of $wL = 30$. In absence of stock loss, this would be the lowest value of R that produces zero stockout rate.

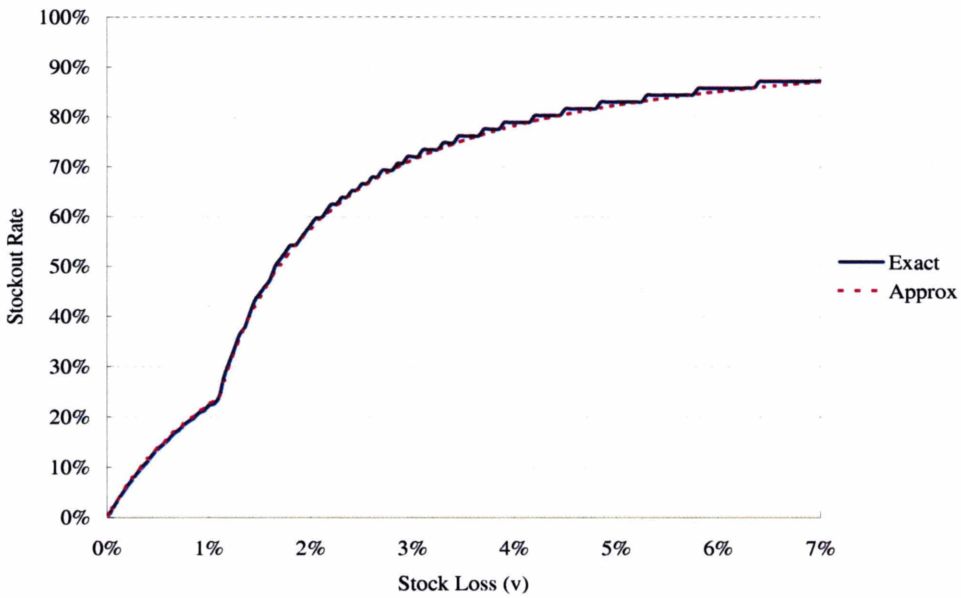

Figure 2-5: Stockout rate for deterministic (Q,R) model subjected to varying stock loss

Stockout rate calculations show that the approximation (dotted line) agrees very closely with the exact (solid line) for the stock loss range shown. When there is no stock loss, stockout rate is zero. With increasing stock loss, however, we observe a rapid rise in stockout rate as seen in the simulation model. Therefore, the deterministic model confirms the finding that when the inventory error is left untreated, system performance is highly sensitive to the inventory inaccuracy created by stock loss. Furthermore, the randomness in demand is not what causes the inventory inaccuracy problem.

One distinct difference between the two curves is the smoothness. While the stockout rate in approximation curve rises smoothly throughout the stock loss shown, the exact calculation curve has small step increases which are followed by almost flat lines at higher values of stock loss. Examining the calculation closely reveals that this abrupt increase occurs whenever the number of cycles in Region B decreases with the increase of stock loss. When the stock loss is high enough to create replenishment freezing, increasing it even further makes the ending inventory record in the last cycle of Region B move closer to the reorder point. When it eventually steps above the reorder point, Region B suddenly loses a cycle and the time of replenishment freeze t_2 decreases. Consequently, the length of flat line beyond t_2 increases abruptly, leading to a sudden rise in stockout rate.

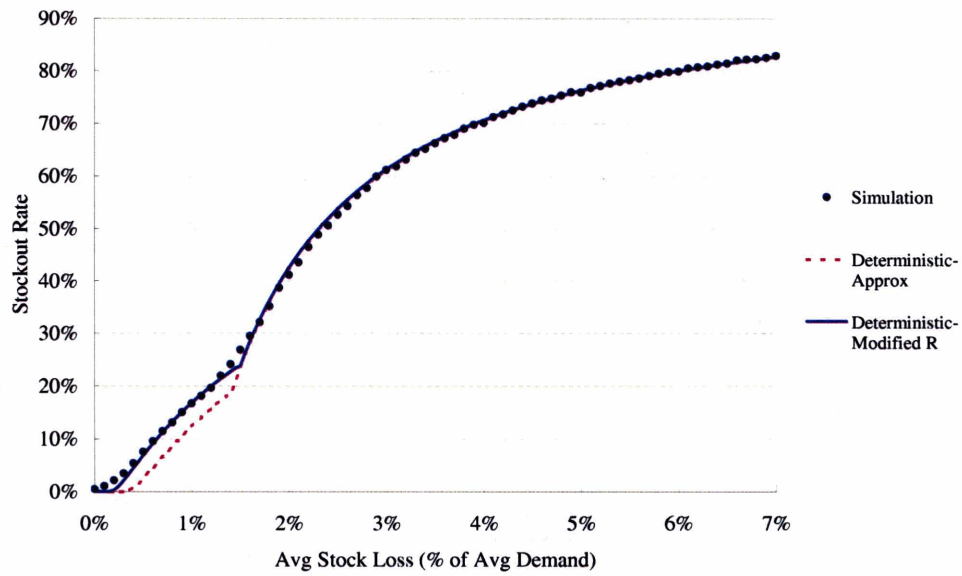

Figure 2-6: Comparison of stockout rate calculations from deterministic and simulation model

Now, we test how well the deterministic demand model predicts the stockout rate of a system subjected to stochastic purchase demand and stock loss demand. In Figure 2-6 is shown the simulation points (from Figure 2-3) along with the approximation for stockout rate calculated in the deterministic model (dotted line). The same values of parameters used in simulation ($R = 41, Q = 50, w = \mu_w = 10, L = 3$) were used for the deterministic model calculation.

The deterministic model makes accurate predictions for stockout rate in stochastic system for high stock losses but it significantly underestimates for stock losses of approximately 1.5% or less. Here we make note of a fundamental difference between the deterministic and stochastic model: whereas in the former order is placed when the sum of inventory record and on-order is *exactly* equal to the reorder point R, in the latter order is often placed at values below the reorder point. This is because in the stochastic model, changes in inventory from one period to the next occurs in multiple quantities. Therefore, the 'true' reorder point in the stochastic model is lower than the designated R, and stockout rate is likely to be higher in the stochastic model. For the cases in which out-of-stocks due to replenishment freeze is not the dominant contributor to the stockout rate (i.e., for small stock losses), this difference makes a noticeable impact on the calculation of stockout rate.

Therefore we look for a correction factor for the reorder point used in deterministic model calculation that will better reflect the actual reorder point in the stochastic model. In absence of stock loss, we find that the average value of inventory record at which ordering takes places is 36.2, which is lower than the set reorder point of 41 by approximately half of the average demand μ_w. Replacing the reorder point R by $R - \frac{w}{2}$ in t_1 expression (Equation (2.6)) and updating the calculation for other variables, we obtain the solid line in Figure 2-6. The result of modifying the reorder point is a marked improvement in the prediction of stockout rate for the entire range of stock losses.

2.4 Sensitivity Analysis

In this section, we investigate in what circumstances the stock loss impacts the system performance the most. Using the simulation model, we conduct a parametric analysis by observing how stockout rate is affected when the lead time L and order quantity Q are varied.

Figure 2-7 illustrates the effect of varying lead time on the stockout rate when the system is subjected to average stock loss demand at 1% of the average purchase demand. The parameters are set to be consistent with what is presented in Section 2.2: $\mu_w = 10, \sigma_w = 2, Q = 50$, and $t_f = 365$. Note that along with the lead time, the reorder point R is also set to provide the same target stockout rate of 0.5% *in the absence of stock loss.* Thus, we assume that the inventory manager, either unaware of the stock loss or ignoring it, simply sets the reorder point based on the purchase demand characteristics and lead time. We have already mentioned that $R = 41$ provided the target stockout rate when the lead time L is 3. For smaller lead times, the variability of the lead time demand is also smaller, and thus the safety stock component of the reorder point can be reduced and still achieve the same target stockout rate — again, provided there is no stock loss. Similarly, for longer lead times, the reorder point will have to increase. However, once we allow uncompensated stock loss, the reduced safety stock associated with shorter lead times leads to much worse performance.

We have seen that the performance of a system with $L = 3$ is highly sensitive to unaccounted stock loss. At shorter lead times, this sensitivity becomes greater. In a system where ordered products are delivered instantly ($L = 0$), it only takes an average stock loss

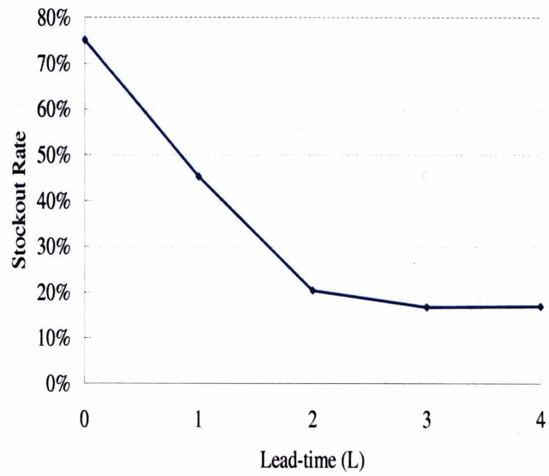

Figure 2-7: Stockout rate vs. L

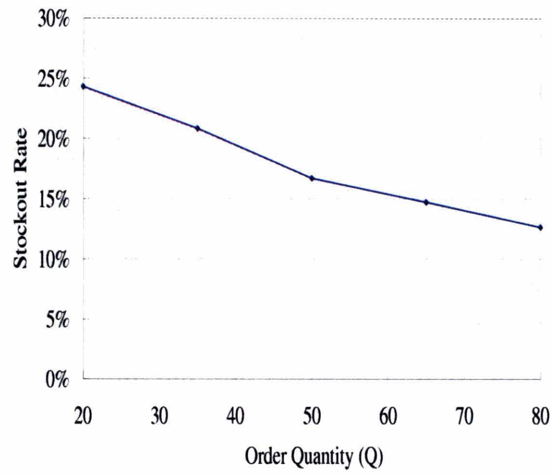

Figure 2-8: Stockout rate vs. Q

of 1% of average demand to render three quarters of the total purchase demand unfulfilled. The reason why such an extreme out-of-stock condition is created is because with zero lead time, Region B in Figure 2-4 does not exist. Instead, in the first cycle after Region A, the system directly enters the replenishment freeze zone. The time of replenishment freeze is on average 95 when $L = 0$, 225 when $L = 1$, and 349 when $L = 2$.

Figure 2-8 is the result of same simulation runs, this time holding $L = 3$ and varying the order quantity Q from 20 to 80. Having a large order quantity reduces stockout rate since the actual inventory is higher on average.

These observations demonstrate the severe consequences of inventory inaccuracy on lean systems characterized by short lead times and frequent ordering of small quantities. At shorter lead times, the desired product availability can be achieved with smaller safety stock if there is no stock loss (thus allowing R to be reduced). However, small safety stock provides little protection against unexpected disturbances in the system. Inventory inaccuracy, which is considered an uncertainty in the system, is likely to wreak far greater havoc on lean systems, and thus maintaining accurate inventory record is critical to reap the benefits lean systems have to offer.

Chapter 3

Inventory Inaccuracy in the Fixed Review Period, Base Stock Policy

Another commonly used inventory policy besides the (Q,R) is the fixed review period, base stock policy. For convenience, this policy will be referred as the base stock policy throughout this chapter. While the (Q,R) policy is efficient in timing of the order and is used for products that are fast-moving, the base stock policy is typically used for slow-moving products that account for a smaller fraction of the revenue.

In this chapter, we investigate how the performance of base stock policy is impacted by inventory inaccuracy caused by stock loss. Also, the cost of this inaccuracy is quantified by the same methods employed in the previous chapter for (Q,R) policy — stochastic simulation and deterministic modelling.

3.1 Mechanisms of the Policy

The mechanisms of the base stock policy differs from the reorder point policy in the quantity and time of order. Whereas in the (Q,R) policy the inventory is reviewed continuously and the order size is fixed, in this policy the inventory is reviewed at fixed times and the order size is varied. Figure 3-1 shows how this policy works. At each review time $r_1, r_2, \ldots, r_k$ which are spaced in interval T, the size of the order is set to bring the on-hand and on-order inventory up to the base stock B.

Under this ordering scheme, the size of order placed at the i^{th} inventory review can be thought of as the sales accumulated since the last review time, denoted $w(r_{i-1}, r_i)$. In this

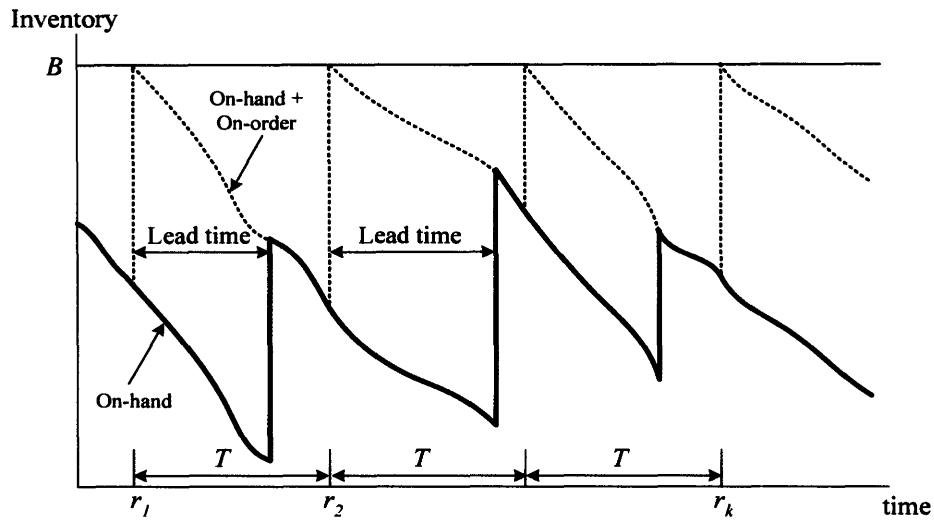

Figure 3-1: Fixed review period, base stock policy

policy, the base stock level B is the decision parameter, and choosing the right base stock is critical in ensuring product availability while keeping the inventory low. Similar to the reorder point in the (Q,R) policy, B is also understood as consisting of two parts as follows:

$$B = \text{(Expected demand during review period and lead time)} + \text{(Safety stock)} \tag{3.1}$$

The safety stock is needed to provide protection from out-of-stocks due to uncertain demand and deliveries.

Because the review is conducted at predetermined times, the base stock policy offers the advantage of shared costs over multiple items, such as cost associated with reviewing orders, order fulfillment, and transportation. Nevertheless, it is not exempt from susceptibility to inventory inaccuracy. In practice, the exact value of the on-hand quantity is not readily available during the inventory review. Instead, much like in the (Q,R) policy, it is implemented in practice based on two accessible pieces of data: the incoming shipment and outgoing sales. The computer system that tracks this estimated inventory (again, called *perpetual inventory*) becomes the guideline in ordering decisions. Unfortunately, when the inventory is subjected to stock loss not reflected in the inventory record, proper replenishment is not made and product availability suffers.

3.2 Stochastic Simulation Model

The assumptions made in simulating an inventory system using the base stock policy are similar to those presented in Section 2.2 for the (Q,R) model. Demand for purchase during each period k, w_k, is assumed to be independent and normally distributed with mean μ_w and standard deviation σ_w. The stock loss demand during each period, v_k is also independent, and has Poisson distribution with mean λ. Furthermore, w_k and v_k are independent. Again, the lead time is fixed and known at L, and the sequence of events in each period does not change: first, review the inventory and place the order to the supplier, secondly, receive the incoming order, and lastly, fill the demand and stock loss. The received quantity in period k and sales during period k are denoted h_k and a_k, respectively. The assumption regarding dividing the available inventory proportionately to the purchase demand and stock loss demand in case of shortage remains unchanged. Therefore, the dynamics of the inventory record $\tilde{x}_k$ and the actual inventory x_k also remains unchanged from Equation (2.4) and (2.5):

$$\begin{aligned}
\tilde{x}_{k+1} &= \tilde{x}_k + h_k - a_k \\
x_{k+1} &= x_k + h_k - a_k - \min(v_k, x_k + h_k - a_k).
\end{aligned}$$

The interval between inventory reviews is T and the first review takes place at time $k = T - L$ (assuming $T \geq L$). The initial value of both the actual inventory and inventory record is $B - \mu_w L$. These selections for the first review time and initial inventory are made to be consistent with the deterministic model to follow in the next section.

In Figure 3-2 is shown the day-ending inventory of a sample simulation run. The value for the parameters are consistent from the (Q,R) simulations in Section 2.2: $\mu_w = 10, \sigma_w = 2, \lambda = 0.2$ (which is 2% of the average demand), $L = 3$, and $t_f = 365$. Again, the target stockout rate is 0.5%, and it was found, through many runs of the simulation, that the base stock level B of 87 achieved this target stockout rate in absence of stock loss. Thus, by keeping B at 87 and introducing stock loss in the system, we are investigating the impact of stock loss if no action is taken to correct the error in inventory record.

We observe a pattern similar to what is found in the (Q,R) policy: the stock loss not recognized by the inventory record causes the actual inventory to be underestimated, and the two inventory curves diverge. As a result, the actual inventory starts to experience

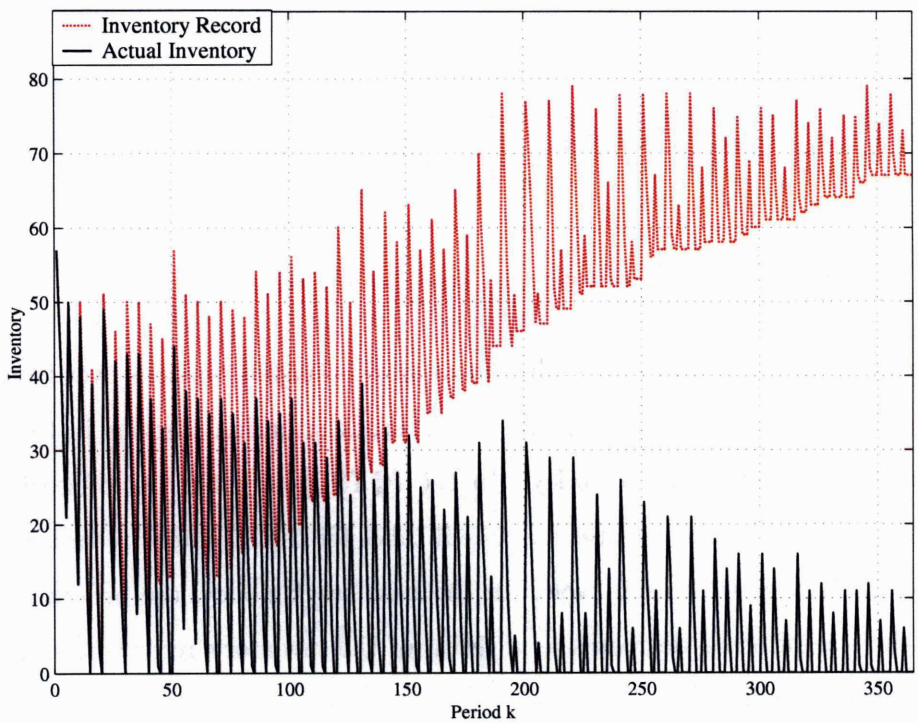

Figure 3-2: A sample simulation run showing inventory evolution of a base stock policy system when subjected to stock loss.

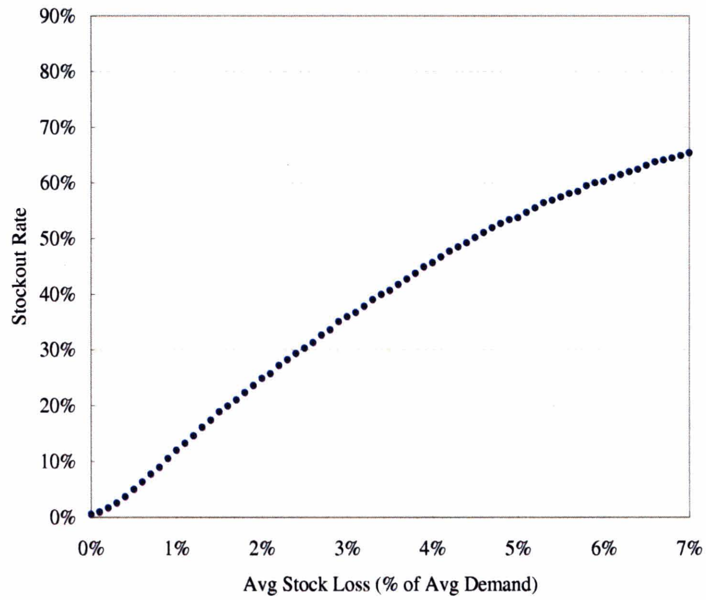

Figure 3-3: Stockout rate vs. stock loss in base stock policy.

more frequent depletion of stock, and the out-of-stock situation worsens over time.

However, the mechanism by which the out-of-stock situation worsens in this policy is different from that of the (Q,R) policy. In the (Q,R) policy, the out-of-stock duration becomes progressively longer because the time it takes for the inventory record to reach the reorder point increases with each cycle. In the base stock policy, the ordering times are fixed, but the size of order decreases as a result of the inventory record moving closer to the base stock B over time.

In Figure 3-3 is shown the relationship between the stockout rate in the base stock policy and the stock loss as the average stock loss is varied from zero to 0.7 (7% of the average demand). The base stock B is maintained at 87, which generates the target stockout rate of 0.5% when there is no stock loss.

The rise in stockout rate is smooth throughout in Figure 3-3, and we do not observe an abrupt change in the slope of the stockout rate curve as seen in the (Q,R) policy (Figure 2-3). The complete shutdown of replenishment like those seen in the (Q,R) policy can also happen in the base stock policy. However, because the order is placed and products are delivered at regular intervals, it takes much longer for the replenishment freeze to occur under comparable settings — the same lead time, demand and stock loss distribution, and average ordering frequency in absence of stock loss. For the replenishment freeze to occur in the base stock policy, the inventory record will have to equal to the base stock level B and the on-order has to be zero simultaneously. For this to happen, either the stock loss rate has to be very high or the operation must continue for an extended period of time.

3.3 Deterministic Model

With the exception of the ordering mechanism, the deterministic demand and stock loss model developed in this section for the base stock policy uses the same set of assumptions made in the (Q,R) deterministic model: demand and stock loss occurs continuously at a constant rate of w and v, respectively, the lead time is fixed at L, and the operation ends at time t_f. Given the base stock level B and the inventory review interval T, we can determine the extent to which the inventory inaccuracy impacts the stockout rate and average inventory.

Figure 3-4 shows how the two on-hand inventories — recorded and actual — change when there is stock loss occurring in the system. The initial inventory is $B - wL$ and the first ordering decision is made at time $T - L$. By setting up the initial conditions this way, the duration of the sawtooth cycles are all equal, and thus the calculation is facilitated.

As the error in the inventory record (dashed line) grows with unrecognized stock loss, the actual inventory (solid line) starts to experience depletion of stock. Whenever there is out-of-stock, no sales take place, and the inventory record stops dropping. This means at the next inventory review, the amount ordered will be less than what it would have been if there were no out-of-stock. With decreasing order size, the out-of-stock duration in the next cycle becomes even longer, which in turn reduces the order size in the next cycle, and so on. The order size approaches zero with infinite operation time.

