

Achieving Cost-effective Supply Chain Agility For The Semiconductor Industry

2005

Mariah Jeffery
University of Central Florida

Find similar works at: <http://stars.library.ucf.edu/etd>

University of Central Florida Libraries <http://library.ucf.edu>

 Part of the [Engineering Commons](#)

STARS Citation

Jeffery, Mariah, "Achieving Cost-effective Supply Chain Agility For The Semiconductor Industry" (2005). *Electronic Theses and Dissertations*. 574.

<http://stars.library.ucf.edu/etd/574>

This Doctoral Dissertation (Open Access) is brought to you for free and open access by STARS. It has been accepted for inclusion in Electronic Theses and Dissertations by an authorized administrator of STARS. For more information, please contact lee.dotson@ucf.edu.

ACHIEVING COST-EFFECTIVE SUPPLY CHAIN AGILITY FOR THE
SEMICONDUCTOR INDUSTRY

by

MARIAH MCMURRAN JEFFERY
B.S. University of Central Florida, 2002
M.S. University of Central Florida, 2003

A dissertation submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy
in the Department of Industrial Engineering and Management Systems
in the College of Engineering and Computer Science
at the University of Central Florida
Orlando, FL

Fall Term
2005

Major Professor: Dr. Renee J. Butler

ABSTRACT

Supply chain agility has been receiving a lot of attention in recent literature as a way for organizations to become more responsive to change and improve their customer service levels. However, agility is typically dealt with qualitatively, and organizations are usually unsure of the steps to take to improve their agility and the customer service level to target. This research studies supply chain agility based on a case study of Intel Corporation, a large semiconductor manufacturer.

Here, agility is defined as the ability to satisfy customer demands by reacting effectively to changes in market stimuli. Reacting effectively does not mean reacting to every change in supply or demand. Doing so means increasing supply chain variability unnecessarily, which is amplified by the bullwhip effect. The essence of supply chain agility is determining the degree to which variability should be managed through artificial means such as safety stock, and appropriate triggers for changing production levels and inventory targets.

The purpose of this research is to examine some of the factors that influence supply chain agility and identify a cost-effective plan for achieving it. The first phase addresses the problem of identifying target inventory and customer service levels based on regression analysis of historical data and financial analysis of inventory holding costs and stock-out costs. The impact of three factors (forecast error, order lead-time, and demand variability) on the relationship between inventory and customer service level is also examined.

The second phase of the research evaluates strategies for production and inventory control with the goal of finding the appropriate trade-off between minimizing cost (of holding inventory and stock-outs) and minimizing variability. Control policies based on the Exponentially Weighted Moving Average (EWMA) control chart with control limits on demand forecasts are proposed to detect when tighter control of processes is necessary. A Monte Carlo supply chain simulation is used to evaluate the performance of these policies under various levels of forecast error and demand variability.

Results indicate that several control chart-based policies outperform Intel's current planning policy in terms of cost without significantly increasing variability. The selection of the appropriate policy must be based on the decision-makers' desire to minimize cost compared to the desire to minimize variability, as each policy results in a trade-off between these two objectives.

This dissertation is dedicated to my husband, Casey, and my parents.

ACKNOWLEDGMENTS

I would like to acknowledge my advisor, Dr. Renee Butler, for seeing me, her first Ph.D. student, through to graduation despite the obstacles. I would also like to thank my committee members for their contributions to this dissertation, especially Dr. Linda Malone, for stepping in at the last minute and providing valuable advice.

I would like to thank several members of Intel's supply chain group, especially Jim Kellso, Rodney Wegemer, and Dennis Arnow for their support of this project and for spending many hours brainstorming with me to define and scope the project.

I would like to thank my parents for instilling in me the importance of education and for pushing me to do my best even when I resisted.

Finally, I would like to thank my husband, Casey, for his moral support and his extremely valuable coding help. Without it I would probably still be trying to get my model to run.