The stockout rate for this policy has been calculated. Let the inventory evolution be divided into three regions as shown in Figure 3-4. Let Region A consist of the sawtooth cycles in which no out-of-stock occurs. Region B consists of the cycles beyond Region A which have out-of-stock duration less than the lead time L. The rest of the cycles, whose out-of-stock duration is greater than L, fall in Region C. Knowing that the approximation method in the (Q,R) deterministic model yields very good results, we present the approximated calculations for the stockout rate and average inventory in this model also. See

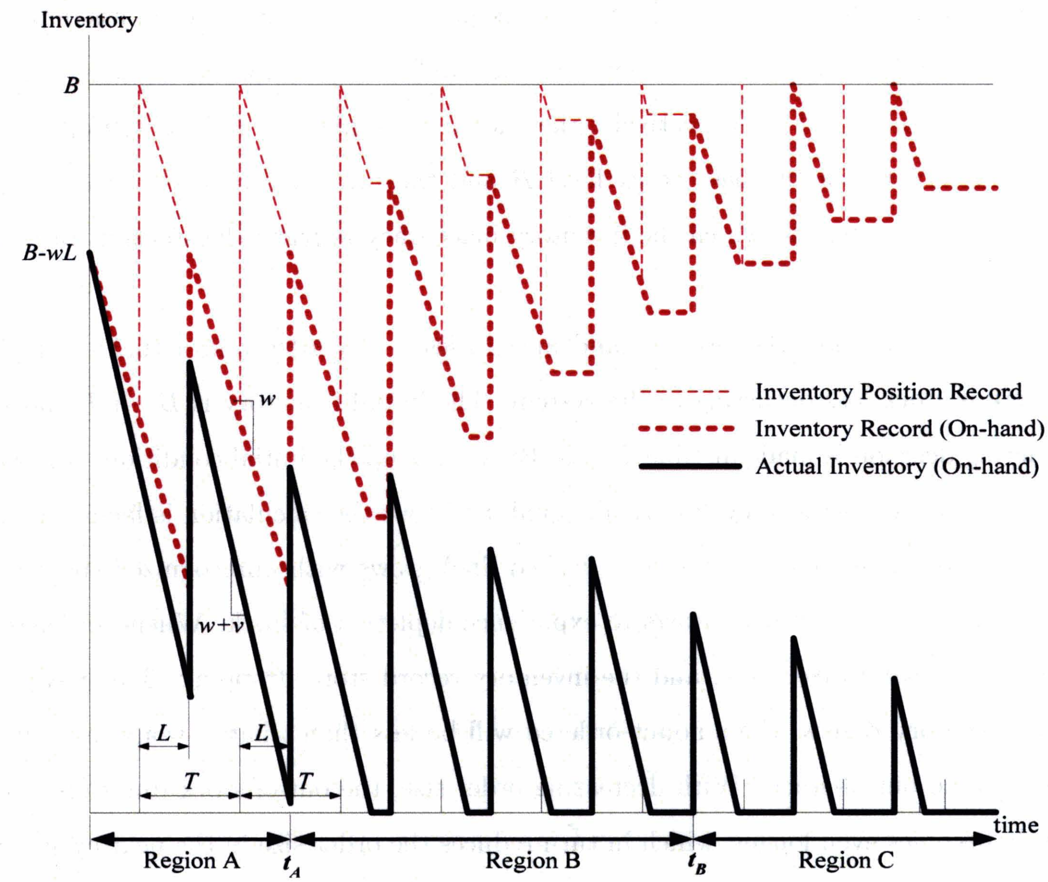

Figure 3-4: Base stock Policy subjected to deterministic demand and stock loss

Appendix B for the details of mathematical steps.

The times that mark the end of Region A and B — t_A and t_B, respectively — are approximated by

$$t_A \approx \frac{B-w(T+L)}{v} \tag{3.2}$$

$$t_B \approx t_A + 2kT, \tag{3.3}$$

where k is related to the number of cycles in Region B and is

$$k = I\left(\frac{\log\left(1-\frac{L}{T}\right)}{\log\left(\frac{w}{w+v}\right)}\right). \tag{3.4}$$

If the operation end time t_f lies in Region C, we need to determine the number of cycles in Region C, denoted n_C, as well. It is approximately

$$n_C \approx I\left(\frac{t_f}{T} - \frac{B-w(T+L)}{vT} - 2k\right). \tag{3.5}$$

The final expression for the stockout rate S_{out} is

$$S_{out} \approx \begin{cases} 0 & \text{if } t_f \le t_A \\ \frac{1}{t_f}\left[(t_f - t_A) - 2T\frac{w}{v} + 2T\left(\frac{w+v}{v}\right)\left(\frac{w}{w+v}\right)^{\frac{t_f-t_A}{2T}+1}\right] & \text{if } t_A < t_f \le t_B \\ \frac{1}{t_f}\left[2kT - 2T\frac{w}{v} + 2T\left(\frac{w+v}{v}\right)\left(\frac{w}{w+v}\right)^{k+1}\right. & \\ \left. +n_C T - (T-L)\left(\frac{w}{v} - \frac{w+v}{v}\left(\frac{w}{w+v}\right)^{n_C+1}\right)\right] & \text{if } t_f > t_B. \end{cases} \tag{3.6}$$

Using this expression for stockout rate, we can examine how well the model built under the assumption of deterministic demand and stock loss can predict the stockout rate of a system subjected to stochastic demand and stock loss. Figure 3-5 compares the results of the stochastic simulation model discussed in the previous section and the deterministic model. The parameter values are set to be consistent with the simulation: $w = \mu_w = 10, T = 5, L = 3$, and $t_f = 365$.

For small stock loss, the stochastic model has a slightly higher stockout rate than the deterministic model. This is because in the former, the base stock level B was purposely set to generate a finite target stockout rate of 0.5%, and thus the out-of-stocks are purely due

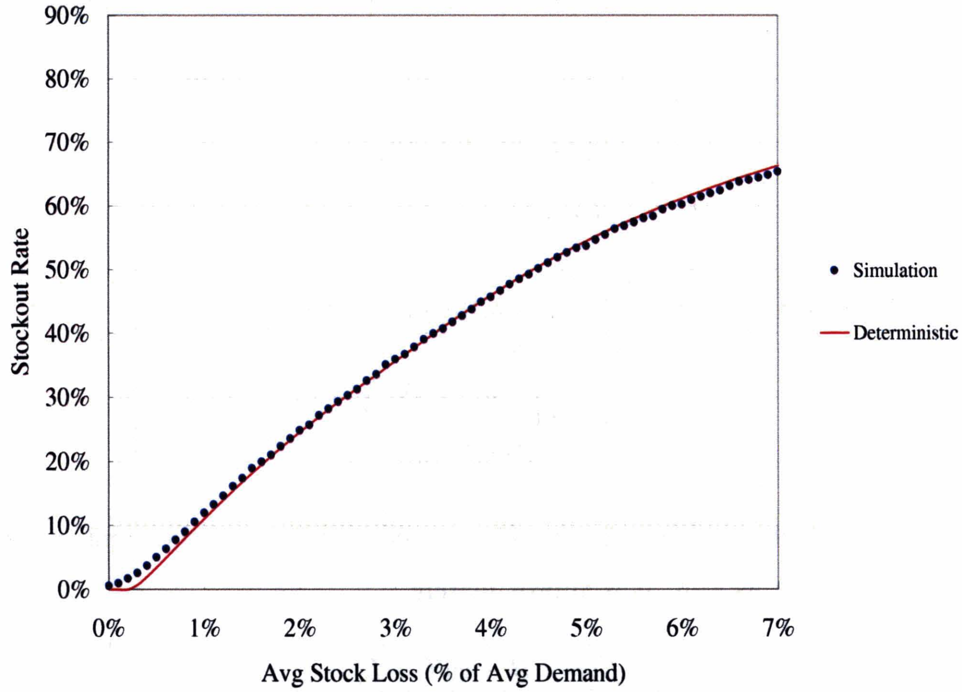

Figure 3-5: Comparison of stockout rate calculations from deterministic and simulation model (base stock policy)

to random demand. In the deterministic model, however, the base stock level of 87 is higher than the minimum base stock that will provide zero stockout rate, which is $w(T+L) = 10(5+3) = 80$. For this reason, the deterministic model produces zero stockout rate until the stock loss rate is 0.2% of the demand rate. With increasing stock loss, the out-of-stocks due to inventory inaccuracy becomes dominant, and the deterministic model makes more accurate predictions.

What both the simulation and deterministic models confirm, however, is that the impact made on the inventory system performance by inaccurate information can be very high in the base stock policy as well as in the (Q,R) policy. A stock loss as small as 1% of demand creates stockout rate well above 10%. At 2% stock loss, more than a quarter of the purchase demand is lost.

3.4 Comparison of the Impact of Inaccuracy on the Two Policies

Having seen how the performance of both the (Q,R) and base stock policy can be highly susceptible to inventory inaccuracy, we now investigate which of the two inventory policies is more sensitive to the inaccuracy created by stock loss. A basis for comparison is made by choosing the parameters in the policies in such a way that *in absence of stock loss*, both policies produce identical inventory curves in the deterministic model. This is achieved by setting the parameters as follows:

$$\begin{aligned} B &= R + Q \\ T &= \frac{Q}{w}. \end{aligned} \tag{3.7}$$

Since the demand rate w is a given parameter, the two equations with two unknowns above will allow us to determine the unique combination of B and T in the base stock policy for any given R and Q in the (Q,R) policy, and vice versa. Suppose the demand rate is 10 units per time and the parameters in the (Q,R) policy are set at $R = 30, Q = 50$, and $L = 3$. Equation (3.7) yields $B = 80$ and $T = 5$. Moreover, the initial inventory needs to be the same. If the initial inventory in the (Q,R) policy is $R + Q - wL = 50$, then it would be the same in the base stock policy at $B - wL = 50$. Figure 3-6 shows how the inventory changes over time in each of the two policies if there is no stock loss occurring in the system.

Since the two curves are identical over time, the two policies produce identical performance — stockout rate, average inventory, and order frequency — in absence of stock loss. Notice the R in the (Q,R) and B in base stock are set to be the minimum possible values which achieve zero stockout rate. Now, by introducing a stock loss in the system, the impact of inventory inaccuracy on the performance of these two policies can be compared.

Stockout rate in the (Q,R) and base stock policy model for varying stock loss is plotted in Figure 3-7. The (Q,R) policy has considerably higher stockout rate for stock losses greater than 1%. This is due to the replenishment freeze that occurs in the (Q,R) model shortly after the stock loss of 1%. The base stock policy, on the other hand, does not experience the replenishment freeze, producing much lower stockout rate.

The base stock policy outperforms the (Q,R) even when the replenishment freeze doesn't occur in the (Q,R) policy (stock losses less than 1%). Note the difference in mechanism by

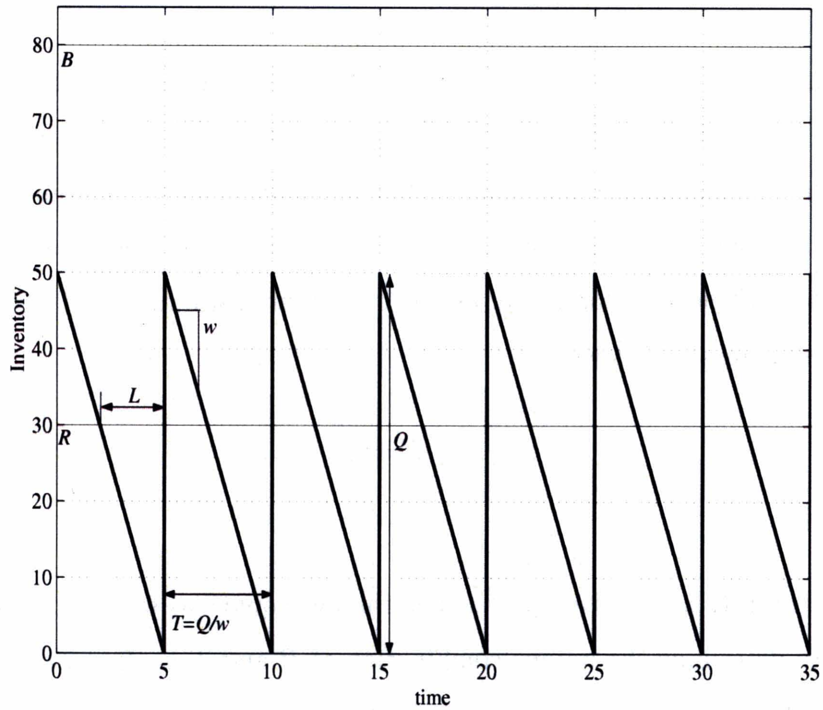

Figure 3-6: Inventory evolution of the (Q,R) and base stock policy in absence of stock loss and randomness

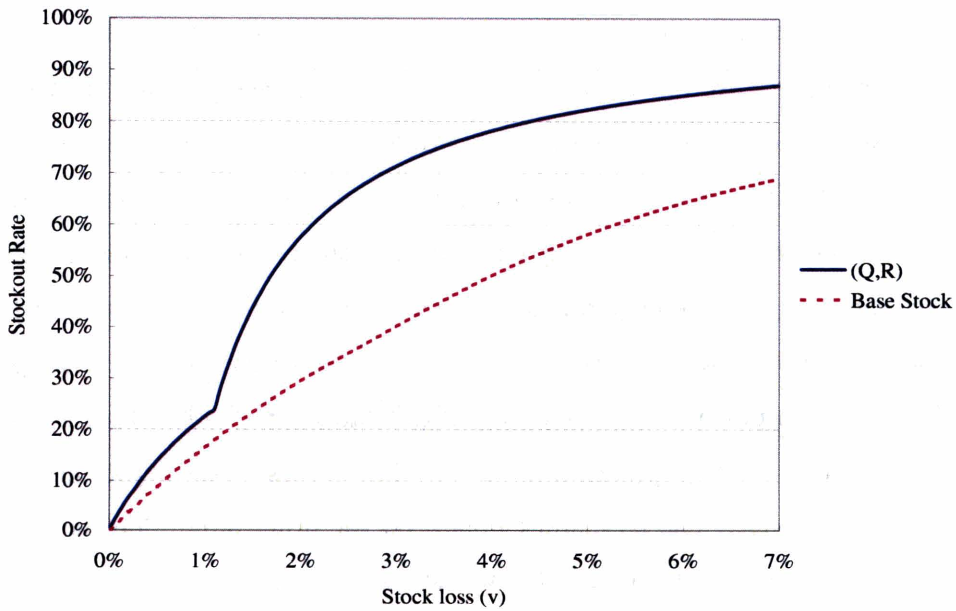

Figure 3-7: Stockout rate comparison of the (Q,R) and base stock policy

which out-of-stock situation worsens with time for a given stock loss. In the (Q,R) policy, the order interval increases with each cycle in Region B while the order quantity is fixed. In the base stock policy, the order quantity decreases with each cycle while the order interval stays constant. The results indicate that the out-of-stocks created due to delay in ordering are greater than the out-of-stocks due to decreasing order quantity. The base stock policy achieves lower stockout rate by ensuring delivery of products at regular intervals during the operation.

Chapter 4

Compensation Methods for Inventory Inaccuracy

In the previous two chapters, it was found that stock loss unknown to the inventory managers can seriously affect the availability of on-hand stock. One underlying assumption used in the models, however, was that nothing is done to correct the inventory record error. The management may not be aware of the stock loss, or may simply choose to ignore it in designing the inventory policy.

In this chapter, we examine various techniques inventory managers can use to compensate for the inventory record error. The possible methods of controlling the error are many, but we describe some of the representative ones here and assess the improvements each method can make in bringing the in-stock performance to the desired level.

4.1 Compensation Methods and Model

Consider the simulation exercise used in Section 2.2 to assess the impact of unknown stock loss on the performance of the (Q,R) policy if no corrective actions are taken. We use this model as a basis for testing how well each error-adjustment method performs in compensating for the inventory error. By using the same set of assumptions, we can examine how much improvement is made from the no-correction case by each compensation technique. Assumptions of the model are restated for convenience:

- The demand for purchase during each period k, w_k, is assumed to be independent and normally distributed (truncated) with mean μ_w and standard deviation σ_w.

- The demand for stock loss in period k, v_k, is also independent and identically distributed, and is generated from a Poisson distribution with mean λ.
- The lead time is known and fixed at L.
- The demand occurring when there is zero actual on-hand inventory is lost (no backlog).

The sequence of events in each period is assumed to be as follows:

1. The inventory record is reviewed and an order is placed to the supplier.
2. The incoming order is received.
3. Sales and stock loss take place.

The dynamics of the inventory record, $\widetilde{x}_k$, and actual inventory, x_k, is influenced by the receipt quantity h_k, sales a_k, and stock loss v_k according to the relationship

$$\widetilde{x}_{k+1} = \widetilde{x}_k + h_k - a_k \tag{4.1}$$

$$x_{k+1} = x_k + h_k - a_k - \min(v_k, x_k + h_k - a_k). \tag{4.2}$$

4.1.1 Safety Stock

Safety stock is often used as a protection against uncertainties in variables in inventory operations, such as the demand and supplier lead time. It can be extended to serve as a buffer against uncertainty in the inventory record.

In (Q,R) policy, the level of safety stock is determined by setting the reorder point R (Section 2.1). Since the reorder point consists of the expected demand during lead time and safety stock, to provide a buffer against inventory error would require increasing the reorder point to a level higher than that needed to cover the variability in purchase demand. In the numerical example shown in Section 2.2, the reorder point of 41 achieved 0.5% stockout rate when there is no stock loss occurring in the system. Since the expected purchase demand during lead time is 30 ($\mu_w L = 10 \cdot 3 = 30$), a safety stock of $41 - 30 = 11$ units was required to provide this target stockout rate. To cover the additional uncertainty in the inventory record, a higher safety stock would be required. Thus, to see the benefit of carrying higher safety stock, we simulate the (Q,R) policy with R higher than 41.

4.1.2 Manual Inventory Verification

One of the most commonly used techniques for mitigating the inventory error is manually counting the items in the facility and correcting the inventory record. The inventory managers can choose to verify the inventory for a part of the entire SKU more frequently than the required yearly audit. This frequency may depend on various elements, such as the availability of the labor and product characteristics, including the profit margin, sales velocity, and whether the products are highly prone to stock loss and other causes of inventory error.

We assume manual verification is done at predetermined, regular intervals, such as every month or every six months. In the simulation, the inventory record is set to equal to the actual on-hand at the end of the period when verification is done. It is assumed that the manual count is done with perfect accuracy.

4.1.3 Manual Reset of the Inventory Record

If a direct measurement of the on-hand inventory is not available, inventory managers can gather and monitor the available data and search for any patterns that may be indicative of the presence of serious inventory error. In the (Q,R) policy, for example, we saw that if the inventory error grows enough, it will eventually reach a point where the inventory record stays above the reorder point and no replenishment is made. In that situation, the daily POS (Point-of-Sales) reading will simply be zero every day. Knowing that this is an unlikely event under normal operations, the inventory manager can choose to manually reset the inventory record to zero, thereby allowing the automated replenishment system to start placing orders again.

To simulate this compensation method, we set the inventory record to zero at the end of each period whenever zero sales is observed. Since the probability of zero demand is very small ($7.4 \cdot 10^{-7}$) in the truncated normal distribution for purchase demand with $\mu_w = 10$ and $\sigma_w = 2$, zero sales would be a strong indication of the existence of an out-of-stock condition.

4.1.4 Constant Decrement of the Inventory Record

If the inventory manager is aware of the presence of stock loss and also knows its stochastic behavior, another way to compensate for the error is to decrement the inventory record by the average stock loss demand each period. Since the actual value of the stock loss at each period is unknown, simply decrementing the record will still not eliminate the error in the inventory record. However, over time, this corrective action can be expected to perform better than leaving the inventory record unadjusted.

In the simulation, an additional step at the end of each period is added to decrease the inventory record by the estimated daily stock loss demand λ. The actual inventory and record now change according to

$$\widetilde{x}_{k+1} = \widetilde{x}_k + h_k - a_k - \lambda \tag{4.3}$$

$$x_{k+1} = x_k + h_k - a_k - \min(v_k, x_k + h_k - a_k). \tag{4.4}$$

4.1.5 Auto-ID

The technology under development at the Auto-ID Center differs fundamentally from the current inventory systems in that it provides a direct measurement of the stock quantity using RFID readers and tags. To preserve generality, we refer as 'Auto-ID' all means of automatically obtaining the direct measurement of the stock quantity without having to count the items manually. Here we assume the Auto-ID provides a perfectly accurate measurement of the actual inventory and examine how it improves the inventory system performance. In the next chapter, we consider a more realistic case in which the Auto-ID makes inaccurate measurement of the actual inventory and develop an optimal ordering policy for various measurement performances.

Auto-ID is simulated by setting the inventory record to equal to the actual inventory at the end of each period. Thus, the ordering decisions are made with the perfect knowledge of the on-hand quantity. The equations describing the inventory estimate and the actual inventory are

$$\widetilde{x}_{k+1} = x_{k+1} \tag{4.5}$$

$$x_{k+1} = x_k + h_k - a_k - \min(v_k, x_k + h_k - a_k). \tag{4.6}$$

It should be noted that new compensation methods can be created by combining two or more of the techniques described above. For instance, manual verification of the inventory can be conducted along with carrying a higher safety stock.

4.2 Results and Discussion

Figure 4-1 shows how the inventory system performance changes when the compensation methods are implemented. The same parameters from the numerical example in Section 2.2 are used: $\mu_w = 10, \sigma_w = 2, L = 3$, and $Q = 50$. The average stock loss λ is held constant at 0.1, which is 1% of average demand for purchase. The reorder point R is varied in steps of 2 around the base value of 41 (which produces the target stockout rate of 0.5% in the absence of stock loss) for each compensation. The rationale behind varying the reorder point is that in the (Q,R) policy, R is the decision parameter, and it is the responsibility of the inventory manager to select the R that produces the most desirable performance (the best combination of average inventory and stockout rate) for each compensation technique.

The figure plots stockout rate against average inventory for each compensation method. Since the desired goal is to obtain a low stockout rate at minimal inventory, the stockout rate-inventory pair is chosen as the performance measure. The vertical distance between the curves is the difference in average inventory required to attain a particular stockout rate. Therefore, for a given stockout rate, the compensation technique with the lowest inventory would be the best-performing one. Notice that by increasing the reorder point higher than the base value 41, we are also testing how each compensation technique performs in conjunction with carrying higher safety stock.

The 'No Compensation' curve represents the case in which nothing is done to correct the inventory error caused by stock loss other than varying R (Section 4.1.1). This curve thus serves as a basis from which improvements made by each compensation method can be observed. The rightmost data point in this curve corresponds to the lowest reorder point, and thus has the highest stockout rate. As R increases, stockout rate improves at the expense of inventory.

In the manual inventory verification method (Section 4.1.2 — represented by 'Verify Twice'), counting is assumed to be conducted twice a year. The result shows that even the infrequent inventory record reconciliation of every six months improves the performance

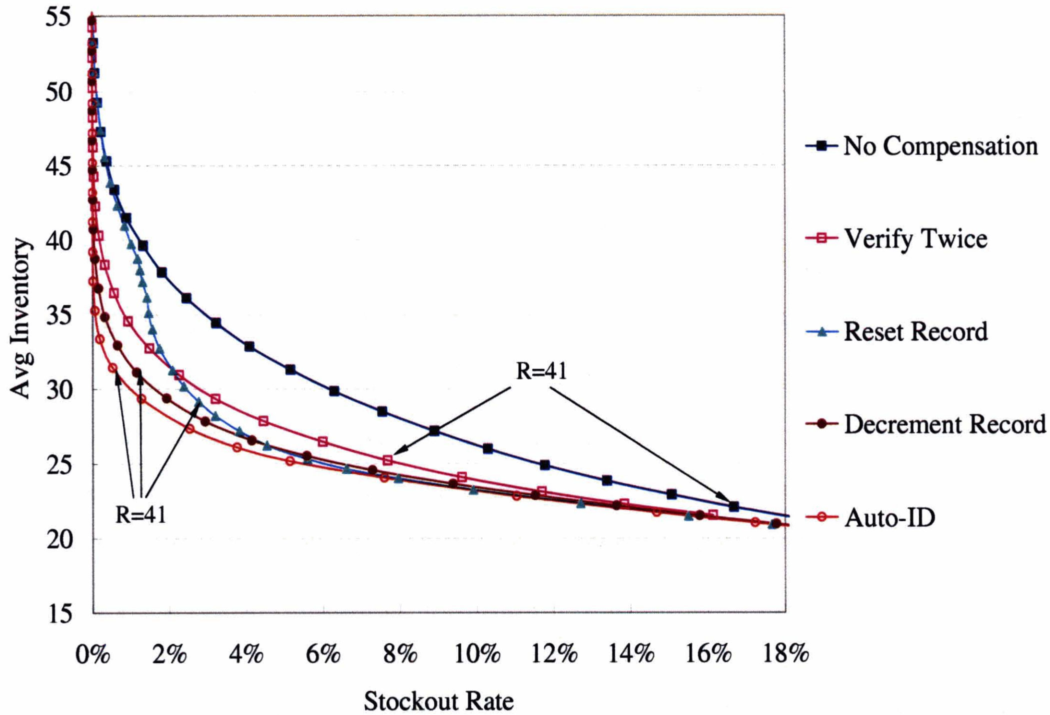

Figure 4-1: Stockout rate and average inventory for various compensation methods

dramatically, primarily it prevents replenishment freeze.

The 'Reset Record' curve is the result of resetting the inventory record to zero when sales are zero (Section 4.1.3). Notice that the vertical distance from the 'No Compensation' curve is large for low reorder points but is almost zero for high reorder points. This is because at low reorder points, inventory is small on average and zero sales occur frequently. Thus, the POS provides useful information needed to correct the inventory record error. At high reorder points, zero sales are infrequent, and the system behaves close to the 'No Compensation' case.

The strategy of decrementing the inventory record daily by the average stock loss (Section 4.1.4 — shown by the 'Decrement Record' curve) performs remarkably well in improving the stockout rate-inventory compromise from the no compensation case. Simply reducing the inventory record value by a constant amount each period still leaves errors in the record, but over time the record is able to track the actual inventory much more closely and keep the out-of-stocks low.

As expected, having a perfectly accurate knowledge of the on-hand inventory (Section 4.1.5 — the 'Auto-ID' curve) achieves the best stockout rate-inventory compromise: Auto-ID is able to attain the lowest inventory for any given stockout rate. The benefit of having the accurate knowledge of on-hand inventory becomes greater as the desired target stockout rate becomes smaller.

The effect of carrying higher safety stock can be observed from the 'No Compensation' curve. Since in the absence of stock loss the minimum reorder point required to achieve 0.5% stockout rate is 41, any reorder point higher than this can be considered safety stock for protecting the system from inventory record error and stock loss. When the stock loss demand is 1% of average purchase demand, the reorder point must be increased to at least 73 to maintain the stockout rate at 0.5%. This means the safety stock will have to increase by more than three days' worth of average purchase demand. Starting the inventory operation with higher reorder point allows more time for the actual inventory to stay above zero. However, as the gap between the actual inventory and inventory record grows and out-of-stocks begin to occur, this compensation method takes no further action to correct the error. In fact, when the stock loss is higher at 3%, the reorder point must be much higher at 145 (including a safety stock of more than eleven days' worth of average demand). This indicates that at high stock losses, the inventory required to maintain the low target stockout rate becomes prohibitive. Therefore, merely stocking up the facility with extra inventory to provide a buffer against uncertainty in inventory accuracy and stock loss is not an effective way to treat the problem.

4.3 Limitations of Each Method

The results reveal that if the stochastic behavior of stock loss is known, a significant improvement in performance can be achieved by compensating for the inventory record error. We also have seen that in some instances, such as decrementing the inventory record by average stock loss, a dramatic improvement can be made even without Auto-ID. However, the stockout rate-inventory performance is not the only measure that has to be taken into account in selecting the appropriate compensation method.

Higher safety stock, as we have seen, keeps the lost sales to a minimum only for very small stock losses, and does so at the price of carrying high inventory. For inventory inaccuracy caused by nonzero-mean error such as the stock loss considered here, this is not a desirable solution.

Manual verification of the inventory record has a number of disadvantages as well. It is costly to implement, especially in low-margin, high-competition environments where the availability of workforce is limited. In addition, manually verifying the entire facility requires shut down of the operation, which leads to loss of revenue. Targeting only a portion of the entire SKUs and cycle counting them is an alternative, but often items cannot be found in the designated locations when they are misplaced by shoppers or employees. In mass merchandise retailing environments where there are hundreds of thousands of individual items at any time, finding the items of interest during the cycle count alone becomes a challenging task. If the possibilities of mis-labeling and mis-counts are also considered, there is no guarantee that the manual counts will accurately reflect the true on-hand inventory.

The method of resetting the inventory record to zero bears the danger of false positives. This is true especially for low demand products, for which zero sales does not necessarily mean zero inventory. Incorrectly setting the inventory record to zero results in over-stocking the inventory. In our example, the use of this compensation makes sense since the probability of zero purchase demand is extremely small. For products with much lower demand, the inventory record should be reset only if a number of consecutive zero sales days are observed. However, determining the number of such days to wait until reset requires a sophisticated analysis.

Decrementing the inventory record, while simple in concept and effective in keeping the stockout rate low in our model, presents a few disadvantages as well. First of all, implementing this method can be expected to face cultural barriers in the organizations. The perpetual inventory has always been discrete, nonnegative integers. Under this method, however, the computer record could be negative and non-integer depending on how it is implemented.

More important is the sensitivity of the system performance to the stock loss demand estimate used in decrementing the inventory record. Earlier in Chapter 2, it was pointed out that even a small level of stock loss can create high stockout rates. This is tantamount to saying that if the estimated stock loss demand is slightly lower than the actual stock loss demand, the stockout rate will also be high. To study this sensitivity, another set of simulations were run. Whereas in the previous simulation the average stock loss demand was assumed to be known exactly, now we assume the inventory manager makes an incorrect estimate of the true stock loss demand. Figure 4-2 shows how the stockout rate and average inventory change when the average stock loss demand is estimated to be 3% of average

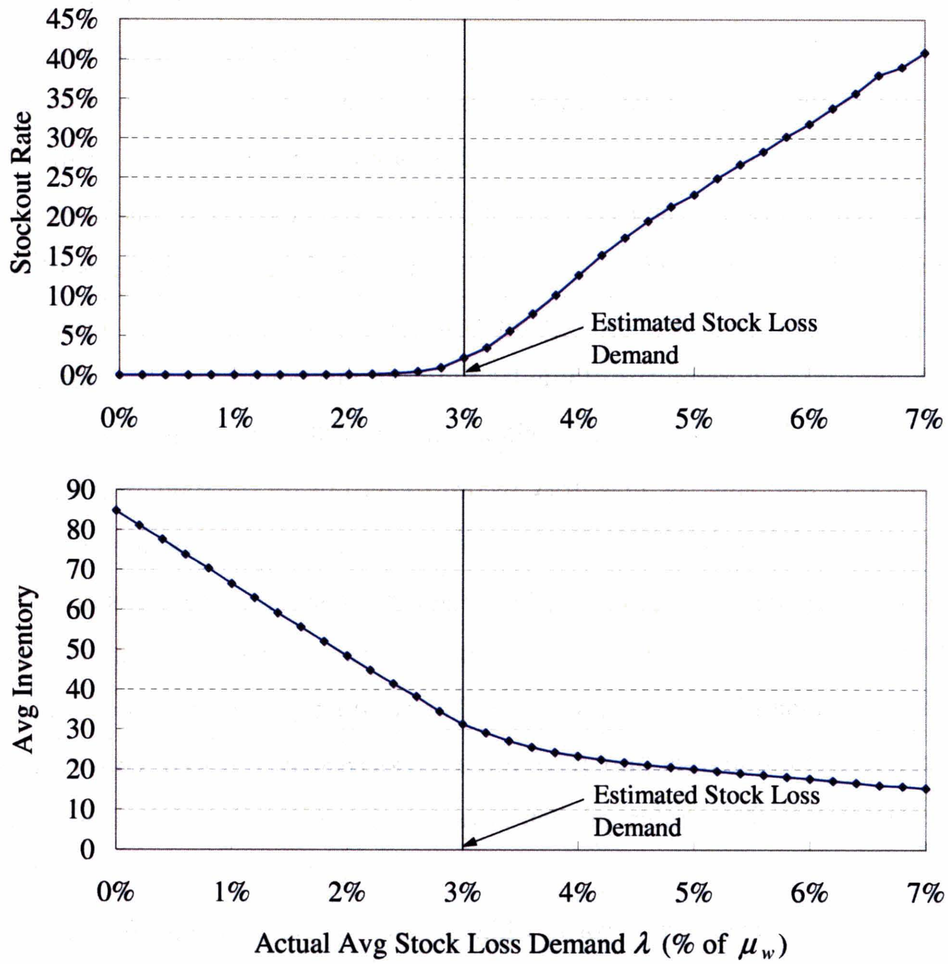

Figure 4-2: Stockout rate and average inventory when the estimated stock loss is incorrect in using the inventory record decrement strategy

demand but the true average stock loss demand varies from 0 to 7%.

If the estimated average stock loss demand of 3% is equal to the actual stock loss demand occurring in the facility, then this compensation method performs well in adjusting the inventory error and achieves a relatively low stockout rate of 2.2%. However, as the amount by which the estimated value underestimates the actual stock loss demand grows, stockout rate rises rapidly. In other words, to the right of 3% stock loss, stockout rate performance exhibits sensitivity similar to what is observed in situations where the inventory error is left uncorrected. There is a difference in how the stockout rate rises, however. In the case where no compensation is applied, stockout rate rises even more sharply when the stock loss is high enough to create replenishment freeze. Here, we do not observe such change in the stockout rate curve slope: the stockout rate rises more or less at a constant rate with increasing actual stock loss demand. This is because even when the system freezing takes place, decrementing the inventory record daily will eventually bring the inventory record below the reorder point, thus setting the replenishment back into action. Therefore, another benefit of the decrement strategy is that it prevents the replenishment freeze from taking place, and thus eliminates the extreme out-of-stocks.

The performance of the system suffers if the estimated stock loss demand is higher than the actual stock loss demand as well. In this case, stockout rate drops to zero, but the average inventory in the system rises rapidly as actual stock loss decreases. When the estimate is off by 2%, the average inventory is more than twice as high as estimating the stock loss correctly at 3%. Therefore, the ability of the inventory record decrement strategy to effectively compensate for the inventory error depends critically on the accuracy of the stock loss estimate. Even a small deviation from the actual stock loss demand will result in either high stockout rate or unnecessary inventory in the facility.

Auto-ID requires high up-front investment in RFID readers and tags, in addition to the costs involved in design and execution of real-time inventory tracking algorithm and software. Moreover, being an emerging technology still under development, there is no guarantee that Auto-ID will work perfectly and provide an exact account of the actual stock quantity in the store.

Chapter 5

Inventory Control Using Auto-ID — Perfect State Information

The radio frequency identification (RFID) technology was presented in the introductory part of this thesis as a motivation for research in inventory inaccuracy. By automatically identifying the presence of physical objects using the RFID tags and readers, the technology can potentially provide what is coveted by inventory managers and what has been assumed in many inventory models — accurate knowledge of the on-hand stock quantity. Furthermore, the benefit of this knowledge in terms of improving the inventory system performance for (Q,R) policy was examined in the previous chapter.

We now probe deeper into the problem of compensating for inventory record error using Auto-ID[1]. In this chapter, we assume Auto-ID works perfectly and provides accurate information of the on-hand quantity. Under this assumption, we find the optimal ordering policy for an inventory system subjected to stock loss. In the next chapter, we make a more realistic assumption that Auto-ID does not work perfectly and has errors in the measurement of on-hand quantity. We present a model of the measurement error, the optimal ordering policy, and a few suboptimal control schemes that are more readily implementable.

Under the assumption that the stock quantity is known exactly, the optimal ordering policy for a system subjected to demand has been treated in many works [Ber00]. Here, we examine how the optimal policy changes when a stock loss is introduced in the system. Specifically, we examine the impact of stock loss on the structure of the ordering policy and

[1]As mentioned earlier, by 'Auto-ID' we refer to all means of automatic capture of the on-hand quantity. Thus, it is not restricted to the technology under development at the Auto-ID Center.

look for the means to compute the optimal profit. This would also allow us to know the upper bound of the performance by Auto-ID and provide a basis of comparison for Auto-ID systems with varying degrees of measurement errors.

5.1 Discrete Time and Continuous State

Consider the inventory control problem of determining the optimal order quantity at each period over N finite horizon so as to maximize the profit when demand for purchase and stock loss is present. Let us denote the following variables of the system:

x_k : inventory at the beginning of period k,

u_k : quantity ordered in period k,

w_k : demand for purchase during period k,

v_k : demand for stock loss during period k.

The variable x_k is the system's state. It evolves over time under the influence of decision made at each period u_k and the random parameters w_k and v_k.

In each period k, the events occur in the following order:

1. Inventory is reviewed and order is placed to the supplier
2. Ordered quantity is received immediately (zero lead time)
3. Demand for purchase is filled
4. Stock loss takes place

We further assume that the demand $w_0, w_1, \ldots, w_{N-1}$ and the stock loss $v_0, v_1, \ldots, v_{N-1}$ are independent random variables that take values from a bounded interval, and that excess demand is lost (i.e., no backlog).[2] Thus the inventory evolves according to the discrete-time equation

$$x_{k+1} = \max(0, x_k + u_k - w_k - v_k), \qquad k = 0, 1, \ldots, N-1 \tag{5.1}$$

Note the max term is necessary since no backlog (i.e. no negative inventory) is allowed in the model.

[2]Note that this formulation is equivalent to the inventory problem in which there are two or more customer price classes, since the stock loss can be treated as just another demand but with different cost implications [CKL88].

We search for an optimal ordering policy that maximizes the total profit raised over time. Profit per stage is defined as revenue minus the total expenditure, and since maximizing the profit is equivalent to minimizing the negative of the profit, we express the cost per stage g_k, which is function of x_k, u_k, w_k, and v_k, as follows

$$\begin{aligned} g_k(x_k, u_k, w_k, v_k) &= -\text{profit} \\ &= [\text{Expenditure}] - [\text{Revenue}]. \end{aligned} \tag{5.2}$$

The expenditure consists of two components: first, the purchasing cost cu_k where c is the per unit price paid to the supplier, and second, the inventory holding cost $h\max(0, x_k + u_k - w_k - v_k)$ where h is the per unit cost of holding inventory left at the end of the period. Similarly, the revenue generated at each period is the selling price p multiplied by the number of units sold. The number of units sold in period k is either equal to the demand w_k (if there is sufficient stock to meet the demand) or the inventory available just prior to meeting the demand $x_k + u_k$ (if there is out-of-stock). The per stage cost, g_k, then becomes

$$g_k(x_k, u_k, w_k, v_k) = cu_k + h\max(0, x_k + u_k - w_k - v_k) - p\min(w_k, x_k + u_k). \tag{5.3}$$

Given an initial inventory x_0 and zero terminal cost, the optimization problem is to minimize the total expected cost accumulated over the N periods, denoted $J_0(x_0)$

$$J_0(x_0) = E\left\{\sum_{k=0}^{N-1} g_k(x_k, u_k, w_k, v_k)\right\} \tag{5.4}$$

by properly choosing the set of order quantities $u_0, u_1, \ldots, u_{N-1}$ subjected to the constraint $u_k \geq 0$ for all k.

Dynamic programming is used as an optimization algorithm since we are searching for an optimal decision to be made in stages. The appropriateness of dynamic programming is evident as the policy is implemented in closed-loop form — i.e., placing the order u_k is postponed until the last possible moment (time k) so that all the information that becomes available since time 0 can be utilized [Ber00]. Since there is no penalty for delaying the decision until time k, we can take advantage of the latest information available, which is the demand for purchase in the past periods, $w_0, w_1, \ldots, w_{k-1}$, and the stock level at time 0 up to k. This is different from an open-loop minimization scheme in which all order quantity would be determined at time 0 without seeing the subsequent information.

According to dynamic programming, the optimal cost $J_0^*(x_0)$ for a given initial state x_0 is equal to $J_0(x_0)$ generated at the last step of the following expression, which proceeds backward in time from period $N-1$ to 0:

$$J_N(x_N) = 0, \tag{5.5}$$

$$J_k(x_k) = \min_{u_k} \underset{w_k, v_k}{E} \{g_k(x_k, u_k, w_k, v_k) + J_{k+1}(x_{k+1})\}, \quad k = 0, 1, \ldots, N-1 \tag{5.6}$$

where the expected value is taken with respect to the probability distribution of w_k and v_k. $J_k(x_k)$ is the 'cost-to-go' from time k to the ending time N, and consists of the per stage cost g_k and the cost-to-go from the immediately following time $k+1$. The set of the order quantities $u_0^*, u_1^*, \ldots, u_{N-1}^*$ that minimizes the right side of Equation (5.6) is the optimal policy.

Substituting g_k in the above expression with Equation (5.3) and x_{k+1} with Equation (5.1), the DP algorithm becomes

$$J_N(x_N) = 0,$$

$$\begin{aligned} J_k(x_k) = \min_{u_k} \underset{w_k, v_k}{E} \{&cu_k + h\max(0, x_k + u_k - w_k - v_k) - p\min(w_k, x_k + u_k) \\ &+ J_{k+1}(\max(0, x_k + u_k - w_k - v_k))\}, \quad k = 0, 1, \ldots, N-1 \end{aligned} \tag{5.7}$$

If the expected value of the quantity in the bracket above (hence the cost-to-go function $J_k(x_k)$) is convex in x_k and has a minimum with respect to x_k, denoted by S_k, the optimal policy has the form

$$u_k^*(x_k) = \begin{cases} S_k - x_k & \text{if } x_k < S_k, \\ 0 & \text{otherwise.} \end{cases} \tag{5.8}$$

It turns out that convexity holds. Refer to Appendix C for the complete proof.

The optimal quantity then follows the order-up-to-level policy which attempts to maintain the inventory at the target value S_k. If the beginning inventory at period k, x_k, is lower than S_k, the optimal order quantity would be what brings x_k up to S_k. If it is greater than S_k, then no order is placed (See Figure 5-1). This is equivalent to the result from the classical inventory control problem in which the system is subject to only the demand for purchase and there is no stock loss [Ber00].

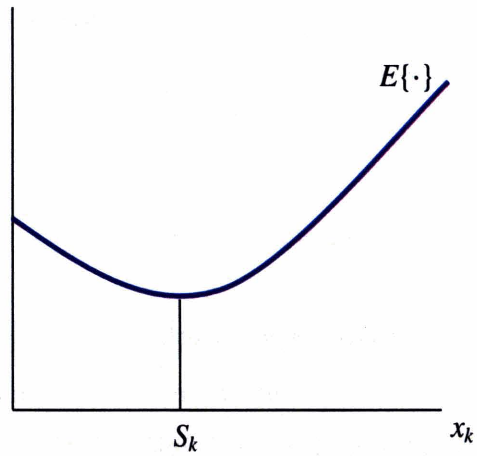

Figure 5-1: Structure of the cost-to-go functions

5.2 Discrete Time and Discrete, Finite State

In the preceding section and in most of the research literature, inventory is represented by a continuous, real variable. In this section, stock is measured in whole units (such as water bottles or razor blades). In addition, we assume that there is a physical upper bound on the possible number of units that can be stored in the facility at one time. Thus, the state space is defined in this section to be a finite set of integers rather than real numbers. We do this in order to compare results with those of Chapter 6 in which these assumptions are required to make the problem numerically tractable.

If the maximum allowed stock is n, the dynamics of the system described in the previous section (Equation (5.1)) changes to

$$x_{k+1} = \max\big(0, \min(n, x_k + u_k) - w_k - v_k\big). \tag{5.9}$$

The min term places a constraint on the total available inventory just prior to filling the demand, $x_k + u_k$, to be no greater than n. The per stage cost becomes

$$g_k(x_k, u_k, w_k, v_k) = cu_k + h\max(0, \min(n, x_k + u_k) - w_k - v_k) - p\min(w_k, \min(n, x_k + u_k)) \tag{5.10}$$

and the dynamic programming equation becomes

$$\begin{aligned} J_N(x_N) &= 0, \\ J_k(x_k) &= \min_{u_k} \mathop{E}_{w_k,v_k} \Big\{ cu_k + h\max(0, \min(n, x_k+u_k) - w_k - v_k) \\ &\quad - p\min(w_k, \min(n, x_k+u_k)) + J_{k+1}(\max(0, \min(n, x_k+u_k) - w_k - v_k)) \Big\}. \end{aligned} \tag{5.11}$$

Since the state space is a finite set of integers from 0 to n, the expectation can be carried out with respect to w_k and v_k. Assuming w_k and v_k can take values from the finite set $[0, m]$ and $[0, q]$, respectively, the above algorithm can be expressed as

$$\begin{aligned} J_N(x_N) &= 0, \\ J_k(x_k) &= \min_{u_k} \sum_{i=0}^{m} \sum_{j=0}^{q} \Big\{ cu_k + h\max(0, \min(n, x_k+u_k) - i - j) \\ &\quad - p\min(i, \min(n, x_k+u_k)) \\ &\quad + J_{k+1}\big(\max(0, \min(n, x_k+u_k) - i - j)\big) \Big\} P(w_k = i) P(v_k = j). \end{aligned} \tag{5.12}$$

The optimal policy and the corresponding optimal cost can be obtained by first minimizing the right-hand side of Equation (5.12) for every possible value of the state x_{N-1} to obtain the optimal order quantity for period $N-1$, $u^*_{N-1}(x_{N-1})$. Knowing $u^*_{N-1}(x_{N-1})$, $J_{N-1}(x_{N-1})$ can be computed and used in the minimization of the right-hand side of the same Equation (5.12) for period $N-2$, which is carried out for every possible value of x_{N-2}. In the same manner, the cost-to-go function $J_{N-3}(x_{N-3})$ and the optimal quantity $u^*_{N-3}(x_{N-3})$ are obtained, and so forth, until the optimal cost $J^*_0(x_0)$ is computed. Since the profit is defined to be negative of the cost, the optimal profit achieved by the optimal policy is -$J^*_0(x_0)$.

The expected value of optimal profit calculated by the DP algorithm for an example numerical problem is shown in Figure 5-2. It is assumed that the maximum allowed inventory, n, is 10, and the cost parameters of $c = 1, h = \frac{.25c}{365}, p = 1.25$ are used. Under these settings, products are sold at a value 25% higher than the purchase price, and the inventory holding cost is 25% of the purchase price per year. The demand w_k is assumed to have binomial

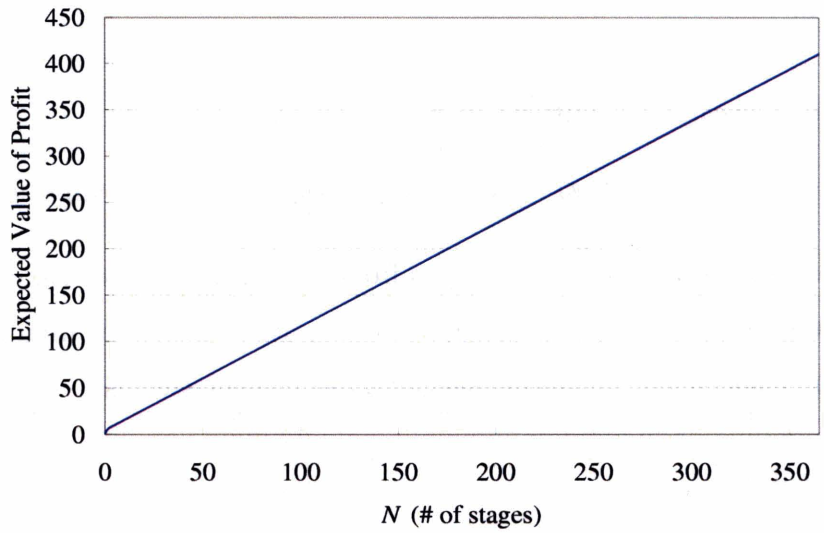

Figure 5-2: Expected profit computed by dynamic programming

distribution characterized by the parameters n_b and p_b as follows[3]

$$P(w_k = i) = \binom{n_b}{i} p_b^i (1 - p_b)^{n_b - i} \qquad\qquad i = 0, 1, \ldots, m. \tag{5.13}$$

The average demand is set at 5 by selecting n_b and p_b to be 10 and 0.5, respectively. Poisson distribution is used for the stock loss, v_k, with the mean λ equal to 3% of the average demand ($\lambda = 0.3$). The maximum possible value of both the demand and stock loss — m and q, respectively — are equal to n. The initial state, x_0, has uniform probability distribution over the values [0,n].

The optimal policy for the $N = 365$ problem is shown in Table 5.1, which lists the optimal order quantity for each period k and for each possible value of beginning inventory x_k. For example, if in period $k = 10$ the beginning inventory $x_{10} = 0$, the optimal order quantity is 7. If in period $k = 364$ the beginning inventory $x_{364} = 4$, the optimal order quantity is 2. It can be seen that the order-up-to-level policy is preserved in the discrete, finite state inventory system with stock loss. The order-up-to-level S_k is, however, non-

[3]The binomial distribution is chosen because it closely approximates the normal distribution, and thus consistency with the models in the previous chapters is provided.