TABLE OF CONTENTS

LIST OF TABLES	vii
LIST OF FIGURES	viii
CHAPTER 1 INTRODUCTION	1
1.1 Supply Chain Agility	2
1.2 Problem Motivation	3
1.3 Purpose	4
1.4 Organization	5
CHAPTER 2 LITERATURE REVIEW	7
2.1 Supply Chain Agility	8
2.1.1. Related Constructs	9
2.1.2. Factors that Influence Agility	10
2.1.3. Modeling Agility	12
2.2 Service Level Models	17
2.3 Methods for Detecting/Reacting to Change	19
2.3.1. System Dynamics and Control Theory	20
2.3.2. Control Charts	22
2.4 Summary of Relevant Literature and Contribution of this Study	24
CHAPTER 3 RELATIONSHIP BETWEEN INVENTORY, CUSTOMER SERVICE LEVEL, AND COST	26
3.1 Introduction to Case Study	26
3.2 Regression Modeling	27
3.2.1. Defining the Customer Service Level and Inventory Relationship	28
3.2.2. Analysis of Factors Affecting the Service Level and Inventory Relationship	30
3.3 Customer Service Level Cost Equation	34
3.4 Determination of Minimum Cost Service Level	37
3.5 Conclusions	38
CHAPTER 4 INVENTORY CONTROL POLICIES	40
4.1 Inventory Planning Policies	41
4.2 Monte Carlo Supply Chain Simulation Model	43
4.3 Generation of Forecast Data	49
4.3.1. Overall Cost and Variance Results by Policy	53
4.3.2. Paired Comparison of Policies	66
4.3.3. Summary of Results	72
4.3.3.1 General Results	73
4.3.3.2 Intel Specific Results	74
CHAPTER 5 SUMMARY AND CONCLUSIONS	76
5.1 Summary and Contribution	76
5.2 Future Research	77
APPENDIX – CODE FOR SUPPLY CHAIN MODEL	79
LIST OF REFERENCES	98

LIST OF TABLES

<u>Table</u>	<u>Page</u>
Table 1. Characteristics of Lean and Agile Supply Chains (adapted from Christopher and Towill, [2002]).	10
Table 2. Summary of Supply Chain Agility Literature.....	17
Table 3. Model Validation Results	34
Table 4. Planning Policies Evaluated by Supply Chain Model	42
Table 5. Summarized Results for All Products, Forecast Errors	55
Table 6. Effect of λ Parameter and Size of Control Limits on Number of Out of Control Points Detected by Control Chart.....	60
Table 7. Rankings of Cost and Variance by Policy and Level of Forecast Error	61
Table 8. Customer Service Level Resulting from Each Policy by Level of Forecast Error	62
Table 9. Comparison of Cost and Variance Rankings by Product Variability Classification.....	64
Table 10. Comparison of Weekly vs. Monthly Planning Policies.....	67
Table 11. Comparison of Monthly vs. Monthly to Weekly Control Chart with $\lambda=0.3$, $\sigma=1$ Planning Policies.....	68
Table 12. Comparison of Monthly-to-Weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits vs. Monthly-to-Weekly Control Chart with $\lambda=0.1$ and 1σ Control Limits...	69
Table 13. Comparison of Monthly vs. Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits Planning Policies	70
Table 14. Comparison of Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits vs. Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1.5σ Control Limits	71
Table 15. Comparison of Biweekly vs. Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits	72

LIST OF FIGURES

<u>Figure</u>	<u>Page</u>
Figure 1. Weeks of Inventory vs. Expected Service Level with 95 Percent Confidence (Product Group 1).....	29
Figure 2. Weeks of Inventory vs. Expected Service Level with 95 Percent Confidence (Product Group 2).....	30
Figure 3. Relationship between order lead-time and inventory required to achieve a 95 percent service level with 95 percent confidence limits.....	31
Figure 4. Relationship between forecast error and inventory required to achieve a 95 percent service level with 95 percent confidence limits.....	32
Figure 5. Relationship between coefficient of variation of demand and inventory required to achieve a 95 percent service level with 95 percent confidence limits.	33
Figure 6. Sensitivity Analysis on the Effect of Product Margin on Ideal Service Level...36	
Figure 7. Cost and Service Level versus Weeks of Inventory (Product Group 1).....	37
Figure 8. Cost and Service Level versus Weeks of Inventory (Product Group 2).....	38
Figure 9. Case Supply Chain Node Diagram.....	43
Figure 10. Cost and Variance Results for Frequency-Based Policies	54
Figure 11. Effect of the λ Parameter on Cost.....	57
Figure 12. Effect of the Sigma Multiplier Parameter on Cost	57
Figure 13. Effect of In Control and Out of Control Updating Frequency on Cost.....	58
Figure 14. Effect of the λ Parameter on Variance	58
Figure 15. Effect of the Sigma Multiplier Parameter on Variance	59
Figure 16. Effect of In Control and Out of Control Updating Frequency on Variance.....	59