Table 5.1: Optimal order quantity

x_k	k		
	0, 1, ..., 363	364	365
0	7	6	4
1	6	5	3
2	5	4	2
3	4	3	1
4	3	2	0
5	2	1	0
6	1	0	0
7	0	0	0
8	0	0	0
9	0	0	0
10	0	0	0

stationary over the periods according to

$$S_k = \begin{cases} 7 & \text{for} \quad 0 \leq k \leq 363 \\ 6 & \text{for} \quad k = 364 \\ 4 & \text{for} \quad k = 365. \end{cases} \tag{5.14}$$

This is because the terminal cost is defined to be zero and thus greater penalty is placed on leftover inventory towards the end of the operation. For periods $k \leq 363$, the effect of this increased penalty is not seen since the leftover inventory in a period can be used to fulfill the demand in the following period. Moreover, the cost parameters c, h, and p are constant, making the order-up-to-level S_k in these periods to be stationary.

This optimal policy structure also explains why the expected value of profit, shown in Figure 5-2, is not linear with respect to N. For short horizons, the order-up-to-level is lower than in longer horizons, resulting in different ordering and inventory cost and thus different expected profit per period. However, for long horizons, the order-up-to level is constant, and with the cost parameters also being constant, a linear relationship yields a good approximation of the expected total cost.

Chapter 6

Inventory Control Using Auto-ID — Imperfect State Information

In the preceding chapter, we considered an inventory control problem in which the state of the system — the on-hand stock quantity — is known accurately. Under the assumption that Auto-ID can obtain a perfectly accurate measurement, the perfect state information dynamic programming allows us to find the optimal ordering policy for an inventory system subject to stock loss. However, this assumption is unrealistic. All real systems that make observation of the system variables are subjected to measurement errors, and Auto-ID is not an exception. In this chapter, we define and model an imperfect state information system in which the Auto-ID measurement is erroneous in nature and obtain an optimal control scheme. Unfortunately, the fast-growing state dimension of the imperfect state information problem limits us to solve only small size, simplified problems. Later we present sub-optimal control schemes that do not require a computationally burdensome optimization and yet are able to achieve a near-optimal performance.

6.1 Optimal Control for Auto-ID with Measurement Errors

To see how Auto-ID measurements can have errors, a closer look at how the RFID technology works is needed.

The communication between RFID reader and tag begins when the reader sends an electromagnetic wave at a certain frequency and power (See Figure 6-1). If the tag is within the reading range of the reader, which can typically vary from a few inches to yards, the

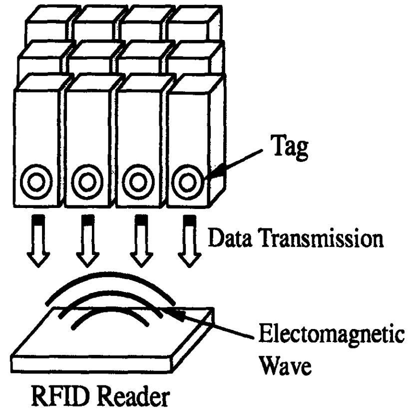

Figure 6-1: Communication between the RFID readers and tags

electromagnetic wave powers the tag, which transmits back to the reader the information stored in the microchip (such as the unique identification code). By gathering all the signals transmitted from the tags to all the readers in the facility, the physical inventory reading is obtained.

Unfortunately, the data transmission between the reader and tags is subject to the uncertainties in the media or channels through which the signal has to pass. First of all, the reading range of the reader can fluctuate depending on the environmental factors, leading to possible failure to capture some of the tagged objects. Even if all the tagged objects reside in the reading range, signal interference and distortion can arise due to many factors such as the material content of the tagged object, their orientation, spacing between the objects, and etc. For example, a high read rate can be achieved by the RFID reader for products such as paper towel and sponge. For products with high fluidic and metallic content — such as soup cans and liquid detergents — the read rate can be much lower. The result of this data corruption is a discrepancy between the actual stock and the measurement reported by Auto-ID.

One important characteristic of this measurement error is that while the RFID system can miss reading some of the tagged objects, it does not identify an object that does not exist. Therefore, the measurement by Auto-ID can only understate the actual quantity.

We model this inventory system using a class of control problem called imperfect state information problems. Whereas before the state was assumed to be known exactly, now we assume what can be accessed instead is a stochastically uncertain observation about the stock quantity in each period. Let us denote the measurement of the state at time k made by Auto-ID as z_k.

Since the state x_k is not accessible, we need to define a new state at time k to be the set of all available information the knowledge of which can be of benefit in looking for the optimal order quantity. In fact, the set of all available information at time k includes all the observations from time 0, $(z_0, z_1, \dots, z_k)$, and all the past controls $(u_0, u_1, \dots, u_{k-1})$. Call this the *information vector* I_k

$$\begin{aligned} &I_k = (z_0, z_1, \dots, z_k, u_0, u_1, \dots, u_{k-1}), \quad k = 1, 2, \dots, N-1 \\ &I_0 = z_0. \end{aligned} \tag{6.1}$$

Next, the cost per stage needs to be reformulated as a function of the variables of the new system. This can be done by taking the conditional expected value of the original cost per stage, $g_k(x_k, u_k, w_k, v_k)$, given the information vector I_k and the control u_k. The reformulated cost per stage, $\tilde{g}_k(I_k, u_k)$, is

$$\tilde{g}_k(I_k, u_k) = \underset{x_k, w_k, v_k}{E} \{g_k(x_k, u_k, w_k, v_k) | I_k, u_k\}. \tag{6.2}$$

Using the expressions (6.1) and (6.2), the DP expression can be written as

$$\begin{aligned} &J_N(I_N) = 0, \\ &J_k(I_k) = \min_{u_k} \underset{x_k, w_k, v_k, z_{k+1}}{E} \{g_k(x_k, u_k, w_k, v_k) + J_{k+1}(I_k, z_{k+1}, u_k) | I_k, u_k\}, \\ &\qquad k = 0, 1, \dots, N-1. \end{aligned} \tag{6.3}$$

Thus the imperfect state information problem has been reformulated as a problem with perfect state information. The optimal policy $\{u_0^*, u_1^*, \dots, u_{N-1}^*\}$ for this problem is obtained in the same manner by performing minimization backwards in stages.

Even though the imperfect state information DP provides a means to determine the optimal policy for a problem subjected to measurement errors, it requires highly intensive computation. This difficulty arises from the dimension of state I_k that expands over time. As a new measurement is made at each stage, the dimension of I_k grows, and even with small number of possible values that z_k and u_k can take, the total possible values of I_k increases rapidly with time. For example, if z_k and u_k are each allowed to take ten different values, I_N can have 10^{19} possible values by the tenth stage ($N = 10$).

Therefore it is of interest for us to look for an alternative system state which would have smaller dimension than I_k and yet provide all the essential content of I_k that is necessary for control purposes. Such quantities are known as *sufficient statistics* [Ber00]. Among many, one sufficient statistic that is useful for our model is the conditional probability distribution of the true state x_k given the information vector I_k, denoted P_k. P_k is a vector whose elements are conditional probabilities that x_k is equal to i, $i = 0, 1, \ldots, n$, given the information vector I_k. Denoting this conditional probability by $p_k^{(i)}$, the vector P_k is written as

$$P_k = \begin{bmatrix} \text{Prob}(x_k = 0|z_0, \ldots, z_k, u_0, \ldots, u_{k-1}) \\ \text{Prob}(x_k = 1|z_0, \ldots, z_k, u_0, \ldots, u_{k-1}) \\ \vdots \\ \text{Prob}(x_k = n|z_0, \ldots, z_k, u_0, \ldots, u_{k-1}) \end{bmatrix} = \begin{bmatrix} p_k^{(0)} \\ p_k^{(1)} \\ \vdots \\ p_k^{(n)}. \end{bmatrix} \tag{6.4}$$

Using the basic laws of probability, we can show that the elements of the above vector, and thus the vector itself, can be generated recursively. That is, the state at stage $k+1$, P_{k+1}, can be expressed as function of the state in the previous stage P_k, the control applied in the previous stage, u_k, and the measurement available at present stage, z_{k+1}. See Appendix D.1 for the mathematical steps. The result is

$$\begin{aligned} p_{k+1}^{(j)} &\equiv \text{Prob}(x_{k+1} = j|z_0, z_1, \ldots, z_{k+1}, u_0, u_1, \ldots, u_k) \\ &= \frac{\sum_{i=0}^{n} p_k^{(i)} \pi_{i,j}(u_k) r_{z_{k+1},j}}{\sum_{s=0}^{n} \sum_{i=0}^{n} p_k^{(i)} \pi_{i,s}(u_k) r_{z_{k+1},s}} \\ &= \Phi_k(P_k, u_k, z_{k+1}) \end{aligned} \tag{6.5}$$

where $\pi_{i,j}(u_k)$ and $r_{i,j}$ are state transition probability and measurement probability, respectively, defined as

$$\begin{aligned} \pi_{i,j}(u_k) &\equiv \text{probability that } x_k \text{ moves from } i \text{ to } j \text{ when control } u_k \text{ is applied} \\ &= \text{Prob}(x_{k+1} = j|x_k = i, u_k) \end{aligned} \tag{6.6}$$

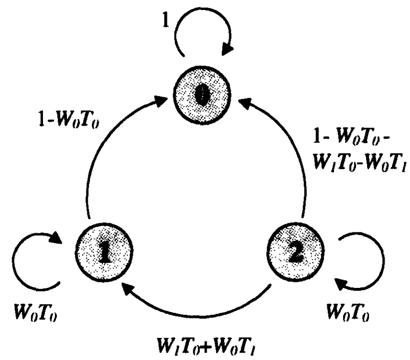

Figure 6-2: Markov chain diagram for state transition

and

$$r_{i,j} \equiv \text{probability that measurement } z \text{ is } i \text{ when the actual stock quantity } x \text{ is } j$$
$$= \text{Prob}(z_k = i | x_k = j). \tag{6.7}$$

Therefore, the system state at a given period is some function Φ of the previous state, the control applied in the previous state, and the current measurement. Using this recursive relationship in Equation (6.5), we can construct the system state at any period from a given initial state P_0.

The state transition probability $\pi_{i,j}(u_k)$ can be determined from the probability distributions of w_k and v_k and the system dynamics described by Equation (5.9). To illustrate how this can be done, let us use a simple example in which x_k, u_k, w_k, and v_k all can take values only from a set of finite integers {0,1,2}. Further denote by W_i the probability that the demand for purchase $w_k = i$, $i = 0, 1, 2$, and denote by T_i the probability that the demand for stock loss $v_k = i$. Let $u_k = 0$. Applying the dynamics of the system expressed in (5.9), the probabilities of x_k moving from one value to another can be shown in the Markov chain diagram in Figure 6-2.

For example, if x_k is 1 at the beginning of time k and no order was placed during the period k ($u_k = 0$), then the probability that x_{k+1} will remain unchanged is 1 only when there is zero demand and zero stock loss. This transition probability is $\pi_{1,1}(0) = W_0T_0$. Similarly, x_k will drop to 0 at the beginning of $k+1$ if the sum of w_k and v_k is at least one. This transition probability is $\pi_{1,0}(0) = 1 - W_0T_0$. The probability that x_{k+1} will increase to 2 is zero since nothing is ordered. Other transition probabilities can be determined similarly.

The transition probabilities can be summarized by the transition probability matrix $\Pi(u_k)$, and for the simple example shown in Figure 6-2, it is

$$\begin{aligned}\Pi(u_k=0) &= \begin{bmatrix}\pi_{0,0}(0) & \pi_{0,1}(0) & \pi_{0,2}(0)\\ \pi_{1,0}(0) & \pi_{1,1}(0) & \pi_{1,2}(0)\\ \pi_{2,0}(0) & \pi_{2,1}(0) & \pi_{2,2}(0)\end{bmatrix}\\ &= \begin{bmatrix}1 & 0 & 0\\ 1-W_0T_0 & W_0T_0 & 0\\ 1-W_0T_0-W_1T_0-W_0T_1 & W_1T_0+W_0T_1 & W_0T_0\end{bmatrix}\end{aligned} \tag{6.8}$$

. The transition probabilities for other values of u_k can be computed similarly.

Using the newly defined system state P_k, the DP equation becomes

$$\begin{aligned}&J_N(P_N)=0\\ &J_k(P_k)=\min_{u_k}\underset{x_k,w_k,v_k,z_{k+1}}{E}\{cu_k+h\max(0,\min(n,x_k+u_k)-w_k-v_k)\\ &\qquad\qquad -p\min(w_k,\min(n,x_k+u_k))+J_{k+1}(\Phi_k(P_k,u_k,z_{k+1}))|I_k,u_k\}\\ &\qquad\qquad k=0,1,\dots,N-1.\end{aligned} \tag{6.9}$$

Note that whereas before the system state at time k was collection of all the available prior and present measurements and controls, now the system state is the conditional probability of the actual stock quantity given the information vector. The per stage cost remains the same as in the problem using I_k as the state, but the cost-to-go term now involves the recursively-generated P_k. Carrying out the expected values and expressing in matrix form, the above algorithm can be expressed as

$$\begin{aligned}&J_N(P_N)=0,\\ &J_k(P_k)=\min_{u_k}\Bigg[\sum_{i=0}^{n}\sum_{j=0}^{m}\sum_{l=0}^{q}p_k^{(i)}W_jT_lg_k(x_k=i,u_k,w_k=j,v_k=l)\\ &\qquad +\sum_{j=0}^{n}\sum_{s=0}^{n}\sum_{i=0}^{n}p_k^{(i)}\pi_{i,s}(u_k)r_{j,s}J_{k+1}(\Phi_k(P_k,u_k,z_{k+1}=j))\Bigg]\\ &\qquad =\min_{u_k}\Big[P_k'G_k(u_k)+\sum_{i=0}^{n}J_{k+1}(R_i'*\Pi(u_k)'P_k)\Big],\end{aligned} \tag{6.10}$$

where

$$G_k(u_k) = \begin{bmatrix} \sum_{j=0}^{m} \sum_{l=0}^{q} W_j T_l g_k(x_k = 0, u_k, w_k = j, v_k = l) \\ \sum_{j=0}^{m} \sum_{l=0}^{q} W_j T_l g_k(x_k = 1, u_k, w_k = j, v_k = l) \\ \vdots \\ \sum_{j=0}^{m} \sum_{l=0}^{q} W_j T_l g_k(x_k = n, u_k, w_k = j, v_k = l) \end{bmatrix}, \tag{6.11}$$

R_i is the ith row of the measurement matrix R whose elements are the measurement probabilities $r_{i,j}$, expressed as

$$R = \begin{bmatrix} r_{0,0} & \cdots & r_{0,n} \\ \vdots & \ddots & \vdots \\ r_{n,0} & \cdots & r_{n,n} \end{bmatrix}, \tag{6.12}$$

and the operator $*$ represents scalar multiplication of the matrices.

Thus, a DP problem has been formulated to find the optimal policy for an inventory control problem in which the inventory manager has access to the measurement of the actual state provided by the Auto-ID system.

Solving the imperfect state information problem in Equation (6.10) requires a different technique from the one employed for the perfect state information problem in Equation (5.12). Whereas in the latter the state is discrete and can take on a finite set of integers, in this problem the state is a vector of real numbers representing probabilities. Therefore, rather than proceeding backward in time from period $N-1$ to 0, we start at the initial stage $k=0$ where we can compute the state P_0 — which is the conditional probability that the initial inventory, x_0, is equal to $i, i = 0, 1, \ldots, n$ given the measurement z_0 — and move forward.

We compute the vector P_0 based on the given problem data — namely, the probability distribution of the actual stock quantity x_0 and the measurement matrix R — and use it to determine $J_0(P_0)$ in Equation (6.10) for all the possible values of u_0. The summation loop in the last line of the expression, however, requires computing P_1 and solving for $J_1(P_1)$, again for all the possible values of the control u_1. Solving for $J_1(P_1)$ in turn requires calculation of P_2, and proceeding forward, minimization would reach the stage $N-1$, at which point $J_{N-1}(P_{N-1})$ can be determined. This allows the summation loop in the previous stages to be calculated, and the minimization of $J_0(P_0)$ can be complete for the first u_0 under consideration. In the same manner, we repeat this procedure for all the possible values of u_0.

This highly iterative computation indicates that even with the reduction of state space achieved by introducing the new state P_k, the problem of intensive computation required to carry out the DP algorithm still persists. According to Equation (6.10), each period requires $(n+1)^2$ matrix operations, resulting in $(n+1)^{2N}$ matrix operation to complete the optimization. Ideally, we'd like to obtain the solution for a problem with the horizon of $N = 365$ periods. This is not computationally feasible even with today's fastest computers, and thus we solve a five stage problem ($N = 5$) and look for any insights regarding the optimal profit and the structure of the optimal policy.

We present the solution to the imperfect state information problem of Equation (6.10) using the example from the preceding section. All the problem data are identical: $n = 10, c = 1, h = \frac{.25c}{365}, p = 1.25$. The same probability distributions for the demand w_k and stock loss v_k are also used. In the imperfect state information problem, however, one additional set of data has to be defined — the measurement matrix R. Earlier it was noted that Auto-ID system can only underestimate the actual stock quantity by failing to detect some of the objects. Therefore, the stochastic behavior of the measurement error is modelled as a geometric probability distribution, expressed as follows:

$$\begin{aligned} r_{i,j} &= \text{Prob}(z_k = i | x_k = j) \\ &= \begin{cases} \frac{(1-p_g)^{(j-i)} p_g}{\sum_{s=0}^{n} U(j-s)(1-p_g)^{(j-s)} p_g} & \text{if } i \leq j, \\ 0 & \text{otherwise,} \end{cases} \end{aligned} \tag{6.13}$$

where p_g is the geometric probability constant. The numerator in the first relation is the geometric distribution, and the denominator normalizes the probabilities. The closer the p_g is to 1, the better the accuracy of the Auto-ID measurement.

Table 6.1 shows the measurement matrices for $p_g = 0.7$ and $p_g = 0.3$. For example, when $p_g = 0.7$, the probability that Auto-ID measurement is 4 when the actual stock quantity is 5 is $r_{4,5} = 0.21$. Notice the columns of the matrices add to 1.

Figure 6-3 shows the expected value of optimal profit for Auto-ID systems with measurement accuracy of $p_g = 1$ (which corresponds to a perfect state information system), $p_g = 0.7$, and $p_g = 0.3$ when the horizon N is varied from 1 to 5 stages.

The result shows that even with poor measurement accuracies, Auto-ID can perform close to the perfect state information case. When Auto-ID can make a perfectly accurate

Table 6.1: Measurement probabilities for $p_g = 0.7$ and $p_g = 0.3$

z_k \ x_k	0	1	2	3	4	5	6	7	8	9	10
0	1.00	0.23	0.06	0.02	0.01	0.00	0.00	0.00	0.00	0.00	0.00
1		0.77	0.22	0.06	0.02	0.01	0.00	0.00	0.00	0.00	0.00
2			0.72	0.21	0.06	0.02	0.01	0.00	0.00	0.00	0.00
3				0.71	0.21	0.06	0.02	0.01	0.00	0.00	0.00
4					0.70	0.21	0.06	0.02	0.01	0.00	0.00
5						0.70	0.21	0.06	0.02	0.01	0.00
6							0.70	0.21	0.06	0.02	0.01
7								0.70	0.21	0.06	0.02
8									0.70	0.21	0.06
9										0.70	0.21
10											0.70

(a) $p_g = 0.7$

z_k \ x_k	0	1	2	3	4	5	6	7	8	9	10
0	1.00	0.41	0.22	0.14	0.09	0.06	0.04	0.03	0.02	0.01	0.01
1		0.59	0.32	0.19	0.12	0.08	0.05	0.04	0.03	0.02	0.01
2			0.46	0.28	0.18	0.12	0.08	0.05	0.04	0.03	0.02
3				0.39	0.25	0.17	0.11	0.08	0.05	0.04	0.03
4					0.36	0.24	0.16	0.11	0.08	0.05	0.04
5						0.34	0.23	0.16	0.11	0.07	0.05
6							0.33	0.22	0.15	0.11	0.07
7								0.32	0.22	0.15	0.10
8									0.31	0.22	0.15
9										0.31	0.21
10											0.31

(b) $p_g = 0.3$

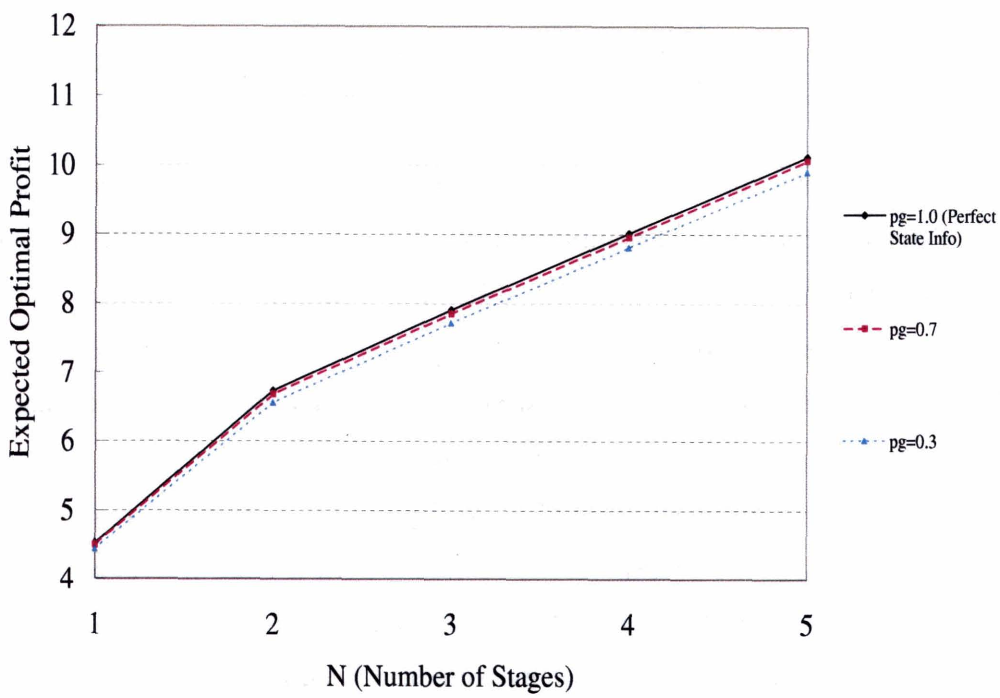

Figure 6-3: Expected profit for various measurement accuracies

observation of the actual stock quantity only 70% of the times for most possible values of x_k, it is able to attain a profit almost equal to that of perfectly accurate measurement system. When $p_g = 0.3$, the optimal profit is only 2.2% lower than that of perfect state information system for $N = 5$. This ability of the poor measurement systems to perform well is due to knowing the dynamics of the system and the stochastic behavior of the demand and stock loss.

However, this ability comes at a price. We know from the results of the perfect state information problem that the optimal order quantity follows a nice and simple order-up-to-level policy (shown in Table 5.1). However, as the measurement accuracy decreases, this simple ordering structure no longer holds. Table 6.2 shows the optimal policy that achieves the maximum profit that appears in Figure 6-3 for $p_g = 0.7$ and $p_g = 0.3$ system when the inventory horizon $N = 5$. The table summarizes the quantity that should be ordered to the supplier for every possible value of the measurement z_k, and it turns out that this quantity also depends on the previous measurement z_{k-1} for $k \geq 1$. The optimal quantities for the first four periods ($k = 0, 1, 2, 3$) are shown for the $p_g = 0.7$ system and first three periods ($k = 0, 1, 2$) for the $p_g = 0.3$ system.

First of all, we see that to achieve the optimal profit, the inaccurate measurement systems in general order less than the perfect information system for the same measurements obtained. For example, when the measurement at the first time period $k = 0$ is 4, the perfect information system should order 3 units, while the $p_g = 0.7$ system should order 2 units and the $p_g = 0.3$ system should order only 1 unit. This is because the inaccurate systems underestimate the actual stock quantity, and to maintain the same target inventory (the order-up-to-level), the inaccurate systems need to order less than the perfect information system would order, given the same measurements.

Another noticeable difference in the optimal policies is in the structure. In the $p_g = 0.7$ case, the optimal policy starts with an order-up-to-level of 6 at $k = 0$. However, as time passes, the order-up-to-level policy is no longer valid. Also note that whereas in the perfect state information case the optimal order quantity depended only on the measurement made at the present stage, in the imperfect state information case it depends also on the previous measurements as well. This is seen when an unusually high measurement in one stage is followed by a low measurement in the next. For instance, when the measurement of $z_0 = 10$ is made in stage $k = 0$, the optimal order quantity in the next stage for measurement of

Table 6.2: Optimal order quantities for $p_g = 0.7$ and $p_g = 0.3$

z_k	k=0	k=1				k=2				k=3		
		If $0 \leq z_0 \leq 7$	If $z_0=8$	If $z_0=9$	If $z_0=10$	If $0 \leq z_1 \leq 7$	If $z_1=8$	If $z_1=9$	If $z_1=10$	If $0 \leq z_2 \leq 8$	If $z_2=9$	If $z_2=10$
0	6	7	6	6	5	7*	6	6	5	6	5	4
1	5	6	5	5	5	6	5	5	5	5	5	4
2	4	5	5	4	4	5	5	4	4	4	4	4
3	3	4	4	4	3	4	4	4	3	3	3	3
4	2	3	3	3	3	3	3	3	3	2	2	2
5	1	2	2	2	2	2	2	2	2	1	1	1
6	0	1	1	1	1	1	1	1	1	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0	0

(a) $p_g = 0.7$ * 6 if $z_0=8$

z_k	k=0	k=1					k=2					
		If $0 \leq z_0 \leq 6$	If $z_0=7$	If $z_0=8$	If $z_0=9$	If $z_0=10$	If $0 \leq z_1 \leq 3$	If $z_1=4,5,7$	If $z_1=6$	If $0 \leq z_1 \leq 8$	If $z_1=9$	If $z_1=10$
0	5	6	5	5	4	3	6	6	6	5	4	3
1	3	5	4	4	4	3	5	5	5	4	4	3
2	2	4	4	3	3	3	4	4	4	4	3	3
3	1	3	3	3	3	2	3	3	3	3	3	2
4	1	2	2	2	2	2	2	2	2	2	2	2
5	0	1	1	1	1	1	1	1	2	1	1	1
6	0	0	0	0	0	0	0	1*	1	1	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0	0

(b) $p_g = 0.3$ * 0 if $z_0=9$

$z_1 = 0$ is 5, compared to the optimal quantity of 7 if the previous measurement z_0 is seven or less. This is due to the fact that Auto-ID cannot report a value that is greater than the actual inventory, and a measurement of 10 — the maximum possible inventory — is always accurate. Since the probability that the actual inventory will move from 10 to 0 is very unlikely, the controller simply does not believe that the measurement of $z_1 = 0$ is accurate, and thus recommends a quantity less than it would if the previous measurements were lower.

The optimal policy for the $p_g = 0.3$ case exhibits similar characteristics, with a much further deviation from the order-up-to-policy and a stronger dependence of the optimal order quantity on the previous measurements. In addition, the order quantities are generally even smaller. Thus, the ability of systems with poor measurement accuracies to achieve optimal performance comes at the price of a complex inventory policy.

Although Auto-ID systems with inaccurate measurements require complex policies to achieve optimality, there may exist alternative solutions that are much simpler in concept and implementation and still enable near-optimal performance. Therefore, it is of interest to examine how sensitive the performance of Auto-ID system with measurement error is to the type of the ordering policies used. We know there exists an optimal policy for each Auto-ID system characterized by the measurement inaccuracy parameter p_g. What we intend to find out is how much the performance degrades when other non-optimal policies are applied. We test the sensitivity in three different ways.

One way to test the sensitivity is making the systems with measurement inaccuracy follow the simple order-up-to-level policy derived for the perfect state information system. This may illustrate a situation in which the inventory manager, either not aware of the measurement errors in the Auto-ID system or ignoring the presence of error, chooses to follow the perfect state information policy. In the next method, we compute, for Auto-ID systems with various p_g, the expected profit that results from using the optimal policy obtained for $p_g = 0.3$. This may correspond to a situation in which the inventory manager is conservative about estimating the measurement accuracy of the Auto-ID system and decides to adopt the policy for a low-end measurement performance.

Lastly, we use the method of estimating the actual inventory by taking the expected value of the conditional probability distribution of the actual inventory given in Equation (6.5) and follow the perfect state information policy by treating this estimate to be

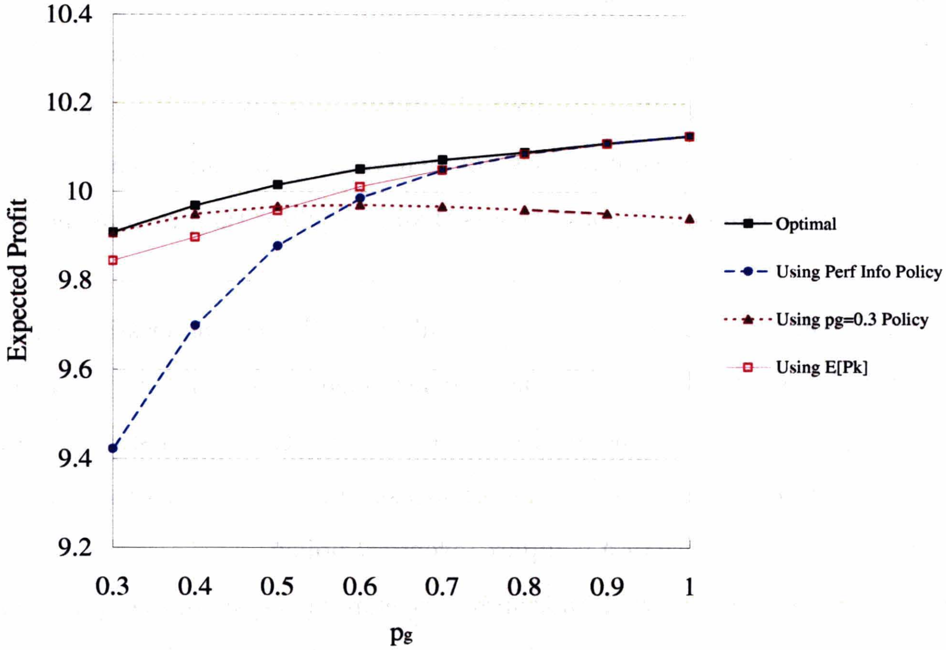

Figure 6-4: Expected profit of imperfect measurement systems subject to various policies ($N = 5$)

exact. The estimated inventory at the beginning of period $k+1$, denoted $\widehat{x}_{k+1}$, is obtained by the matrix multiplication

$$\widehat{x}_{k+1} = [0 \ 1 \ \dots \ n]P_{k+1} \tag{6.14}$$

where the conditional probability distribution of actual inventory P_{k+1} can be calculated from the conditional probability distribution at the beginning of the previous period k, P_k, the order quantity u_k, and the measurement observed at the beginning of present period z_{k+1} according to Equation (6.5).

This method is similar to what is called *certainty equivalent* control. For several types of problems involving a linear system and a quadratic cost, the system estimates are incorporated into the control law as if they are perfect measurements of the state, and the resulting optimization is the same as for the corresponding deterministic problem [Ber00].

Figure 6-4 summarizes the changes in expected profit due to these tests for Auto-ID systems with p_g varying from 0.3 to 1.0. The horizon is $N = 5$, and the same numerical example presented earlier in this section is used here.

The 'Optimal' curve represents the optimal expected profit achievable by each imperfect measurement system from following its own optimal policy. This upper bound on performance is compared with the expected profit computed for the same set of imperfect measurement systems when they follow the order-up-to-level policy obtained for the perfect state information system shown in Table 5.1 ('Using Perf Info Policy' curve). What we observe is a decrease in the profit that becomes more and more noticeable as the measurement accuracy decreases.

The upper bound is also compared to the expected profit achieved when the systems with various p_g follow the optimal policy of the worst measurement accuracy $p_g = 0.3$ (shown by the 'Using $p_g = 0.3$ Policy' line). Similar to the case where the order-up-to-level policy is used, the drift from the optimal performance grows as the measurement accuracy moves away from $p_g = 0.3$. However, the biggest performance decrease in using the $p_g = 0.3$ policy, which takes place at $p_g = 1.0$, is much less than the biggest performance decrease in using the perfect state information policy. This is because in the $p_g = 0.3$ system, the actual stock quantity is greater than the measurement provided by the Auto-ID system (Auto-ID can only underestimate the actual quantity), and if the inventory manager believes the measurement is perfect and follows the perfect information system ordering policy (illustrated by the 'Using Perf Info Policy' curve), quantities less than the optimal would be ordered consistently. Furthermore, the penalty of lost sales is higher than the penalty of storing inventory, and thus using the perfect information system ordering policy for $p_g = 0.3$ system results in a greater loss in expected profit than using the optimal policy for $p_g = 0.3$ for the $p_g = 1$ system.

What is common in both methods, however, is flatness of the expected profit curve around the p_g for which the optimal policy is used. This means, for example, the inventory manager can incorrectly assume the p_g of the Auto-ID system to be 0.5 when the true p_g is 0.3 and still attain a near-optimal profit.

Overall, the best performing of the three methods is using the expected value of the conditional probability distribution as an estimate for the actual inventory ('Using E[P_k]' curve). This state estimation technique attains the expected profit very close to the upper boundary for the entire range of p_g considered.

Having been inspired by the performance of this state estimation technique, we explore further the applicability of this compensation technique in the following section.

6.2 Sub-optimal Compensation Using State Estimation

The dynamic programming exercise in the previous section demonstrates that the optimal performance and the ordering policy can be obtained for inventory systems with inaccurate measurements. However, due to the high computational power requirements, the size of the problem (namely, the horizon and maximum allowed inventory) had to be limited and oversimplifications in the model (such as zero lead time) had to be made. Therefore, for more realistic models, we are interested in searching for control schemes that do not require computationally burdensome optimization and yet are able to attain near-optimal performance.

One of the advantages of the state estimation technique presented in the previous section is the reduction of this computational burden. Notice in this control technique no optimization is done. Instead, the state, which is the conditional probability distribution of the actual inventory given the available information, is constructed at each stage using expression (6.5), and by computing the expected value of this state, a dramatic improvement in the estimate of the state is achieved. This relatively simple computation makes possible the control of inventory problems with much longer horizon, much higher upper limit on the stock quantity, and finite lead time.

Consider, for example, the (Q,R) model used in the earlier chapters, but now the state is no longer exactly known and instead the only available information are the inaccurate Auto-ID measurements in the past and present period $(z_0, z_1, \ldots, z_k)$ and the past order quantity $(u_0, u_1, \ldots, u_{k-1})$. We can build a state estimation technique that is conceptually identical to what is presented in the previous chapter with minor modifications. Assuming the lead time is known and fixed at L and the receipt quantity is exactly equal to the order quantity, our information vector at period $k+1$ would now consist of all the past shipment quantities, $h_0, h_1, \ldots, h_k$ (which are equal to the quantities ordered L lead time periods ago), and the past and present measurements $z_0, z_1, \ldots, z_{k+1}$

$$I_{k+1} = (h_0, h_1, \ldots, h_k, z_0, z_1, \ldots, z_{k+1}). \tag{6.15}$$

This is simply the information vector used in the zero lead time model in the previous section with the order quantity u_k replaced by the shipment quantity h_k.

The conditional probability distribution of the actual inventory in period $k+1$ would then be computed using the expression

$$p_{k+1}^{(j)} = \frac{\sum_{i=0}^{n} p_k^{(i)} \pi_{i,j}(h_k) r_{z_{k+1},j}}{\sum_{s=0}^{n} \sum_{i=0}^{n} p_k^{(i)} \pi_{i,s}(h_k) r_{z_{k+1},s}} \tag{6.16}$$

which, similarly, is equivalent to Equation (6.5) with u_k replaced by h_k. The state transition probability $\pi_{i,j}(h_k)$ can be determined from the probability distribution of the demand w_k and the stock loss v_k and the system dynamics in (2.5)

$$x_{k+1} = x_k + h_k - a_k - \min(v_k, x_k + h_k - a_k).$$

The sequence of events is same as that of Section 2.2 with additional steps at the beginning of each period. In period $k+1$, events occur in the following steps:

1. The measurement of the actual inventory, z_{k+1}, is obtained from Auto-ID system.
2. The conditional probability distribution of the actual inventory, P_{k+1}, is computed using (6.16).
3. The estimate of the inventory is obtained by taking the expected value of P_{k+1}.
4. Quantity Q is ordered to the supplier if the estimate is less than the reorder point R.
5. The incoming shipment is received.
6. The demand for purchase and stock loss are met.

Therefore we have designed an inventory error compensation method that uses the probability distribution of the stock quantity under the presence of erroneous Auto-ID measurements, and its performance can be compared with that of other compensations presented in Chapter 4. This is accomplished by using the same numerical example: the demand for purchase during period k is normally distributed with mean $\mu_w = 10$, standard deviation $\sigma_w = 2$, and the stock loss demand has Poisson distribution with mean $\lambda = 0.1$ (1% of average demand). The lead time is $L = 3$ and $Q = 50$.

One more parameter needs to be specified — p_g. We test an extreme case of the measurement performance and set p_g to be 0.05. To illustrate the measurement performance of this system, the cumulative measurement probability when the actual inventory is 50 is shown in Figure 6-5. Under this setting, the Auto-ID measurement is off by more than 12 items with the probability of 0.5.

In Figure 6-6 are the sample simulation runs showing the effect of using the state estimation technique. The reorder point is set at 41 so that the target stockout rate of 0.5% is

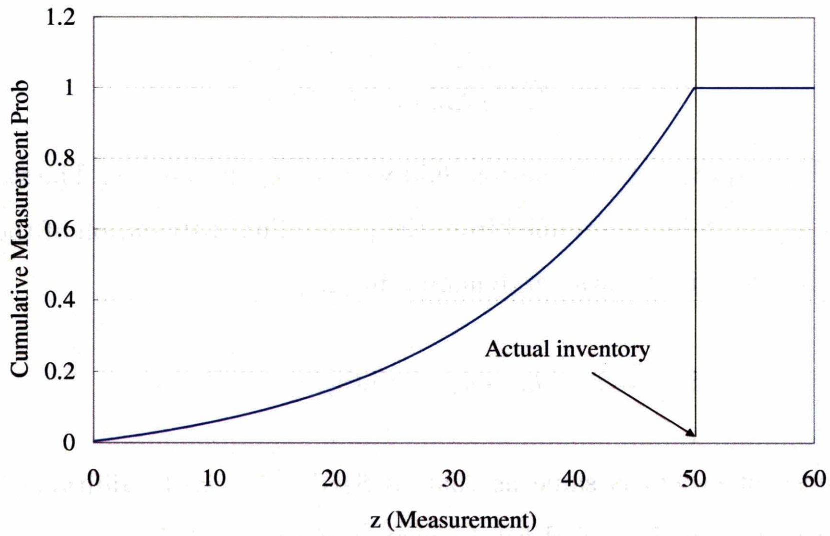

Figure 6-5: Cumulative Auto-ID measurement probability (Actual inventory is 50)

achieved in absence of stock loss and under perfect state information. When the Auto-ID system has the poor measurement performance of $p_g = 0.05$, the inventory record severely underestimates the actual inventory frequently (shown in Figure 6-6(a)), resulting in higher average inventory level than desired. When this poor measurement is 'filtered' by using the state estimation technique, the inventory record is able to track the actual inventory much more closely (Figure 6-6(b)). The overall improvement in the performance achieved by this state estimation technique is shown in Figure 6-7, which plots the average inventory against the average stockout rate for various reorder points. Each data point is the result of 500 repetitions of independent simulations. This is the same type of plot in Section 4.2 used to compare the various compensation methods. In fact, the performance of system when the Auto-ID is working perfectly ('Perfect State Info' curve) and when no correction is made to the inventory record ('No Compensate' curve) is reproduced here for comparison purposes, along with the result of decrementing the inventory by the average stock loss ('Decrement Record' curve). The Auto-ID system with $p_g = 0.05$ ('pg=0.05' curve) makes only limited improvement from the worst case scenario due to the poor measurement performance. In fact, its performance is noticeably worse than the simple technique of decrementing the inventory record each period. However, when this corrupt Auto-ID measurement is filtered using the probability distribution of the actual inventory (represented by the 'pg=0.05, Filtered' curve), dramatic improvement in the inventory-stockout rate compromise is seen, as

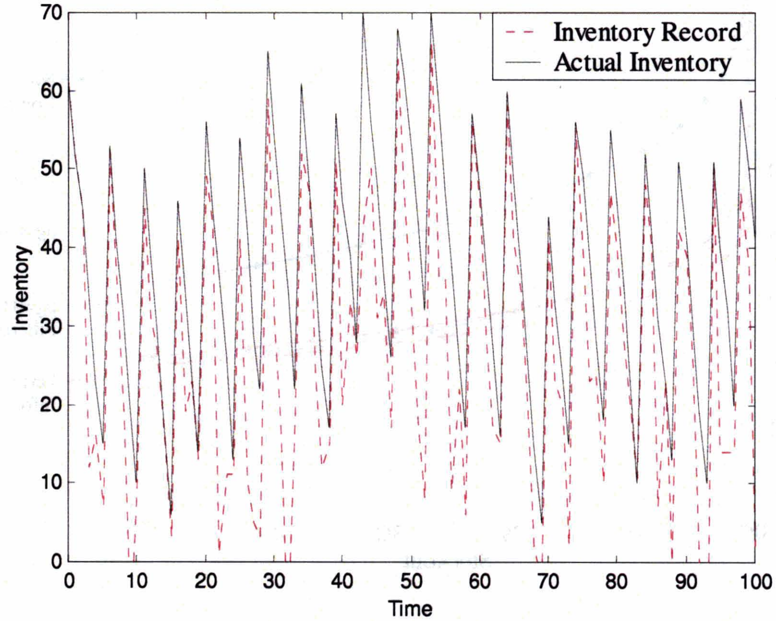

(a) Auto-ID with p_g=0.05

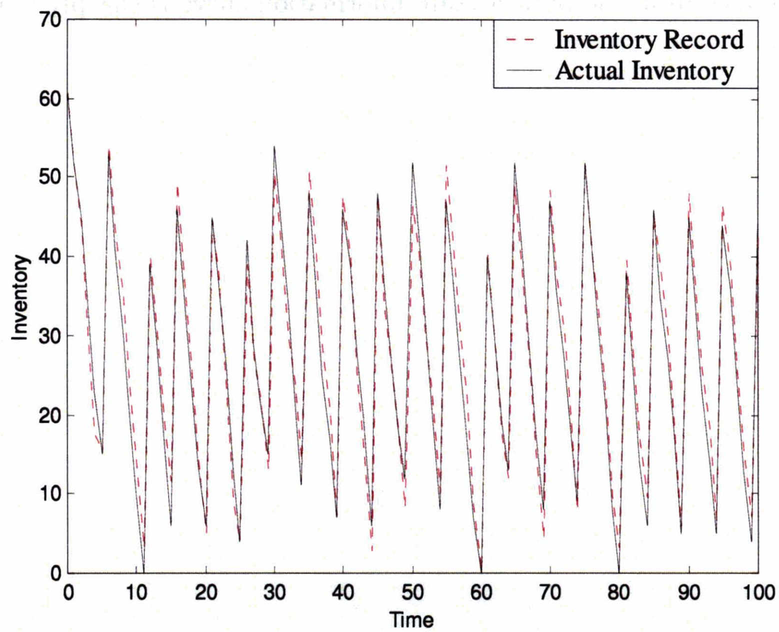

(b) Auto-ID with p_g=0.05 (Using state estimation technique)

Figure 6-6: Sample simulation runs showing the effect of state estimation technique

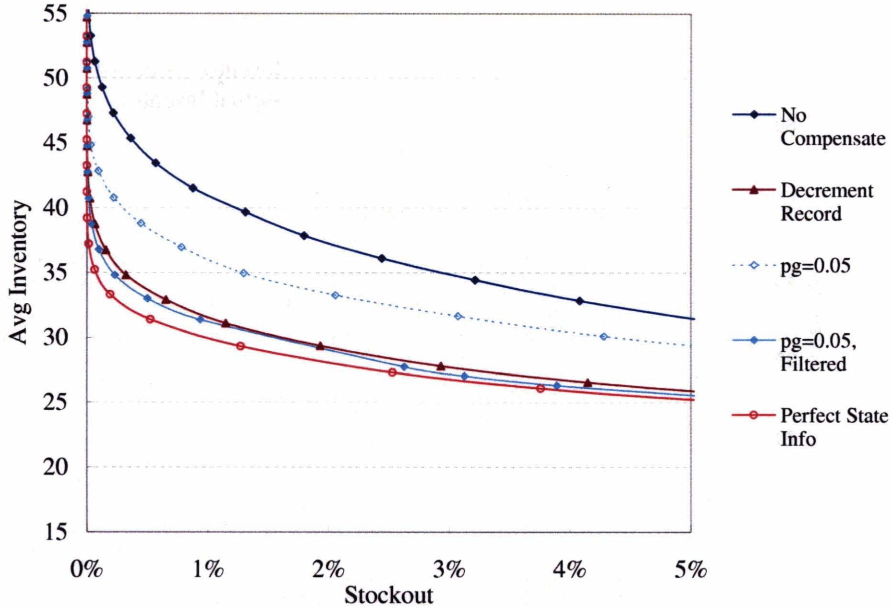

Figure 6-7: Overall improvement in performance obtained by the state estimation technique.

the vertical distance from the perfect state information curve is sharply reduced.

Chapter 7

Sub-Optimal Inventory Control Without Auto-ID

In the previous chapter, it was found that even when the Auto-ID system falls short of providing the perfectly accurate measurement, the inventory error can be effectively compensated through better estimation of the actual inventory. This was accomplished by computing the probability distribution of the actual inventory in each period, simply taking its expected value, and using it as if it is an accurate read of the actual inventory in making ordering decisions.

In this chapter, we apply this control scheme to today's typical inventory environment in which there is no Auto-ID, but the only available data is the incoming shipment quantity and the sales data (namely, the perpetual inventory system). It was shown in Chapter 2 that if nothing is done to adjust for the inventory inaccuracy, the recurring stock loss would create a growing gap between the inventory record and the actual inventory and lead to performance degradations. It is of interest to see how the state estimation technique based on the available data in today's inventory control environment can make the inventory record better track the actual inventory.

This control scheme basically is an addition to the various compensation methods presented in Chapter 4. In fact, we subject this control scheme to the same model and testing, and compare its performance under the same set of criteria.

7.1 The State Estimation Technique

As described in the introduction, many of today's automatic replenishment systems are designed to update their computerized inventory records whenever there there is incoming shipment from the suppliers or items are sold. The transaction for the former is done typically through the purchase order documents (which may be in hardcopy or electronic form) and the latter through barcode scan data at the checkout counters (also known as POS — Point of Sales). Under this 'perpetual inventory' system, the actual inventory in period $k+1$, x_{k+1}, and the inventory record, $\tilde{x}_{k+1}$, change according to (2.4) and (2.5), respectively:

$$\tilde{x}_{k+1} = \tilde{x}_k + h_k - a_k \tag{7.1}$$

$$x_{k+1} = x_k + h_k - a_k - \min(v_k, x_k + h_k - a_k) \tag{7.2}$$

where h_k is the receipt quantity, a_k is the sales quantity, and v_k is the demand for stock loss in period k.

The goal is to build a state estimation technique analogous to what is used for the Auto-ID system (Section 6.2) applicable for this perpetual inventory system. One clear difference lies in the available information. At the beginning of period k, the information that can aid the inventory manager's ordering decision is the collection of the past incoming shipment quantity (assuming this is equal to the order quantity) and the sales quantity. Summarizing this data is the information vector I_k which is expressed as

$$I_k = (h_0, h_1, \dots, h_{k-1}, a_0, a_1, \dots, a_{k-1}). \tag{7.3}$$

The next task is to come up with an expression for the conditional probability distribution of actual inventory given the information vector, P_k,

$$P_k = \begin{bmatrix} \text{Prob}(x_k = 0 | h_0, \dots, h_k, a_0, \dots, a_{k-1}) \\ \text{Prob}(x_k = 1 | h_0, \dots, h_k, a_0, \dots, a_{k-1}) \\ \vdots \\ \text{Prob}(x_k = n | h_0, \dots, h_k, a_0, \dots, a_{k-1}) \end{bmatrix} = \begin{bmatrix} p_k^{(0)} \\ p_k^{(1)} \\ \vdots \\ p_k^{(n)}. \end{bmatrix} \tag{7.4}$$

It turns out that much like for the case involving the Auto-ID measurement z_k, the above probability distribution can be generated recursively using the probability distribution at the previous period and the incoming shipment and sales observed in the previous period. In fact, the expression for $p_k^{(j)}, j = 0, 1, \ldots, n$, that constitutes the above matrix can also be obtained using the basic probability laws, and the final expression is

$$\begin{aligned} p_{k+1}^{(j)} &= \frac{\sum_{i=0}^{n} p_k^{(i)} \Pr(x_{k+1} = j | x_k = i, h_k, a_k) \Pr(a_k | x_k = i, h_k)}{\sum_{i=0}^{n} p_k^{(i)} \Pr(a_k | x_k = i, h_k)} \\ &= \Phi(P_k, h_k, a_k). \end{aligned} \tag{7.5}$$

See Appendix D.2 for the detailed mathematical steps of derivation.