CHAPTER 1

INTRODUCTION

Supply chain agility is studied in this dissertation. Supply chain agility involves developing a strategy to respond to uncertainty and changes in market conditions, demand and supply over time. The supply chain strategy must be adaptive to respond to changes, but not be too sensitive to chase noise (variability) in the system. The ideal supply chain agility level is based on a trade-off between costs of holding inventory and costs of poor customer satisfaction, including lost sales. This approach protects companies from holding large inventories that risk obsolescence, damage and devaluation. In addition, agile strategies focus on determining the appropriate customer service level target.

In this dissertation, a regression modeling approach is developed that identifies the relationship between inventory, cost and customer service level. Understanding these relationships is then used to determine a policy for production and inventory control that results in the ability to react effectively to change but minimizes supply chain variability.

In addition, a case study of the approach in the semiconductor industry is presented. The approach provides a means to identify the relationship between inventory, cost and customer service level for the particular company and to develop and evaluate an inventory policy accounting for the effects of uncertainty in the supply chain.

In this chapter, an introduction to supply chain agility is presented and the supply chain issues important in the semiconductor industry are specifically discussed, which provides the motivation for this research.

1.1 Supply Chain Agility

Supply chain management offers great potential for organizations to reduce costs and improve customer service performance. Supply chain agility, in particular, has been receiving a great deal of attention in recent literature for its potential to provide a competitive advantage to organizations that determine successful and yet cost-effective agile strategies.

Due to the increasing complexity of supply chains, high cost of holding inventory and penalties for stock-outs, and non-stationary nature of demand, it is important to select the best possible inventory strategy; one that minimizes the probability of stock-outs, while keeping inventory holding costs as low as possible. However, organizations in these complex environments often base target inventory and customer service levels on experience and gut feelings rather than mathematical models. According to Ettl et al. (2000, pg. 216), “a common problem for asset managers is not knowing how to quantify the trade-off between service levels and the investment in inventory required to support those service levels.” They are also unclear about where to invest resources and efforts to improve delivery performance to customers and whether their investments in resources to improve these factors will pay off.

In addition, while many companies have focused on agility, the concept is still somewhat ambiguous, which makes the determination of a unified general strategy toward improving agility a difficult task. “Companies are starting to become aware of the importance of agility but have not yet linked the concept to concrete actions” [Katayama and Bennett, 1999]. While agility means being able to react effectively to change, in a highly stochastic environment such as the semiconductor industry, noise can disguise true changes. Therefore, knowing when to react can be quite challenging.

1.2 Problem Motivation

This work is motivated by interaction with Intel Corporation, a large semiconductor manufacturer. Semiconductor supply networks are becoming increasingly complex and dynamic. The challenges of semiconductor logistics include the high cost of inventory, short product lifecycles, increasing customer expectations, and a widely dispersed supply chain [Maltz et al, 2000]. A semiconductor supply network can be classified as a multi-echelon supply network due to the multiple tiers that are controlled by a single organization. The network starts with the first tier of raw materials, which includes bare silicon wafers that are shipped from subcontractors to wafer fabrication facilities (fabs) for manufacturing. The manufacturing process, which consists of hundreds of process steps with re-entrant process flows, during which each wafer is subdivided into dies of integrated circuits, takes an average of 10 weeks.