According to the above equation, the construction of P_{k+1} involves P_k and two conditional probabilities $\Pr(x_{k+1} = j | x_k = i, h_k, a_k)$ and $\Pr(a_k | x_k = i, h_k)$. The former is the transition probability of the actual state x_k given the knowledge of h_k and a_k and is analogous to $\pi_{i,j}$ in Equation (6.5). Both of these conditional probabilities can be computed from the problem data, which would contain the distribution for the purchase demand w_k and the stock loss demand v_k, and the dynamics of the system described by Equation (7.2).

We can now use the (Q,R) simulation model used in Chapter 4 and Section 6.2 to see how well the state estimation technique using the currently available data in today's perpetual inventory systems compares with that of others. The sequence of events is as follows:

1. The probability distribution of the actual inventory, P_{k+1}, is computed using Equation (7.5).
2. The estimate of the inventory, $\tilde{x}_{k+1}$, is obtained by taking the expected value of P_{k+1}.
3. The order of quantity Q is placed to the supplier if $\tilde{x}_{k+1} \leq R$.
4. The incoming shipment is received.
5. The demand for purchase and stock loss are met.

7.2 Results and Discussion

We use the same numerical example of the previous sections: the demand during period k is normally distributed with the mean $\mu_w = 10$ and standard deviation $\sigma_w = 2$, and the stock loss demand has Poisson distribution with mean $\lambda = 0.1$ (1% of average demand). The lead time is $L = 3$ and $Q = 50$. Figure 7-1 shows sample simulation runs when nothing is done to correct the inventory record error (7-1a) and when the inventory record's estimate is enhanced using the control scheme outlined in the previous section (7-1b).

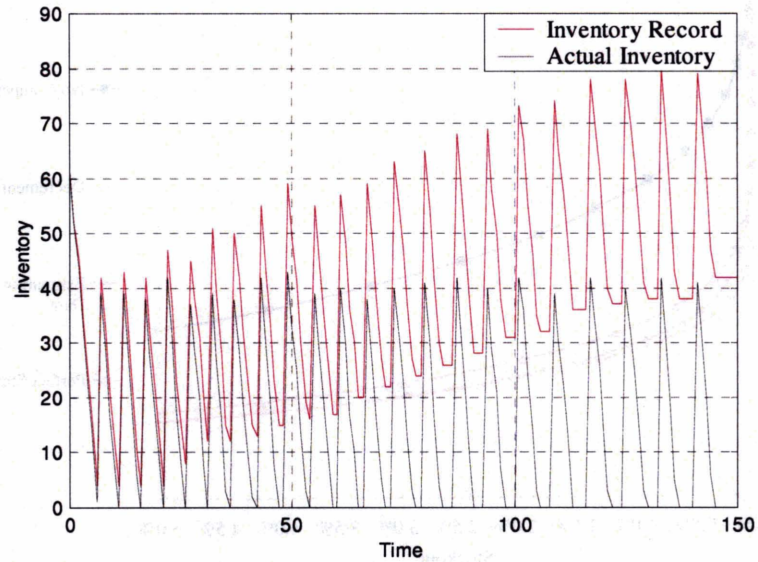

(a) No compensation

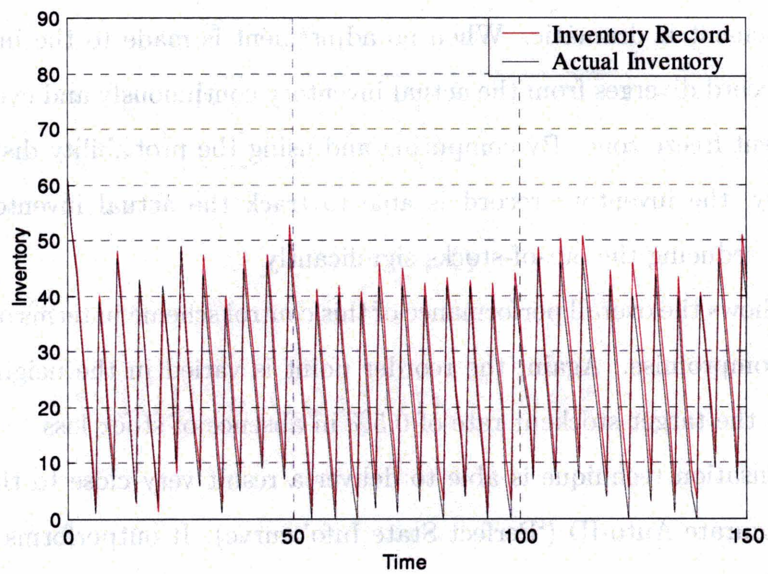

(b) Compensation using the state estimation technique

Figure 7-1: Sample simulation runs showing the effect of improving the estimate of inventory using currently available data

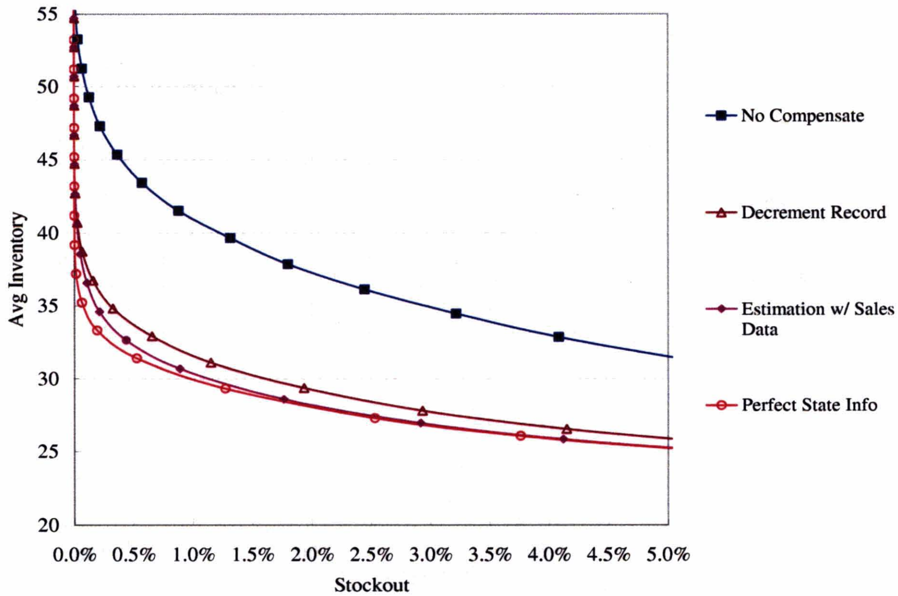

Figure 7-2: Overall improvement in performance made by the state estimation technique using currently available data

The improvement is dramatic. When no adjustment is made to the inventory record, the inventory record diverges from the actual inventory continuously and eventually reaches the replenishment freeze zone. By computing and using the probability distribution of the actual inventory, the inventory record is able to track the actual inventory much more closely, thereby reducing the out-of-stocks significantly.

Figure 7-2 shows the overall performance of this control scheme in terms of the inventory-stockout rate compromise. Again, the reorder point is varied in the neighborhood of 41, which produces the target stockout rate of 0.5% in absence of stock loss.

This compensation technique is able to deliver a result very close to that achieved by the perfectly accurate Auto-ID ('Perfect State Info' curve). It outperforms the strategy of decrementing the inventory record by the average stock loss ('Decrement Record' curve) for all the reorder points considered. Since this control scheme requires a far less costly implementation than Auto-ID, the investment required by Auto-ID systems may not be justified if the measurement accuracy is not high enough.

Nevertheless, the ability of this compensation method to perform well depends on one critical factor: the model using state estimation, as well as other compensation techniques discussed in previous chapters, needs accurately estimated system parameters. Consider,

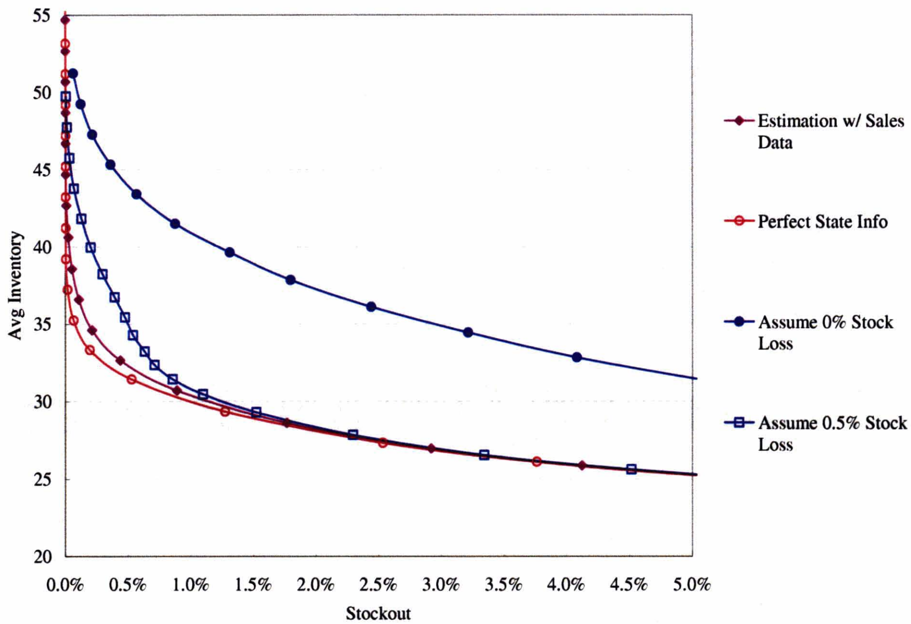

Figure 7-3: Sensitivity of the performance of state estimation technique on the estimate of the stock loss

for example, the assumptions regarding the stock loss. We have seen that the perpetual inventory system is highly sensitive to the rate of stock loss if no correction is made on the inventory record. This can be interpreted to saying that even a slightly inaccurate estimate of the stock loss is likely to result in serious performance decrease. Figure 7-3 shows how much the performance of this compensation technique suffers when the actual stock loss (λ in the Poisson distribution used for the stock loss) is 1% of the average demand for purchase, but it is incorrectly estimated at 0% ('Assume 0% Stock Loss' curve) and at 0.5% ('Assume 0.5% Stock Loss' curve). If the stock loss estimate is off by as small as 1%, a sharp degradation in the performance is observed.

The Auto-ID is exempt from this sensitivity on the stock loss estimate since its measurement is dependent only on the actual inventory at the time of the read and is thus unaffected by the stock loss.

Chapter 8

Conclusion and Further Research

8.1 Conclusion

Motivated by an emerging product identification and tracking technology under investigation at the Auto-ID Center at Massachusetts Institute of Technology, this research attempts to assess the value of accurate account of what products are where in what quantity throughout the various points in supply chain. However, in order to quantify this impact, it was necessary to first understand what the causes of the inventory errors are, the mechanism by which the inaccurate inventory record degrades the inventory system performance, and what the cost of that inaccuracy is. Understanding this baseline performance enables us to quantify the incremental benefit of the Auto-ID technology.

Among many different causes of the inventory error, the unknown stock loss, defined as disappearance of items (such as due to theft) not detected by the information system, is identified as a serious potential threat to inventory system performance and becomes the focus of this research. Analytical and simulation modelling of two commonly used inventory policy in practice — the (Q,R) policy and the fixed review period, base stock policy — demonstrate that the inventory system performance is highly sensitive to the inaccuracy caused by stock loss. That is, even a small level of stock loss accumulated over time can lead to inventory inaccuracy that disrupts the replenishment process and creates severe out-of-stocks. The consequences of untreated inventory error can be especially severe in the (Q,R) policy where the inventory record can stay above the reorder point and halt the replenishment process completely. This modelling work also provides a managerial insight: that revenue losses due to out-of-stocks can far outweigh the losses due to the disappearing items themselves, and thus calls for the inventory managers to pay a close attention to maintaining the inventory record accuracy to effectively control the stock loss problem.

Parametric analysis further reveals that the sensitivity of the inventory system performance to the inventory record error is heightened in lean systems characterized by short lead times and frequent ordering of small quantities. Stripped of the ability to deal with uncertainties in the operation, performance of lean system is likely to be hampered far more by inventory inaccuracy, and thus maintaining accurate inventory record is critical to reap the benefits lean systems have to offer.

The benefit of the accurate inventory information, such as provided by the Auto-ID technology, is indeed substantial. The simulation model (demonstrated in Chapter 4) shows that once furnished with accurate stock quantity every period, the inventory system has the resilience to counter the damaging effects of unknown stock loss and maintain a performance level (characterized by the inventory-stock loss compromise) very close to that of a system without any stock loss.

Comparison of the Auto-ID technology to other compensation techniques demonstrate that there are much simpler and less costly ways to effectively control the inventory error, such as decrementing the inventory record by the mean of the stock loss every period. However, due to the high sensitivity of the system performance to the stock loss, accurate estimate of the system parameters and the dynamics are required for successful execution. Furthermore, it is found that classic ways of buffering against operational uncertainties, such as carrying higher safety stock, is not a desirable solution to the inventory inaccuracy problem caused by the stock loss.

Next, we examine the performance of the Auto-ID system whose measurement of the actual stock quantity is not perfectly accurate. This is motivated by the fact that the technology is still under development in the industry and is yet to reach a full maturity in operational environment. We formulate the inventory management problem as an imperfect state information dynamic programming problem and we obtain the optimal inventory policy. This analysis reveals that even with a poor measurement performance, optimal profit close to the perfect measurement case can be achieved. However, this comes at a price. As the measurement performance gets poorer, the structure of the optimal policy becomes more complex. The optimal order quantity develops the tendency to depend not only on the current period measurement, but on the previous periods as well.

Another challenge in the implementation of the optimal inventory policy for Auto-ID systems with measurement error is the small scope of the problem due to the computational limitations. To address this problem, we examine a state estimation technique that uses the probability distribution of the actual stock, the quantity ordered, and the Auto-ID measurement, and discover that it is able to substantially simplify the computation of the order quantity and yet achieve a near-optimal performance.

Moreover, motivated by the superb performance of this state estimation technique, we also look for a similar sub-optimal control scheme which does not require the Auto-ID measurement, but utilizes the information currently available in today's typical inventory operations — incoming and outgoing shipment quantity. We again discover that this state estimation technique is able to achieve a performance very close to that using Auto-ID measurement. However, it also requires a very accurate knowledge of the system parameters due to the highly sensitive nature of the system on inventory inaccuracy caused by stock loss.

These observations lead us to the conclusion that the selection of an appropriate control technique to combat the inventory inaccuracy caused by unknown stock loss requires careful evaluation of the dynamics of the system and the system parameters, including the purchase demand, stock loss demand, and the behavior of the Auto-ID measurements. If these parameters can be estimated with confidence, then there exist compensation techniques that are much simpler and less costly than Auto-ID and still achieve near-optimal performance. However, in environments where the system behavior is governed by high uncertainty, which is true in most real-life situations, the accurate stock quantity information provided by Auto-ID (or, an estimate of the quantity based on an Auto-ID system with measurement errors) can be of a great benefit.

8.2 Further Research

This research attempts to contribute to the field of inventory management and control by challenging a fundamental assumption commonly used in many inventory models: perfectly accurate knowledge of inventory system information is available and accessible for decision making. In particular, we relax the assumption that the knowledge of the actual stock quantity, which is the single most important piece of information required in inventory policies, is provided accurately.

Such deviation from the classical inventory models opens the doors for many research areas other than the ones covers in this thesis. We summarize some of the potential key research problems that build or expand upon this thesis in the area of information inaccuracy in inventory systems:

Modelling other causes of inventory inaccuracy: In this research, unknown stock loss was identified as one of the primary causes of inventory inaccuracy and became the focus of the study. However, other types of causes of inventory inaccuracy exist as described in Chapter 1, and some may be more important than unknown stock loss depending on the nature of the product and operations.

Selecting other causes of inaccuracy is likely to result in different conclusions about the mechanism by which the error creates disruptions in the operations, sensitivity of the system performance on the inventory inaccuracy, performance of various compensation techniques, and etc.

Supply-chain-wide impact of inventory inaccuracy: No inventory-carrying facility is a stand-alone system in the supply chain. Therefore, disruptions caused by inventory inaccuracy at one point in supply chain are likely to propagate to the upstream and downstream trading partners. For example, when a retailer at the far end of the supply chain suffers from replenishment freeze due to prolonged period of untreated inventory inaccuracy, its supplier will be impacted by this disruption and experience unusual ordering patterns. Modelling a multi-stage inventory system and looking for a macroscopic understanding of the impact of inventory inaccuracy is a potential area of investigation. This investigation takes greater importance today as more and more companies tightly integrate their supply chain processes with their trading partners to maximize efficiency.

It is also of interest to examine the impact of information inaccuracy in transaction between trading partners. By nature, supply chain partners operate on frequent hand-off of materials and goods, and every handoff is susceptible to information inaccuracy. Investigation would include how such transaction errors can accumulate over time and through what mechanism they can create disruptions in ordering and supplying behavior in the partners.

Expansion in dynamic programming problem: The analytical works involving dynamic programming in Chapter 5 to Chapter 7 were intended to illustrate a way in which an inventory inaccuracy problem can be formulated as a dynamic programming model. For this reason, simplifications were used in the model. However, a more interesting and realistic problem can be formulated by expanding the scope of the model.

For an example, a finite lead time, rather than zero lead time as used in this research, can be incorporated into the model. The dynamic programming formulation would then use state augmentation that involves enlargement of the state space. Unfortunately, this reformulated problem may have very complex state and/or control spaces.

Another example is to add a fixed ordering cost. In the perfect state information problem, this leads to the optimal policy known as the (s, S) policy. This research area could include an investigation of the relevancy of this policy to the imperfect state information problem and identification of the structure of the optimal ordering policy.

Inventory Inaccuracy in MRP and DRP systems: Manufacturing Resource Planning (MRP) and Distribution Resource Planning (DRP) systems require inventory control policies very different from the ones discussed in this thesis. The research questions addressed in this thesis can be studied in the MRP and DRP system: what the causes of errors are, by what mechanism the error degrades the performance, how sensitive the performance is to the error, and what effective compensation techniques are available.

A number of research works appear in the literature that addresses inventory inaccuracy in MRP systems. In particular, Krajewski (1987) uses simulation to assess which factors in a MRP-based production environment (inventory inaccuracy being one of them) have the biggest impact on performance. One possible contribution beyond this work is rather than building a complex and large-scale model, starting with a simple model and looking for insights to explain the relationship between inventory inaccuracy and the performance of the MRP systems.

Appendix A

Calculations for Deterministic Model - (Q,R) Policy

A.1 Exact Calculations

Calculating performance measures of the inventory system of Section 2.3 requires computing the ending times t_A and t_B of Regions A and B, respectively, of Figure 2-4. The schematic is shown again in Figure A.1 for convenience.

t_1

The analysis of Region A and the calculation of times t_A and t_1 require a focus on the actual inventory, since they are defined by the actual inventory reaching zero.

t_A can be determined by finding the number of cycles in Region A, denoted n_A, and the length of each cycle. In the first cycle, the initial actual inventory and the initial inventory record are both $R+Q-wL$. When the inventory record reaches the reorder point R, it has decreased by $(R+Q-wL)-R=Q-wL$. The inventory record further decreases by wL until just before the first order arrives. The inventory record is then $R-wL$ so the total decrease in the first cycle is $(R+Q-wL)-(R-wL)=Q$. When the order of amount Q arrives, the inventory record jumps back up to $R+Q-wL$. This cycle in the inventory record repeats as long as the real inventory is above zero. The length of a cycle is the time required for the demand (at rate w) to consume the amount Q, or Q/w.

When the actual inventory reaches zero, sales are interrupted and a new kind of behavior

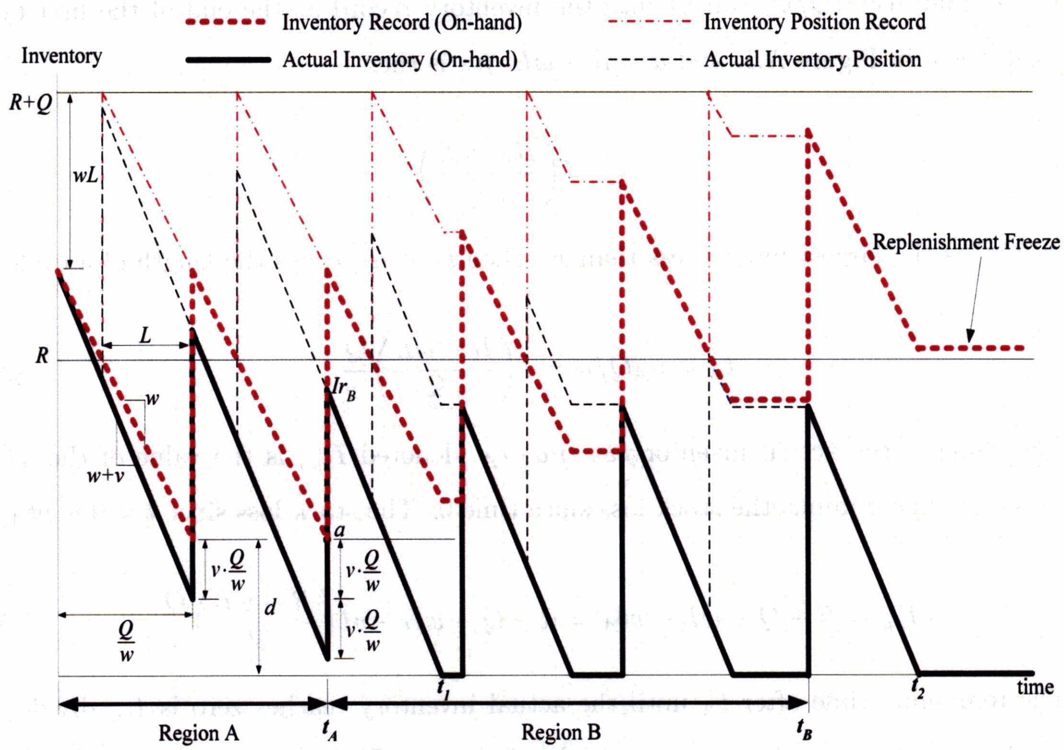

Figure A-1: (Q,R) Policy subjected to stock loss begins. The actual inventory, since it decreases at the rate v faster than the inventory record, drops by an additional amount, $v \cdot$ (length of a cycle) $= v(Q/w)$, in the first cycle. That is, the actual inventory decreases by $Q(1 + v/w)$ in the first cycle. The net change, after the order arrives, is then a decrease of $v(Q/w)$.

As long as the demand and stock loss rate are constant, the actual inventory decreases by this amount during each cycle. However, when the actual inventory reaches zero, the slope of the two inventory curves change — in fact, they go to zero. The number of cycles until this happens, i.e., the number of cycles in Region A, is the largest integer number of times the quantity $v(Q/w)$ can fit into the inventory record at the end of the first cycle. This value is d in Figure A.1, and $d = R - wL$. n_A is then

$$n_A = I\left(\frac{R - wL}{v\frac{Q}{w}}\right) \tag{A.1}$$

where $I(x)$ is the largest integer less than x. Then t_A is n_A times the length of a cycle, or

$$t_A = n_A Q/w = I\left(\frac{R - wL}{v\frac{Q}{w}}\right)\frac{Q}{w}. \tag{A.2}$$

The value of the actual inventory at time t_A, denoted I_{r_B}, is the value of the actual inventory at time 0 minus the stock loss since time 0. The stock loss since $t = 0$ is vt_A, so

$$I_{r_B} = R + Q - wL - vt_A = R + Q - wL - vI\left(\frac{R - wL}{v\frac{Q}{w}}\right)\frac{Q}{w}. \tag{A.3}$$

The remaining time after t_A until the actual inventory reaches zero is I_{r_B} divided by $v + w$, the rate that actual inventory actually decreases. The time of the first out-of-stock, t_1, is t_A plus the time after t_A that the actual inventory reaches zero, or

$$\begin{aligned} t_1 =& t_A + \frac{I_{r_B}}{v + w} \\ =& I\left(\frac{R - wL}{v\frac{Q}{w}}\right)\frac{Q}{w} + \frac{I_{r_B}}{v + w}. \end{aligned} \tag{A.4}$$

t_2

The analysis of Region B and the calculation of times t_B and t_2 require a shift of focus to the inventory record, since they are defined by the inventory record exceeding R.