After manufacturing, wafers are sent to E-test and Sort, where malfunctioning die are identified. They are then stored in an intermediate buffer for work-in-process inventory before being sent to Assembly/Test, where they are separated into individual chips and packaged. Next, inventory is sent to finished inventory warehouses until it is shipped to customers to fulfill awaiting orders. The processing time for the back-end of the supply chain (after manufacturing is completed) is about 10 weeks. Consequently, the total supply chain lead-time is about 20 weeks.

Despite these long lead-times, more than 50 percent of orders are placed within four weeks of their requested delivery dates. Demand is volatile, and since product lifecycles typically last 1.5 years, it is also non-stationary, making both forecasting and

inventory management difficult. Rapid drops in demand can leave companies with excess inventory at the end of a product life cycle, which must be scrapped. Furthermore, rapid increases in demand can lead to stock-outs and lost revenue when customers turn to the competition. For these reasons, it is important to find an inventory control policy that results in the ability to detect and react to changes quickly in this highly variable environment.

1.3 Purpose

The purpose of this dissertation is to examine some of the factors that influence supply chain agility in a highly stochastic environment and identify a cost-effective plan for achieving agility. The following statement summarizes the fundamental research thrust.

This research develops a method for determining and controlling inventory levels for stochastic non-stationary demand and this method has the ability to identify appropriate triggers to warrant changes to inventory and production while minimizing unnecessary reacting to noise.

Specifically, this study seeks to:

- Determine the relationship between customer service level, inventory, and cost;
- Find a cost-effective customer service level and target inventory level for specific products based on their characteristics;
- Identify the factors that have the greatest impact on the inventory/customer service level relationship and examine the effects of these factors; and

- Determine a policy for updating inventory targets that results in the ability to react effectively to change but minimizes supply chain variability.

A case study in the semiconductor industry is used to demonstrate the methodology developed as well as gain insights into the role and impact of agility in the supply chain. Although this research is illustrated with an application in the semiconductor industry, it is also applicable to any organization with similar challenges, i.e., a complex supply chain, long supply chain lead-times, high costs of holding inventory, and demanding customers. Therefore, the ultimate goal of this research is to provide a methodology for achieving supply chain agility that can be applied in the semiconductor and similar industries.

1.4 Organization

This remainder of this dissertation is organized as follows. Chapter 2 provides a review of literature related to supply chain agility and inventory control. This section explores the relation of supply chain agility to similar constructs such as flexibility, lean and responsiveness. A definition of agility is given to provide context for this research. Additionally, Chapter 2 highlights previous work in identifying factors that influence agility and agility modeling. Inventory control models and methods for detecting and reacting to change are also incorporated. Finally, the literature review concludes with identifying the research gap to be addressed in this dissertation.

Chapter 3 presents a research methodology for finding the most cost-effective inventory target and resulting customer service level. First the relationship between

customer service level and inventory is examined. Next, several factors that affect this relationship are identified and their impact is examined and quantified.

Chapter 4 uses the inventory targets identified in Chapter 3 by determining and evaluating policies for resetting the inventory targets. Several control chart-based policies are proposed and evaluated for their ability to detect when tighter control of processes is needed. The cost and variability resulting from each policy is compared.

In both Chapters 3 and 4, a case study based on the experience of the corporate sponsor is utilized to illustrate the methodology. Additionally, the analysis of the case study provides insights on the trade-off between inventory levels and supply chain costs. It is shown how this trade-off can be used to make the supply chain less vulnerable to uncertainty and more likely to effectively meet customer demand.

In Chapter 5, the results and contributions of this research as well as proposed extensions and new directions to explore are summarized.

CHAPTER 2

LITERATURE REVIEW

A supply chain can be defined as “a network of facilities and distribution options that functions to procure materials, transform these materials into intermediate and finished products, and distribute these finished products to customers” [Dong, 2001]. The supply chain encompasses all of the activities necessary to produce a product or to fulfill a customer’s request. Its elements typically include manufacturers, suppliers, transporters, retailers, and customers. Supply chain management involves the management of flows between these stages of the supply chain to maximize total profitability.

Three fundamental sources of uncertainty exist in a supply chain: demand (volume and mix), process (yield, machine downtimes), and supply (part quality, reliability of delivery). Uncertainty is amplified as it propagates upstream (from customer to supplier) through the supply chain due to the bullwhip effect [Lee et al., 1997]. Its consequences include difficulty of accurate supply and demand planning and the necessity of safety stock. An agility-based strategy can improve customer service and minimize the consequences of the bullwhip effect by avoiding reacting to noise and unnecessarily increasing production variability.

In the remainder of this chapter, recent literature is reviewed on supply chain agility, customer service level models, and strategies for reacting to change while minimizing unnecessary variability in the supply chain. In Section 2.1, supply chain agility is defined and compared to other constructs, and several agility modeling approaches are presented. The purpose of this section is to gain a general understanding of agility and how it can be addressed. Next, inventory models that focus on achieving a minimum customer service level are discussed.

Then methods for detecting and reacting to changes in demand while minimizing variability are described. Finally, research gap that has been identified is summarized.

2.1 Supply Chain Agility

In today's rapidly changing environment, it is extremely important for an organization's supply network to be able to quickly recognize and react to change effectively. Doing so can greatly increase an organization's customer service levels, time to market, and competitive advantage. For this reason, agility receives a great deal of attention in supply chain literature as a way for organizations to become more responsive to changes in the business environment [Dong, 2001; Christopher and Towill, 2001; Katayama and Bennett, 1999]. According to Goldman et al. (1994), agility is the competency that sustains world class performance over time.

The "agile enterprise" began in 1991 during a four-month long collaborative agility workshop funded by the United States government with a goal of developing the successor to the Japanese "lean" manufacturing. Here, agility is defined as the ability of an organization to thrive in a continuously changing, unpredictable business environment [Dove, 1999]. Since then, many other definitions for agility have been employed. Katayama and Bennett [1999] define agility as "a set of abilities for meeting widely varied customer requirements for price, specification, quality, quantity, and delivery." The authors define four underlying principles to agility: (i) delivering value to a customer, (ii) being ready for change, (iii) valuing human knowledge and skills, and (iv) forming virtual partnerships.

Uncertainty, or stochasticity, is one of the most important problems in supply chain management [Sabri and Beamon, 2000]. Managing the stochasticity involved in supply networks by planning for and reacting effectively to changes when they occur is the essence of supply

chain agility. Therefore, the definition of agility for this research is “the ability to satisfy customer demands by reacting effectively to changes in market stimuli.” Encompassed in this definition is being able to store the right amount of inventory in the right places to be able to deliver orders on time, and knowing when to react to change by increasing or decreasing inventory levels and/or production in order to match supply with demand. Agility does not mean reacting to every change in supply, demand, and forecast that is experienced in the supply chain because many of these changes are simply noise in the system that can be managed through safety stock, and reacting to each of these changes means increasing the variability of production and other processing steps, which can lead to longer throughput times. Determining the degree to manage variability through safety stock and through production and inventory control is a central goal of this research.

2.1.1. Related Constructs

Several constructs similar to agility, including flexibility, leanness, and responsiveness are given in the supply chain literature. This section is devoted to describing the difference between these constructs.

Slack [1983] defines *flexibility* as “the number of different positions, or flexible options, that can be achieved with existing resources,” in terms of both cost and time. Mahoney and Plossl [1997] outline three types of flexibility that are important in a high-mix, low-volume environment as product mix, volume, and workforce flexibility. Mix flexibility is the ability to manufacture several different products using the same resources while process flexibility is the ability of a production line to handle drastic changes in product mix.

Responsiveness can be defined as the ability to apply knowledge, such as a market opportunity or a competitor's threat, effectively [Dove, 1999]. Knowledge is most valuable at the

time it is acquired and decreases in value from that time; thus it is important to be able to deploy and utilize knowledge quickly and effectively.

Naylor et al. [1999, pg. 108] define *leanness* as “developing a value stream to eliminate all waste, including time, and to ensure a level schedule.” While lean and agile supply chains focus on the reduction of waste and lead-time, a lean supply chain emphasizes smooth demand and level scheduling, while the agile supply chain stresses robustness and the ability to rapidly react to changes in market conditions. Additionally, the top priority in an agile supply chain is excellent customer service while a lean supply chain strives primarily to reduce costs.