In this section, we determine n_B, the number of cycles in Region B, and the length of each cycle. To calculate n_B, we make use of the fact that in Region B, the ending value of an inventory record cycle is higher than that of the previous cycle.

First cycle in Region B The inventory record is $R + Q - wL$ at time t_A, and it drops at rate w as long as the actual inventory remains strictly positive. As we have shown, the actual inventory remains positive for a time period of length $I_{r_B}/(v + w)$, so the inventory record drops by $wI_{r_B}/(v + w)$ between t_A and t_1. Note that this decrease is *less* than the order amount Q since the time between t_A and t_1 is less than a full Region A cycle.

Since the actual inventory is zero immediately after t_1, sales are zero and the inventory record remains constant until the order arrives. The inventory record, between t_1 and when the order arrives, stays at $R + Q - wL - wI_{r_B}/(v + w)$. Since the value of the inventory record just before the order arrives in Region A is $R - wL$, the increase in the inventory record just before the order arrives at the end of the first cycle in Region B is $(R + Q - wL - wI_{r_B}/(v + w)) - (R - wL) = Q - wI_{r_B}/(v + w)$. When the order arrives, both the inventory record and the actual inventory jump by Q.

Later cycles In later cycles, the actual inventory always starts at exactly Q. The time required for the actual inventory to reach zero is $Q/(v + w)$, and it remains at zero until the next order arrives. The inventory record decreases by $wQ/(v + w)$ while the actual inventory is positive, and then, after a period of no sales, it jumps by Q. Therefore, the net change in the inventory record during a Region B cycle other than the first is $Q - wQ/(v + w) = vQ/(v + w)$.

To summarize, the value of the inventory record at t_1, the end of the first cycle of Region B, is

$$R + Q - wL - \frac{wI_{r_B}}{v + w} \tag{A.5}$$

and the inventory record at the end of the i^{th} cycle is

$$R + Q - wL - \frac{wI_{r_B}}{v + w} + (i - 1)\left(\frac{vQ}{v + w}\right). \tag{A.6}$$

Number of cycles Let n_B be the number of Region B cycles until the inventory record is greater than or equal to R at the lowest point of a cycle. For there to be at least one cycle, we must have

$$R + Q - wL - \frac{wI_{r_B}}{v+w} \geq R$$

or

$$Q - \frac{wI_{r_B}}{v+w} \geq wL.$$

Then n_B is given by

$$n_B = \begin{cases} 0 & \text{if } Q - \dfrac{wI_{r_B}}{v+w} < wL \\ 1 + \min n & \text{otherwise} \end{cases}$$

where n is an integer such that

$$R - wL + \left(Q - \frac{wI_{r_B}}{v+w}\right) + n\left(Q - \frac{vQ}{v+w}\right) \geq R$$

or

$$n \geq \frac{wL - \left(Q - \frac{wI_{r_B}}{v+w}\right)}{\left(\dfrac{vQ}{v+w}\right)}.$$

Therefore, n_B is given by

$$n_B = \begin{cases} 0 & \text{if } Q - \dfrac{wI_{r_B}}{v+w} < wL \\ 1 + I\left(\dfrac{wL - \left(Q - \dfrac{wI_{r_B}}{v+w}\right)}{\left(\dfrac{vQ}{v+w}\right)}\right) & \text{otherwise.} \end{cases}$$

Length of Region B Unlike the Region A cycles, the duration of the Region B cycles are not all the same. To determine the total length of Region B, we need to compute the length of its cycles. The length of each cycle is the time it takes for the inventory record to reach the reorder point from the start of the cycle plus the lead time L.

That is, if the inventory record at the start of a cycle is x (which includes the order of size Q that just arrived), the time until the inventory record reaches R is $(x - R)/w$ and the length of the cycle is

$$(x - R)/w + L. \tag{A.7}$$

At the start of the first cycle, the inventory level is the same as in every cycle in Region A: $R + Q - wL$. Therefore, the length of the first cycle is Q/w.

To determine the length of the second cycle, we recall from (A.5) that the inventory record at time t_1 is $R+Q-wL-wI_{r_B}/(v+w)$. Therefore, the inventory record, at the start of the second cycle, just after the order arrives, is $R+Q-wL-wI_{r_B}/(v+w)+Q$. The length of the second cycle is then $(Q-wL-wI_{r_B}/(v+w)+Q)/w+L$ or $Q/w+Q/w-I_{r_B}/(v+w)$.

More generally, (A.6) implies that the inventory record at the start of the i^{th} cycle, for $i \geq 2$, is

$$R + 2Q - wL - \frac{wI_{r_B}}{v+w} + (i-2)\left(\frac{vQ}{v+w}\right)$$

so the length of the i^{th} cycle is, according to (A.7),

$$\frac{1}{w}\left(2Q - \frac{wI_{r_B}}{v+w} + (i-2)\left(\frac{vQ}{v+w}\right)\right). \tag{A.8}$$

The total length of Region B, for $n_B \geq 2$, is therefore

$$\frac{Q}{w} + \sum_{i=2}^{n_B} \frac{1}{w}\left(2Q - \frac{wI_{r_B}}{v+w} + (i-2)\left(\frac{vQ}{v+w}\right)\right).$$

Carrying out the summation above and simplifying, the total length of Region B can be summarized as

$$\begin{cases} 0 & \text{if } n_B = 0, \\ \frac{Q}{w} & \text{if } n_B = 1, \\ (n_B - 1)\left(\frac{2Q}{w} - \frac{I_{r_B}}{v+w}\right) + (n_B-1)(n_B-2)\frac{vQ}{2w(v+w)} & \text{if } n_B \geq 2. \end{cases} \tag{A.9}$$

t_2 is now the sum of t_A, the total length of Region B, and the last in-stock duration that exists immediately after Region B, which is $\frac{I_{rB}}{v+w}$ if $n_B = 0$ and $\frac{Q}{v+w}$ otherwise. This is

$$t_2 = \begin{cases} t_A + \frac{I_{rB}}{v+w} & \text{if } n_B = 0, \\ t_A + \frac{Q}{w} + \frac{Q}{v+w} & \text{if } n_B = 1, \\ t_A + \frac{Q}{w} + (n_B - 1)\left(\frac{2Q}{w} - \frac{I_{rB}}{v+w}\right) + (n_B - 1)(n_B - 2)\frac{vQ}{2w(v+w)} + \frac{Q}{v+w} & \text{if } n_B \geq 2. \end{cases} \tag{A.10}$$

S_{out}

We compute the stockout rate S_{out} to be the fraction of the entire operation time, t_f, occupied by the flat portions of the actual inventory curve in Region B. The length of the flat line in each Region B cycle is found by subtracting from the length of each cycle (which has already been determined in the previous section) the in-stock duration of each cycle, which is $\frac{I_{rB}}{v+w}$ for the first cycle and $\frac{Q}{v+w}$ thereafter. Using (A.8), the length of the flat line in the i^{th} cycle is

$$\frac{1}{w}\left(2Q - \frac{wI_{rB}}{v+w} + (i-2)\left(\frac{vQ}{v+w}\right)\right) - \frac{Q}{v+w} = \left(\frac{Q}{w} - \frac{I_{rB}}{v+w}\right) + (i-1)\frac{Qv}{w(v+w)}.$$

If $t_f \geq t_2$, then the stockout rate is the sum of the lengths of the flat lines in the above expression with $i = n_B$ and the amount by which t_f exceeds t_2, which is

$$\begin{aligned} S_{out} =& \frac{1}{t_f}\left[\sum_{i=1}^{n_B}\left(\frac{Q}{w} - \frac{I_{rB}}{v+w} + (i-1)\frac{Qv}{w(v+w)}\right) + t_f - t_2\right] \\ =& \frac{1}{t_f}\left[n_B\left(\frac{Q}{w} - \frac{I_{rB}}{v+w}\right) + (n_B)(n_B-1)\frac{Qv}{2w(v+w)} + t_f - t_2\right]. \end{aligned} \tag{A.11}$$

If $t_1 \leq t_f < t_2$, then stockout rate takes the form similar to the previous expression, except n_B is replaced by the number of complete Region B cycles that exist prior to the finishing time t_f, denoted by m. Also, the last two terms $t_f - t_2$ are replaced by length of the remaining flat line that may exist between the completion of m cycles and t_f. If m is zero — meaning t_f is located between t_A and the end of the first cycle in Region B — then this remaining flat line is the greater of zero or the quantity that remains when t_A and the in-stock duration $\frac{I_{rB}}{v+w}$ are taken away from t_f. Otherwise, it is the greater of zero or the quantity that remains when the total length of m cycles and the in-stock duration $\frac{Q}{v+w}$ is taken away from t_f. Using (A.9) and (A.11), we write this expression as

$$S_{out} = \begin{cases} \frac{1}{t_f}\max\left[0, t_f - t_A - \frac{I_{r_B}}{v+w}\right] & \text{if } m = 0, \\ \frac{1}{t_f}\left[\left(\frac{Q}{w} - \frac{I_{r_B}}{v+w}\right) + \max\left[0, t_f - t_A - \frac{Q}{w} - \frac{Q}{v+w}\right]\right] & \text{if } m = 1, \\ \frac{1}{t_f}\left[m\left(\frac{Q}{w} - \frac{I_{r_B}}{v+w}\right) + (m)(m-1)\frac{Qv}{2w(v+w)}\right. & \\ \left. + \max\left[0, t_f - t_A - (m-1)\left(\frac{2Q}{w} - \frac{I_{r_B}}{v+w}\right) - (m-1)(m-2)\frac{vQ}{2w(v+w)} - \frac{Q}{v+w}\right]\right] & \text{if } m \geq 2 \end{cases} \tag{A.12}$$

What remains is the expression for m. First, we look for the number of intervals (a real number), denoted by m', that lie between t_A and t_f. Using (A.9), we can write the quadratic equation

$$(m'-1)\left(\frac{2Q}{w} - \frac{I_{r_B}}{v+w}\right) + (m'-1)(m'-2)\frac{vQ}{2w(v+w)} = t_f - t_A. \tag{A.13}$$

Solving this quadratic equation for m', and taking its integer portion (since we are looking for integer number of complete cycles), we obtain for m

$$m = I\left(m'\right) = I\left(\frac{-b + \sqrt{b^2 - 4ac}}{2a}\right) \tag{A.14}$$

where

$$\begin{aligned} a &= \frac{Qv}{w(v+w)} \\ b &= 4\frac{Q}{w} - 2\frac{I_{r_B}}{v+w} - 3\frac{Qv}{w(v+w)} \\ c &= 2\left(\frac{Qv}{w(v+w)} - \frac{Q}{w} - \frac{I_{r_B}}{v+w} - (t_f - t_A)\right). \end{aligned}$$

We complete the calculation of stockout rate by noting that when $t_f < t_A$, $S_{out} = 0$.

A.1.1 Approximate Calculations

The exact calculations shown in the previous section can be simplified significantly by making a set of appropriate approximations. The first approximation begins with the number of cycles in Region A, n_A. Whereas in the exact calculation n_A had to be an integer, we now relax this constraint and use the approximation

$$I\left(\frac{R - wL}{v\frac{Q}{w}}\right) \approx \frac{R - wL}{v\frac{Q}{w}} \tag{A.15}$$

in Equation (A.1). This approximation works well when the stock loss rate v is small and the cycle length $\frac{Q}{w}$ is much smaller than t_f. If this is true, the argument in $I(\cdot)$ in the left side of equation will be large and taking only the integer portion of the argument will be close to the argument itself. n_A and t_A then becomes

$$n_A \approx \frac{R-wL}{v\frac{Q}{w}} \tag{A.16}$$

$$t_A \approx \frac{R-wL}{v}. \tag{A.17}$$

In determining t_1, we make an additional assumption that the beginning actual inventory in the first cycle of Region B, I_{r_B}, is equal to Q. Again, this works well for small v because the amount by which the actual inventory decreases more than the inventory record in each cycle will be small. With this approximation, the expression for t_1 becomes

$$t_1 \approx \frac{R - wL}{v} + \frac{Q}{v + w}. \tag{A.18}$$

Applying this approximation also to n_B, and once again relaxing the constraint that n_B has to be an integer, we obtain

$$n_B \approx \frac{wL(v + w)}{vQ}. \tag{A.19}$$

t_2 now simplifies to

$$\begin{aligned} t_2 \approx& t_A + \frac{Q}{w} + (n_B - 1)\left(\frac{2Q}{w} - \frac{I_{r_B}}{v + w}\right) + (n_B - 1)(n_B - 2)\frac{vQ}{2w(v + w)} + \frac{Q}{v + w} \\ =& t_A + L\left(\frac{1}{2} + \frac{w}{v}\right) + \frac{wL^2(v + w)}{2vQ} + \frac{Q}{v + w}. \end{aligned} \tag{A.20}$$

We proceed further by carrying out these approximations to calculation of stockout rate and arrive at

$$S_{out} \approx \begin{cases} 0 & \text{if } t_f < t_1, \\ \frac{1}{t_f}\left[m(m+1)\frac{vQ}{2w(v+w)}\right] & \text{if } t_1 \leq t_f < t_2, \\ \frac{1}{t_f}\left[\frac{L}{2}\left(\frac{wL(v+w)}{vQ}+1\right)+t_f-t_2\right] & \text{if } t_f \geq t_2. \end{cases} \tag{A.21}$$

where m also changes by the approximations to

$$m \approx \frac{2w+v}{2v} + \sqrt{\left(\frac{2w+v}{2v}\right)^2 + 2\frac{w(v+w)}{Qv}(t_f - t_1)}. \tag{A.22}$$

Note that in approximating S_{out} for the case when $t_1 \leq t_f < t_2$, we also assumed any flat line that may exist beyond m cycles is negligible.

Appendix B

Calculations for Deterministic Model — Base Stock Policy

Define a cycle as the evolution of inventory between arrival of two consecutive orders. A cycle therefore would consist of a single sawtooth curve in Figure 3-4. Similar to what is done in the deterministic (Q,R) model, the inventory evolution is divided into regions of interest. Let Region A consist of cycles for which no out-of-stock occurs. Region B consists of the cycles that follow the end of Region A and for which the out-of-stock duration is shorter than the lead-time L. Region C consists of the rest of the cycles. Let t_A and t_B denote the end of Region A and B, respectively.

t_1

t_1 can be determined by finding the number of cycles in Region A, denoted n_A, the length of each cycle, and the in-stock duration of the first cycle in Region B.

Initially, both the recorded and actual inventory is at $B - wL$, and by the first review time, which takes place at time $T - L$, the inventory record falls by $w(T - L)$. By the time the order arrives, it further falls by wL. The ending inventory is thus $B - wL - w(T - L) - wL = B - w(T + L)$. Since the actual inventory falls at the faster rate of $w + v$ and the time length of the first cycle is T, the actual inventory will be lower than the recorded inventory by vT at the end of the first cycle. The quantity ordered at the first review is the difference between B and the inventory record at time $T - L$, which is computed to be $B - (B - wl - w(T - L)) = wT$. Since there is no out-of-stock in Region A, the rest of the cycles in this region will have the identical order quantity of wT and cycle length of T. The number of cycles in Region A is then the integer number of times the quantity vT can fit in the vertical distance between zero and the ending recorded inventory of the first cycle, and is expressed as

$$n_A = I\left(\frac{B - w(T + L)}{vT}\right) \tag{B.1}$$

where $I(x)$ is the largest integer less than x.

For small stock loss rate v, however, the above can be approximated by

$$n_A \approx \frac{B - w(T + L)}{vT}. \tag{B.2}$$

The ending time of Region A, t_A, is then approximated by multiplying n_A by the cycle length T, expressed as

$$t_A \approx \frac{B - w(T + L)}{v}. \tag{B.3}$$

The in-stock duration of the first cycle in Region B is found first by determining the actual inventory at the beginning of Region B, denoted I_{r_B}. Knowing the initial inventory of this cycle(which is $B - wL$), the amount by which the beginning value of actual inventory drops with each cycle in Region A (which is vT), and the number of cycles in Region A (which is n_A), it becomes

$$\begin{aligned} I_{r_B} &= B - wL - vT \cdot I\left(\frac{B - w(T + L)}{vT}\right) \\ &\approx wT. \end{aligned} \tag{B.4}$$

The time of first out-of-stock, t_1 is then t_A plus the time required for the quantity wT to be consumed at the rate $w + v$, and is expressed as

$$t_1 \approx t_A + \frac{wT}{w + v}. \tag{B.5}$$

S_{out}

To obtain the expression for stockout rate S_{out}, we need to calculate the out-of-stock duration for each cycle beyond Region A. Calculating the out-of-stock duration requires the beginning value of actual inventory of each cycle, which depends on the out-of-stock duration of its previous cycle, and so forth.

Table B.1: Cycle progression in Region B

Cycle in Region B	Beginning Inventory	Out-of-stock Duration	Order Quantity
1st	wT	$T - T\frac{w}{w+v}$	wT
2nd	wT	$T - T\frac{w}{w+v}$	$wT\frac{w}{w+v}$
3rd	$wT\frac{w}{w+v}$	$T - T\left(\frac{w}{w+v}\right)^2$	$wT\frac{w}{w+v}$
4th	$wT\frac{w}{w+v}$	$T - T\left(\frac{w}{w+v}\right)^2$	$wT\left(\frac{w}{w+v}\right)^2$
5th	$wT\left(\frac{w}{w+v}\right)^2$	$T - T\left(\frac{w}{w+v}\right)^3$	$wT\left(\frac{w}{w+v}\right)^2$
6th	$wT\left(\frac{w}{w+v}\right)^2$	$T - T\left(\frac{w}{w+v}\right)^3$	$wT\left(\frac{w}{w+v}\right)^3$
⋮	⋮	⋮	⋮

Region A In the first cycle of Region B, the actual inventory begins at wT and remains in stock for $\frac{wT}{w+v}$. Since we know the time length of each cycle is fixed at T, the out-of-stock duration for this first cycle is $T - \frac{wT}{w+v}$. We also know that inventory is reviewed and order is placed to the supplier $(T - L)$ after the beginning of this cycle, and the size of the order is wT since there has been no out-of-stock in the last cycle of Region A.

In the second cycle, the actual inventory begins with wT (what is ordered in the first cycle), and thus the out-of-stock duration for this cycle is again $T - \frac{wT}{w+v}$. The size of order placed in this cycle is the quantity sold since the last inventory review time, and would be simply wT if there was no out-of-stock in the previous cycle. Since however there was out-of-stock for the duration $T - \frac{wT}{w+v}$, the size of order is $wT - w(T - \frac{wT}{w+v}) = wT\frac{w}{w+v}$.

Calculations for the rest of the cycles in Region B can be carried out in the same manner. Table B.1 summarizes the beginning actual inventory, out-of-stock duration, and quantity ordered in each cycle of Region B. It turns out that the out-of-stock duration changes in cycle pairs.

The total out-of-stock duration in Region B is the sum of the individual durations in the above table, and can be expressed as

$$\sum_{i=1,2,\dots}^{k} 2\left(T - T\left(\frac{w}{w+v}\right)^i\right) \tag{B.6}$$

The out-of-stock duration increases with each cycle, and by the definition of the Region B cycles, the last cycle would be the one having the longest out-of-stock duration equal to or less than the lead time L. We thus look for the quantity k in the above summation that satisfies

$$T - T\left(\frac{w}{w+v}\right)^k = L. \tag{B.7}$$

Table B.2: Cycle progression in Region C

Cycle in Region B	Beginning Inventory	Out-of-stock Duration	Order Quantity
1^{st}	$w(T-L)$	$T-(T-L)\left(\frac{w}{w+v}\right)$	$w(T-L)\left(\frac{w}{w+v}\right)$
2^{nd}	$w(T-L)\left(\frac{w}{w+v}\right)$	$T-(T-L)\left(\frac{w}{w+v}\right)^2$	$w(T-L)\left(\frac{w}{w+v}\right)^2$
3^{rd}	$w(T-L)\left(\frac{w}{w+v}\right)^2$	$T-(T-L)\left(\frac{w}{w+v}\right)^3$	$w(T-L)\left(\frac{w}{w+v}\right)^3$
i^{th}	$w(T-L)\left(\frac{w}{w+v}\right)^{i-1}$	$T-(T-L)\left(\frac{w}{w+v}\right)^i$	$w(T-L)\left(\frac{w}{w+v}\right)^i$

Solving for k, we obtain

$$k = I\left(\frac{\log\left(1-\frac{L}{T}\right)}{\log\left(\frac{w}{w+v}\right)}\right). \tag{B.8}$$

Since the out-of-stock duration increases in pairs (Table B.1), the number of cycles in Region B, n_B, is

$$n_B \approx 2k \tag{B.9}$$

and the ending time of Region B is

$$\begin{aligned} t_B \approx & t_A + n_B T \\ = & t_A + 2kT \end{aligned} \tag{B.10}$$

The total out-of-stock duration in Region B now becomes

$$\begin{aligned} &\sum_{i=1}^{k} 2\left(T - T\left(\frac{w}{w+v}\right)^i\right) \\ =&2\left[kT - T\frac{\frac{w}{w+v} - \left(\frac{w}{w+v}\right)^{k+1}}{1-\left(\frac{w}{w+v}\right)}\right] \\ =&2kT - 2T\frac{w}{v} + 2T\left(\frac{w+v}{v}\right)\left(\frac{w}{w+v}\right)^{k+1} \end{aligned} \tag{B.11}$$

Region B The expressions for Region C are obtained through the same analysis. The starting actual inventory of the first cycle in Region C is the size of order placed in the last cycle in Region B, which is equal to $w(T-L)$ because the last two cycles in Region B are approximated to have out-of-stock duration equal to exactly L. Knowing the beginning inventory, we obtain the progression of the order quantity and out-of-stock duration of the cycles in Region C as shown in Table B.2.

If the end-of-operation time t_f is greater than t_B, then the total out-of-stock duration will have contribution from Region B (which is expressed in Equation (B.11)) and from

Region C cycles. If the total number of cycles is approximated to be $\frac{t_f}{T}$, then the number of cycles in Region C, n_C, is approximated by subtracting from this quantity the number of cycles in Region A and B. Using expressions (B.2) and (B.9), we arrive at

$$n_C \approx I\left(\frac{t_f}{T} - \frac{B - w(T+L)}{vT} - 2k\right). \tag{B.12}$$

The contribution from Region C to the total out-of-stock duration is then

$$\begin{aligned} &\sum_{i=1}^{n_C}\left(T - (T-L)\left(\frac{w}{w+v}\right)^i\right) \\ =&n_C T - (T-L)\left(\frac{w}{v} - \frac{w+v}{v}\left(\frac{w}{w+v}\right)^{n_C+1}\right) \end{aligned} \tag{B.13}$$

If $t_A < t_f \leq t_B$, then the total out-of-stock duration takes the same form as Equation (B.11), with the quantity k, which is half the number of cycles in Region B, replaced by the new number of cycles defined by t_f, which is $\frac{t_f - t_A}{2T}$. The total out-of-stock duration is then

$$(t_f - t_A) - 2T\frac{w}{v} + 2T\left(\frac{w+v}{v}\right)\left(\frac{w}{w+v}\right)^{\frac{t_f-t_A}{2T}+1} \tag{B.14}$$

Knowing the there is no out-of-stock in Region A, and dividing Equation (B.14), (B.11) and (B.14) by t_f, we arrive at the final formulation for the stockout rate S_{out}

$$S_{out} \approx \begin{cases} 0 & \text{if } t_f \leq t_A \\ \frac{1}{t_f}\left[(t_f - t_A) - 2T\frac{w}{v} + 2T\left(\frac{w+v}{v}\right)\left(\frac{w}{w+v}\right)^{\frac{t_f-t_A}{2T}+1}\right] & \text{if } t_A < t_f \leq t_B \\ \frac{1}{t_f}\left[2kT - 2T\frac{w}{v} + 2T\left(\frac{w+v}{v}\right)\left(\frac{w}{w+v}\right)^{k+1} + n_C T\right. & \\ \left. -(T-L)\left(\frac{w}{v} - \frac{w+v}{v}\left(\frac{w}{w+v}\right)^{n_C+1}\right)\right] & \text{if } t_f > t_B \end{cases} \tag{B.15}$$

Appendix C

Proof of Convexity

Given the discrete time and continuous state system described by Equation (5.1)

$$x_{k+1} = \max(0, x_k + u_k - w_k - v_k), \qquad k = 0, 1, \ldots, N-1 \tag{C.1}$$

and the dynamic programming algorithm

$$\begin{aligned} J_N(x_N) &= 0, \\ J_k(x_k) &= \min_{u_k} \mathop{E}_{w_k, v_k} \big\{ cu_k + h\max(0, x_k + u_k - w_k - v_k) - p\min(w_k, x_k + u_k) \\ &\qquad + J_{k+1}(\max(0, x_k + u_k - w_k - v_k)) \big\}, \quad k = 0, 1, \ldots, N-1 \end{aligned} \tag{C.2}$$

we intend to prove that the function $J_k(x_k)$ is convex in x_k. Assume the demand and stock loss are bounded by $w_k, v_k \in [0, b]$. The state and control are constrained by $x_k, u_k \geq 0$.