Christopher and Towill [2002] outline the distinguishing characteristics of lean and agile supply chains, which are summarized in Table 1. The authors suggest that leanness and agility are not mutually exclusive. A hybrid strategy is suggested in which lean principles are applied to higher volume product lines with stable demand and agile principles are applied to more volatile product lines.

Table 1. Characteristics of Lean and Agile Supply Chains (adapted from Christopher and Towill, [2002]).

Distinguishing Attributes	Lean	Agile
Typical Products	Commodities	Fashion Goods, Semiconductors
Demand	Predictable	Volatile
Product Variety	Low	High
Product Life Cycle	Long	Short
Customer Drivers	Cost	Availability
Profit Margins	Low	High
Dominant Costs	Physical	Design, research and development
Stockout Penalties	Long-term contractual	High, often intangible
Information Enrichment	Desirable	Obligatory
Forecasting Mechanism	Algorithmic	Consultative

2.1.2. Factors that Influence Agility

This section explores some of the factors that contribute to an organization’s ability to increase agility. One such factor is the type of manufacturing environment, including the

production lead times. In a make-to-order environment, agility would involve having a flexible production system and cross-trained workers that can adjust quickly to new orders. In a make-to-stock environment with long lead times like the semiconductor industry, agility involves determining appropriate safety stock levels to manage variability, and filtering signals from noise in forecast to determine when to react by adjusting production and inventory.

The inventory policy, e.g., the amount, type, and location of stored inventory, is an important factor that determines whether customers' needs are met. Two inventory policies organizations could follow are to stock as much inventory of every type of product as possible, or to stock no inventory using a Just-in-time philosophy. The first policy would result in extremely high inventory storage costs, cycle time, and risk of product obsolescence, while the Just-in-time policy is unrealistic in industries like the semiconductor industry, where production lead-times are multiples of order lead-time. A balance between the two strategies that facilitates a low cost and high customer service level must be found.

Another factor that can affect the choice of the appropriate inventory strategy is transportation mode and speed. The aspects of transportation used to move raw materials, semi-finished, and finished goods throughout the supply network that are of primary concern are cost, speed, and consistency [Bowersox and Closs, 1996]. Typically, if faster transportation is used to ship finished goods to customers, inventory levels can be lower, but transportation cost could increase.

Unlike inventory and shipping speeds, some of the factors that contribute to agility cannot be directly controlled, such a product volume and variety, forecast accuracy, customer geographic dispersion, and order lead-time. Demand variability is perhaps the most important uncontrollable factor. As demand variability increases, more inventory is necessary to keep

service levels acceptable. Low volume products can often have sporadic demand, which is difficult to manage across a multi-echelon, globally dispersed supply network.

Forecasting is “the fundamental input to planning and coordinating logistical operations” [Bowersox and Closs, 1996]. It is the way an organization seeks to anticipate future uncertainty. The accuracy with which an organization can forecast the demand for its products affects its ability and need to be agile as well as the cost of being agile. Thus, forecast accuracy should be considered when determining the appropriate agility strategy.

Order lead-time supplied by customers also affects the need for agility and impacts other factors than influence agility. When order lead-time exceeds production lead-time, safety stock is minimized and forecasts are applied only to raw material inventory because goods can be made to order. As order lead-time decreases, forecast accuracy becomes more crucial, and organizations must be very agile to satisfy customer demands. All of these factors must be considered when determining an appropriate agile strategy for the products of interest.

2.1.3. Modeling Agility

There is a great deal of recent literature on supply agility, most of which takes a qualitative approach. Below are several papers that concentrate on modeling or measuring supply chain agility.

Mason-Jones and Towill [1999] focus on increasing agility by decreasing total supply chain cycle time, including both process and information lead-times, in order to create the information enriched supply chain. The authors hypothesize that cycle time can be reduced by cooperation and information sharing, and support their hypothesis with simulation models of both typical and “information enriched” supply chains in the fashion industry.