First, we re-write the term representing the revenue, $p\min(w_k, x_k + u_k)$ by

$$\begin{aligned} p\min(w_k, x_k + u_k) &= p\big[-\max(-w_k, -(x_k + u_k))\big] \\ &= p\big[-\max(0, w_k - x_k - u_k) + w_k\big] \\ &= pw_k - p\max(0, w_k - x_k - u_k). \end{aligned} \tag{C.3}$$

Now, $J_k(x_k)$ becomes

$$\begin{aligned} J_k(x_k) &= \min_{u_k} \mathop{E}_{w_k,v_k} \{cu_k + h\max(0, x_k + u_k - w_k - v_k) + p\max(0, w_k - x_k - u_k) - pw_k \\ &\quad + J_{k+1}(\max(0, x_k + u_k - w_k - v_k))\} \\ &= \min_{u_k} \Big[cu_k + h \mathop{E}_{w_k,v_k} \{\max(0, x_k + u_k - w_k - v_k)\} + p\mathop{E}_{w_k} \{\max(0, w_k - x_k - u_k)\} \\ &\quad - p\mathop{E}_{w_k}\{w_k\} \quad + \mathop{E}_{w_k,v_k} \{J_{k+1}(\max(0, x_k + u_k - w_k - v_k))\} \Big]. \end{aligned} \tag{C.4}$$

By introducing the variables $y = x_k + u_k$ and $H(y)$

$$H(y) = h \mathop{E}_{w_k,v_k} \{\max(0, y - w_k - v_k)\} + p\mathop{E}_{w_k}\{\max(0, w_k - y)\}, \tag{C.5}$$

we can write $J_k(x_k)$ as

$$J_k(x_k) = \min_{y \geq x_k} \Big[cy + H(y) + \mathop{E}_{w_k,v_k} \{J_{k+1}(\max(0, y - w_k - v_k))\} \Big] - cx_k - p\mathop{E}_{w_k}\{w_k\}. \tag{C.6}$$

Let the function in the bracket be called

$$G_k(y) = cy + H(y) + \mathop{E}_{w_k,v_k} \{J_{k+1}(\max(0, y - w_k - v_k))\}. \tag{C.7}$$

Suppose $G_k(y)$ is convex in y and has a minimum at S_k. Due to the constraint $y \geq x_k$, the minimizing y is S_k if $x_k < S_k$ and x_k otherwise. According to the reverse transformation $u_k = y - x_k$, the optimal policy takes the form stated in Equation (5.8)

$$u_k^*(x_k) = \begin{cases} S_k - x_k & \text{if } x_k < S_k, \\ 0 & \text{otherwise.} \end{cases} \tag{C.8}$$

We proceed to show by induction that the cost-to-go functions $J_k(x_k)$ (and hence also the functions G_k) are convex.

Since $J_N(x_N)$ is equal to zero, it is convex. The minimum of G_{N-1} can be found by taking the derivative with respect to y:

$$\begin{aligned}\frac{d}{dy}G_{N-1} &= c + \frac{d}{dy}H(y) \\ &= c + \frac{d}{dy}\left[cu_k + h\underset{w_k,v_k}{E}\{\max(0, y - w_k - v_k)\} + p\underset{w_k}{E}\{\max(0, w_k - y)\}\right] \\ &= c + h\frac{d}{dy}\left[\underset{w_k,v_k}{E}\{\max(0, y - w_k - v_k)\}\right] + p\frac{d}{dy}\left[\underset{w_k}{E}\{\max(0, w_k - y)\}\right].\end{aligned} \tag{C.9}$$

We determine the two derivative terms in the right-hand side of the above equation separately. Given the probability density function of the random variables associated with the demand and stock loss, denoted $f_W(w_k)$ and $f_V(v_k)$ respectively, the last derivative term becomes

$$\begin{aligned}\frac{d}{dy}\left[\underset{w_k}{E}\{\max(0, w_k - y)\}\right] &= \frac{d}{dy}\int_{\infty}^{\infty} \max(0, w_k - y) f_W(w_k) dw_k \\ &= \frac{d}{dy}\int_{y}^{b} (w_k - y) f_W(w_k) dw_k \\ &= \frac{d}{dy}\int_{y}^{b} w_k f_W(w_k) dw_k - \frac{d}{dy}\int_{y}^{b} y f_W(w_k) dw_k \\ &= \frac{d}{dy}\int_{y}^{b} w_k f_W(w_k) dw_k - y\frac{d}{dy}\int_{y}^{b} f_W(w_k) dw_k - \int_{y}^{b} f_W(w_k) dw_k.\end{aligned} \tag{C.10}$$

Using the second fundamental theorem of calculus,

$$\begin{aligned}\frac{d}{dy}\left[\underset{w_k}{E}\{\max(0, w_k - y)\}\right] &= -\frac{d}{dy}\int_{b}^{y} w_k f_W(w_k) dw_k + y\frac{d}{dy}\int_{b}^{y} f_W(w_k) dw_k - \int_{y}^{b} f_W(w_k) dw_k \\ &= -y f_W(y) + y f_W(y) - \text{Prob}(y \le w_k \le b) \\ &= -\text{Prob}(y \le w_k \le b).\end{aligned} \tag{C.11}$$

Similarly, the first derivative term in the right-hand side of Equation (C.9) can be determined using double integrals (since there are two random variables associated with the demand and stock loss), and can be shown to equal to

$$\frac{d}{dy}\underset{w_k,v_k}{E}\{\max(0, y - w_k - v_k)\} = \text{Prob}(0 \le w_k + v_k \le y). \tag{C.12}$$

Substituting Equation (C.11) and Equation (C.12) into Equation (C.9), the derivative of

G_{N-1} becomes

$$\frac{d}{dy}G_{N-1} = c - p\text{Prob}(y \le w_k \le b) + h\text{Prob}(0 \le w_k + v_k \le y). \tag{C.13}$$

As $y \to -\infty$, $\frac{d}{dy}G_{N-1}$ approaches $c-p$, and since $c < p$, $\frac{d}{dy}G_{N-1}$ is negative. As $y \to \infty$, $\frac{d}{dy}G_{N-1}$ approaches $c+h$, and is positive. Therefore, G_{N-1} is convex, and given the convexity of J_N, convexity of J_{N-1} is proved.

An optimal policy at time $N-1$ then is given by

$$u^*_{N-1}(x_{N-1}) = \begin{cases} S_{N-1} - x_{N-1} & \text{if } x_{N-1} < S_{N-1}, \\ 0 & \text{otherwise,} \end{cases} \tag{C.14}$$

and the cost-to-go function J_{N-1} takes the value

$$J_{N-1}(x_{N-1}) = \begin{cases} c(S_{N-1} - x_{N-1}) + H(S_{N-1}) - p\underset{w_k}{E}\{w_k\} & \text{if } x_{N-1} < S_{N-1}, \\ H(x_{N-1}) & \text{otherwise.} \end{cases} \tag{C.15}$$

We can repeat these steps to show that for all $k = N-2, \ldots, 0$, if J_{k+1} is convex, then we have

$$J_k(x_k) = \begin{cases} c(S_k - x_k) + H(S_k) - p\underset{w_k}{E}\{w_k\} + E\{J_{k+1}(S_k - w_k - v_k)\} & \text{if } x_k < S_k, \\ H(x_k) - p\underset{w_k}{E}\{w_k\} + E\{J_{k+1}(x_k - w_k - v_k)\} & \text{otherwise.} \end{cases} \tag{C.16}$$

where S_k is the scalar that minimizes $cy + H(y) + E\{J_{k+1}(\max(0, y - w_k - v_k))\}$. Furthermore, J_k is convex, and thus the optimality of the policy 5.8 is proved.

Appendix D

Calculation of P_k

D.1 P_k Using Auto-ID

P_{k+1} is the conditional probability distribution of the actual inventory x_{k+1} given the information vector I_{k+1}, and is a vector whose elements are the individual conditional probabilities $p_{k+1}^{(j)}$, $j = 0, 1, \ldots, n$, defined as

$$\begin{aligned} p_{k+1}^{(j)} &\equiv \text{Prob}(x_{k+1} = j | I_{k+1}) \\ &= \text{Prob}(x_{k+1} = j | z_0, z_1, \ldots, z_{k+1}, u_0, u_1, \ldots, u_k). \end{aligned} \tag{D.1}$$

Using the property of conditional probability, the above can be expressed as

$$\begin{aligned} p_{k+1}^{(j)} &\equiv \text{Prob}(x_{k+1} = j | z_0, z_1, \ldots, z_{k+1}, u_0, u_1, \ldots, u_k) \\ &= \frac{\text{Prob}(x_{k+1} = j, z_{k+1} | z_0, z_1, \ldots, z_k, u_0, u_1, \ldots, u_k)}{\text{Prob}(z_{k+1} | z_0, z_1, \ldots, z_k, u_0, u_1, \ldots, u_k)} \\ &= \frac{\text{Prob}(x_{k+1} = j, z_{k+1} | I_k, u_k)}{\text{Prob}(z_{k+1} | I_k, u_k)}. \end{aligned} \tag{D.2}$$

Again applying the property of conditional probability, the numerator of the above equation can be expressed as

$$\text{Prob}(x_{k+1} = j, z_{k+1} | I_k, u_k) = \text{Prob}(z_{k+1} | I_k, u_k, x_{k+1} = j)\text{Prob}(x_{k+1} = j | I_k, u_k). \tag{D.3}$$

The first probability term on the right-hand side of the equation is simply the measurement probability $r_{z_{k+1},j}$. The second probability term can be expanded using the total probability theorem:

$$\begin{aligned}\text{Prob}(x_{k+1}=j|I_k,u_k) &= \text{Prob}(x_k=0|I_k,u_k)\text{Prob}(x_{k+1}=j|I_k,u_k,x_k=0)+\ldots \\ &\quad + \text{Prob}(x_k=n|I_k,u_k)\text{Prob}(x_{k+1}=j|I_k,u_k,x_k=n) \\ &= p_k^{(0)}\pi_{0,j}(u_k)+\ldots+p_k^{(n)}\pi_{n,j}(u_k) \\ &= \sum_{i=0}^{n} p_k^{(i)}\pi_{i,j}(u_k).\end{aligned} \tag{D.4}$$

The numerator in Equation (D.2) then becomes

$$\text{Prob}(x_{k+1}=j,z_{k+1}|I_k,u_k)=\sum_{i=0}^{n} p_k^{(i)}\pi_{i,j}(u_k)r_{z_{k+1},j}. \tag{D.5}$$

The denominator in Equation (D.2) also can be expanded using the total probability theorem to become

$$\begin{aligned}\text{Prob}(z_{k+1}|I_k,u_k) &= \text{Prob}(x_k=0|I_k,u_k)\text{Prob}(z_{k+1}|I_k,u_k,x_k=0)+\ldots \\ &\quad + \text{Prob}(x_k=n|I_k,u_k)\text{Prob}(z_{k+1}|I_k,u_k,x_k=n)+\ldots \\ &= \sum_{i=0}^{n} p_k^{(i)}\text{Prob}(z_{k+1}|I_k,u_k,x_k=i) \\ &= \sum_{i=0}^{n} p_k^{(i)}\Big[\text{Prob}(x_{k+1}=0|I_k,u_k,x_k=i)\text{Prob}(z_{k+1}|I_k,u_k,x_k=i,x_{k+1}=0)+\ldots \\ &\quad + \text{Prob}(x_{k+1}=n|I_k,u_k,x_k=i)\text{Prob}(z_{k+1}|I_k,u_k,x_k=i,x_{k+1}=n)\Big] \\ &= \sum_{s=0}^{n}\sum_{i=0}^{n} p_k^{(i)}\pi_{i,s}(u_k)r_{z_{k+1},s}.\end{aligned} \tag{D.6}$$

Substituting Equation (D.5) and Equation (D.6) into Equation (D.2), we arrive at the final expression for the conditional probability $p_{k+1}^{(j)}$

$$p_{k+1}^{(j)}=\frac{\sum_{i=0}^{n} p_k^{(i)}\pi_{i,j}(u_k)r_{z_{k+1},j}}{\sum_{s=0}^{n}\sum_{i=0}^{n} p_k^{(i)}\pi_{i,s}(u_k)r_{z_{k+1},s}}. \tag{D.7}$$

D.2 P_k Without Using Auto-ID

At the beginning of period k, the information vector contains all the past sales and shipment receipt data, expressed as

$$I_k = (a_0, a_1, \ldots, a_{k-1}, h_0, h_1, \ldots, h_{k-1}). \tag{D.8}$$

We begin with the definition of P_{k+1} the conditional probability distribution of the actual inventory x_{k+1} given the information vector I_{k+1}. P_{k+1} is a vector whose individual element is

$$\begin{aligned} p_{k+1}^{(j)} &\equiv \text{Prob}(x_{k+1} = j | I_{k+1}), \qquad j = 0, 1, \ldots, n \\ &= \text{Prob}(x_{k+1} = j | a_0, a_1, \ldots, a_k, h_0, h_1, \ldots, h_k) \\ &= \frac{\text{Prob}(x_{k+1} = j, a_k | a_0, a_1, \ldots, a_{k-1}, h_0, h_1, \ldots, h_k)}{\text{Prob}(a_k | a_0, a_1, \ldots, a_{k-1}, h_0, h_1, \ldots, h_k)} \\ &= \frac{\text{Prob}(x_{k+1} = j, a_k | I_k, h_k)}{\text{Prob}(a_k | I_k, h_k)}. \end{aligned} \tag{D.9}$$

The denominator of the last expression can be expressed as follows using the total probability theorem:

$$\begin{aligned} \text{Prob}(a_k | I_k, h_k) &= \text{Prob}(x_k = 0 | I_k, h_k)\text{Prob}(a_k | I_k, h_k, x_k = 0) + \ldots \\ &\quad + \text{Prob}(x_k = n | I_k, h_k)\text{Prob}(a_k | I_k, h_k, x_k = n) \\ &= \sum_{i=0}^{n} p_k^{(i)} \text{Prob}(a_k | x_k = i, h_k). \end{aligned} \tag{D.10}$$

The conditional probability $\text{Prob}(a_k | x_k = i, h_k)$ can be computed from the dynamics of the system (Equation (7.2)) and the problem data, which includes the distribution for w_k and v_k.

Applying the property of conditional probability, the numerator of Equation (D.9) becomes

$$\text{Prob}(x_{k+1} = j, a_k | I_k, h_k) = \text{Prob}(x_{k+1} = j | I_k, h_k)\text{Prob}(a_k | I_k, h_k, x_{k+1} = j). \tag{D.11}$$

The second conditional probability in the right-hand side of the above equation can be expanded using the total probability theorem

$$\begin{aligned}\text{Prob}(a_k|I_k,h_k,x_{k+1}=j) &= \text{Prob}(x_k=0|I_k,h_k,x_{k+1}=j)\text{Prob}(a_k|I_k,h_k,x_{k+1}=j,x_k=0)+\\ &\quad \ldots + \text{Prob}(x_k=n|I_k,h_k,x_{k+1}=j)\text{Prob}(a_k|I_k,h_k,x_{k+1}=j,x_k=n)\\ &= \sum_{i=0}^{n}\text{Prob}(x_k=i|I_k,h_k,x_{k+1}=j)\text{Prob}(a_k|I_k,h_k,x_{k+1}=j,x_k=i).\end{aligned} \tag{D.12}$$

The first probability term in the summation, $\text{Prob}(x_k=i|I_k,h_k,x_{k+1}=j)$, can be expressed as

$$\begin{aligned}\text{Prob}(x_k=i|I_k,h_k,x_{k+1}=j) &= \frac{\text{Prob}(x_k=i,x_{k+1}=j|I_k,h_k)}{\text{Prob}(x_{k+1}=j|I_k,h_k)}\\ &= \frac{\text{Prob}(x_{k+1}=j|x_k=i,I_k,h_k)\text{Prob}(x_k=i|I_k,h_k)}{\text{Prob}(x_{k+1}=j|I_k,h_k)}\\ &= \frac{\pi_{i,j}(h_k)p_k^{(i)}}{\text{Prob}(x_{k+1}=j|I_k,h_k)}.\end{aligned} \tag{D.13}$$

By substituting the above equation into Equation (D.12) and the subsequent resulting expression into Equation (D.11), we arrive at

$$\text{Prob}(x_{k+1}=j,a_k|I_k,h_k) = \sum_{i=0}^{n} p_k^{(i)}\pi_{i,j}(h_k)\text{Prob}(a_k|I_k,h_k,x_{k+1}=j,x_k=i). \tag{D.14}$$

The quantity $\text{Prob}(a_k|I_k,h_k,x_{k+1}=j,x_k=i)$ in the right-hand side of the above equation can also be expressed in terms of the quantities we can compute from the problem data. Specifically,

$$\begin{aligned}\text{Prob}(a_k|I_k,h_k,x_{k+1}=j,x_k=i) &= \frac{\text{Prob}(a_k,x_{k+1}=j|x_k=i,h_k)}{\text{Prob}(x_{k+1}=j|x_k=i,h_k)}\\ &= \frac{\text{Prob}(x_{k+1}=j|x_k=i,h_k,a_k)\text{Prob}(a_k|x_k=i,h_k)}{\pi_{i,j}(h_k)}.\end{aligned} \tag{D.15}$$

Substituting the last expression into Equation (D.14) and Equation (D.10) into Equation (D.9), we arrive at the final expression for $p_{k+1}^{(j)}$

$$
\begin{aligned}
p_{k+1}^{(j)} &= \frac{\sum_{i=0}^{n} p_k^{(i)} \text{Prob}(x_{k+1} = j | x_k = i, h_k, a_k) \text{Prob}(a_k | x_k = i, h_k)}{\sum_{i=0}^{n} p_k^{(i)} \text{Prob}(a_k | x_k = i, h_k)} \\
&= \Phi(P_k, h_k, a_k).
\end{aligned}
\tag{D.16}
$$

References

[AL81] A. A. Arens and J. K. Loebbecke. *Applications of Statistical Sampling to Auditing.* Prentice-Hall, Englewood Cliffs, NJ, 1981.

[Ber] Paul A. Bernard. Cycle counting: The missing link. *Production and Inventory Management*, 25(4).

[Ber00] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control, Vol 1.* Athena Scientific, Belmont, Massachusetts, 2000.

[BHO84] Uttarayan Bagchi, Jack C. Hayya, and J. Keith Ord. Concepts, theory, and techniques: Modeling demand during lead time. *Decision Sciences*, 15:157–176, 1984.

[BIC] K. L. Brown, R. A. Inman, and J. A. Calloway. Measuring the effects of inventory inaccuracy in mrp inventory and delivery performance. *Production Planning and Control*, 12(1).

[Bro01] David L. Brock. The physical markup language: A universal language for physical objects. *White Paper, MIT Auto-ID Center*, February 2001.

[BS83] James R. Buck and Randall P. Sadowski. Optimum stratified sampling. *IIE Transactions*, 15:119–126, 1983.

[Buk84] David W. Buker. Inventory accuracy. *Inventory Management Reprints*, APICS:221–223, 1984.

[Can85] Jim Cantwell. The how and why of cycle counting; the abc method. *Production and Inventory Management Journal*, Second Quarter, 1985.

[Cho86] Vinnie Chopra. How to set up an effective inventory accuracy program. *29th Annual APICS International Conference Proceedings*, page 238, 1986.

[CKL88] M. A. Cohen, P. R. Kleindorfer, and H. L. Lee. Service constrained (s,s) inventory systems with priority demand classes and lost sales. *Management Science*, 34:482499, 1988.

[Coc77] William Cochran. *Sampling Techniques.* John Wiley and Sons, New York, 3 edition, 1977.

[DH59] Tore Dalenius and J.L. Hodges. Minimum variance stratification. *American Statistical Journal*, pages 198–203, 1959.

[Ede84] Martin P. Edelman. Cycle counting for distribution inventories. *Readings in Production and Inventory Control and Planning*, APICS:29, 1984.

[EGR93] Ricardo Ernst, Jose-Luis Guerrero, and Alan Roshwalb. A quality control approach for monitoring inventory stock levels. *J. Opl Res. Soc.*, 44(11):1115–1127, 1993.

[Eur01] ECR Europe. *Shrinkage: Introducing a Collaborative Approach to Reducing Stock Loss in the Supply Chain.* ECR Europe, Brussels, Belgium, 2001.

[Fre80] R. L. French. Accurate wip inventory - a critical mrp system requirement. *Production and Inventory Management*, 21:17–22, 1980.

[Gra87] Richard T. Graff. The limitations of cycle counting. *Production and Inventory Management Journal*, Fourth Quarter:39–43, 1987.

[Har98] Marilyn K. Hart. Improving inventory accuracy using control charts. *Production and Inventory Management*, Third Quarter(11):44–48, 1998.

[Hol03] Richard C. Hollinger. *2002 National Retail Security Survey Final Report.* University of Florida, Gainesville, Florida, 2003.

[HS96] W. J. Hopp and M. L. Spearman. *Factory Physics.* McGraw Hill, Boston, Massachusetts, 1996.

[IM72] Donald L. Iglehart and Richard C. Morey. Inventory systems with imperfect asset information. *Management Science*, 18(8):388–394, 1972.

[Jos00] Yogeshi V. Joshi. Information visibility and its effect on supply chain dynamics. Masters thesis, Massachusetts Institute of Technology, June 2000.

[KKRW] L. J. Krajewski, B. E. King, L. P. Ritzman, and D. S. Wong. Kanban, mrp, and shaping the manufacturing environment. *Management Science*, 33(1).

[MD86] Richard C. Morey and David A. Dittman. Optimal timing of account audits in internal control. *Management Science*, 32(3):272–282, 1986.

[MG87] James W. Martin and Peter S. Goodrich. Minimizing sample size for given accuracy in cycle counting. *Production and Inventory Management Journal*, Fourth Quarter:24–27, 1987.

[Mor85] Richard C. Morey. Estimating service level impacts from changes in cycle count, buffer stock, or corrective action. *Journal of Operations Management*, 5(4):411–418, 1985.

[Nee87] P.S. Neely. Simple mathematics applied to inventory accuracy. *Production and Inventory Management*, 28(3):64–68, 1987.

[RDT01] A. Raman, N. DeHoratius, and Z. Ton. Execution: The missing link in retail operations. *California Management Review*, 43(3):136–152, 2001.

[Red84] John G. Reddock. How to count a cycle. *Readings in Production and Inventory Control and Planning*, APICS:28, 1984.

[Sch01] Tom Scharfield. An analysis of the fundamental constraints on low cost passive rfid system design. *The Auto-ID Center, MIT*, 2001.

[You] Jan B. Young. The limits of cycle counting. *29th Annual APICS International Conference Proceedings*, APICS.