Towill and McCullen [1999] propose a strategy for improving agility by focusing on four principles: selection of appropriate control techniques, reduction of material and information flow lead-times to become less sensitive to forecast accuracy, sharing information between echelons of the supply chain, and elimination of echelons or interfaces wherever possible. This strategy is supported by a case study application, which shows a statistically significant reduction in the variability of quantities ordered and resulted in an increase in customer service levels. A time series analysis on data collected before and after implementing these agility strategies validates the improvement of the company's supply chain.

Swafford et al. [2001] present a model for global supply chain agility based on related constructs. Global supply chain agility is defined as a measure of the supply chain's ability to efficiently adapt to a rapidly changing global competitive environment to provide products and/or services. The goal of the research is to determine if an organization's global supply chain agility is defined by elements of flexibility, and if agility impacts competitive performance. First, several constructs are proposed as components of global supply chain agility. Next, the relationship between the construct of global supply chain agility and two dimensions of performance are examined, as well as the effect of the global competitive environment on global supply chain agility.

Ramasesh et al. [2001] develop a modeling framework for an agile manufacturing system, consisting of the supply sources, network or manufacturing facilities, and distribution outlets. The focus is on the assessment of system performance under change, the comparison of various system configurations with various agility levels, and financial justification of agility investments. In the framework, agility is linked to a set of aggregate performance measures.

Power and Sohal [2001] identify characteristics common to agile organizations by surveying companies that are given an agility rating. Independent variables are defined related to management style, computer-based technologies, resource management, and supplier involvement, among others. Responses include customer satisfaction, process changeover times, productivity, delivery performance, technological competitiveness, and product innovation. The most important subset of these variables is identified via factor analysis. Next, 1,000 Australian manufacturing companies, each of which is labeled “more agile” or “less agile,” answer a 246-question survey. The focus of the survey is to determine the degree to which the organization has achieved “best practice” in regard to the identified factors. Answers are given using a five-point Likert scale. Based on these answers, the relationship between the dependent and independent variables is modeled using multiple regression. Results show that agile organizations are more customer-focused, communicate more with their suppliers, and use technology to promote productivity more than the less agile organizations.

Goranson [2000] proposes a framework for measuring the structural agility of an organization based on its interaction with suppliers and customers. The framework, based on speech-act theory, uses two metrics of business communication: a distance metric and a time delay metric. While the distance metric is a function of the number of arcs connected to the node, the time delay metric and is the sum of the number of loops in the system. The author concludes that agility decreases with the number of nodes and loops in the system because complexity results in increased reaction time and difficulty in changing business processes.

Giachettia et al. [2003] propose a mathematical measurement framework for the flexibility and agility of an organization’s manufacturing system. Measurements of agility are based on four general dimensions: enriching the customer, cooperating to compete, organizing to

master change and uncertainty, and leveraging the impact of people and information. Some of the measures included are delivery flexibility, the ability to move planned delivery dates forward and accommodate special orders; and volume flexibility (the ability of a manufacturing system to be profitable at many output levels). The model also includes four indicators of cost, time, and scope proposed by Dove [1995]. All measurements are based on an interval scale.

Sarkis et al. [1995] describe a method for financially justifying agility by considering the strategic and long-term benefits. The methodology consists of five integrated phases: Identify System Impact, Identify Transition Impact, Estimate Costs and Benefits, Perform Decision Analysis, and Audit Decision. Various performance metrics are used to estimate the financial impacts of the selected strategy, while an activity-based costing approach is used to estimate its cost. These impacts are translated into financial criteria such as net present value or return on investment, and an alternatives comparison matrix is utilized to compare the various alternatives and their financial implications.

Shaw et al. [2002] present a methodology for measuring the impact of disturbances in the supply chain and benchmarking a firm's ability to react to these changes. Three categories of disturbances are considered: an upstream failure in the supply chain, a failure in the production system, and an unusual variation in demand. The goals of the methodology are to maximize customer service (measured by on-time deliveries) and minimize cost. The authors propose several variables and metrics, usually based on a ratio of two variables, to describe an organization's responsiveness to the three disturbances that are considered. The cost of this responsiveness is estimated by considering the cost of maintaining buffer stocks, a flexible workforce, and other factors.

Dubelaar et al. [2001] develop a regression model using retail data in order to quantify the relationships between inventory levels, service (availability) and sales. Data are collected both from a database of past orders and inventory levels and a survey of customers. Inventory is the response variable, while product variety, competition and demand uncertainty are independent factors studied. Results indicate that demand uncertainty is the most important of the factors studied in determining required safety stock. This work is similar to the work presented in this dissertation in that it explores the relationship between inventory and other independent factors. However, a major difference is that the availability of inventory in retail is the primary determinant of demand while in this research demand is independent of availability.

The recent supply chain agility literature is summarized in Table 3. Although there are many definitions and approaches for modeling agility, common themes among them are the ability to anticipate and respond quickly and effectively to change, improving customer service levels, and determining appropriate inventory levels. This research will address each of these aspects of agility.

The remainder of this chapter describes literature associated with each specific modeling phase of the research, identifying an ideal service level, and determining an effective inventory control policy that minimizes overall supply chain variability.

Table 2. Summary of Supply Chain Agility Literature

Paper	Modeling/Analysis Approach	Contribution
Ramasesh et al. [2001]	Linking agility to a set of performance measures	Methodology for assessing system's performance under change, comparing agility of various system configurations
Swafford et al. [2001]	Qualitatively model agility based on related constructs	Determines effects of global supply chain agility on performance
Mason-Jones and Towill [1999]	Simulation and case study of the fashion industry	Approach for decreasing total supply chain cycle time by information sharing and cooperation
Power and Sohal [2001]	Survey, factor analysis, regression	Determines characteristics of agile organizations, provide a basis for future research on agility
Dubelaar et al. [2001]	Regression, customer surveys	Determine relationship between inventory and independent variables
Sarkis et al. [1995]	Estimate Net Present Value and Return on Investment of agility alternatives	Framework for financial justification of agility
Shaw et al. [2002]	Interviews, development of graphical tool to identify strengths and weaknesses, case studies	Approach for the measurement and benchmarking of response capabilities
Goranson [2000]	Speech-act theory	Framework for measuring the structural agility based on interaction with suppliers and customers
Giachettia et. al [2003]	Develop ordinal scale to represent factors identified in literature review	Framework for measuring the flexibility and agility of a manufacturing system

2.2 Service Level Models

While the previous section discusses supply chain agility in a broad sense, this section is devoted to examining literature relating to one of the major goals of agility, achieving a high customer service level, in greater detail. A summary of inventory models that include customer service level in the objective function or constraints is presented. Some of these models have an objective to minimize costs (inventory) in the supply chain, subject to service level constraints [Ettl et al., 2000; Liu et al., 2004; Inderfurth and Minner, 1998], while others maximize service level with cost or inventory constraints [Schwarz et al., 1985; Lagodimos, 1992; De Kok and Verrijdt, 1995]. A few select models are described in more detail in the remainder of this section.

Ettl et al. [2000] present a supply chain model based on a bill of material (BOM). The objective is to minimize the total average dollar value of inventory in the supply chain as an

objective function, subject to service level constraints, in order to determine safety stock in several supply chain locations. Non-stationary demand is considered. No justification is given for why this particular service level constraint is selected.

Schwarz et al. [1985] present a model for determining a continuous review inventory policy for a one-warehouse, N -retailer distribution system. The objective is to maximize customer service level, and although supply is considered to be unlimited, the model is subjected to an inventory (cost) constraint for which no explanation is provided. Stationary demand with known mean and standard deviation is assumed.

Sabri and Beamon [2000] present a supply chain model for simultaneous strategic and operational planning that includes both multiple echelons as well as multiple supply chain objectives. The strategic sub-model's goal is to optimize the supply chain configuration and material flow by minimizing cost, while ensuring a sufficient amount of volume flexibility, subject to capacity constraints and customer demand. The operational sub-model takes inputs from the strategic model's solution and optimizes inventory lot sizes, reorder points, and safety stock. The objective function incorporates the tradeoff between cost, customer service level, and flexibility by utilizing weights representing the relative importance of these factors, which are determined by the decision-maker. The shortcoming of this paper is that the approach's success is based heavily on the decision-maker's choice of weights, and therefore, it is unlikely to achieve the most cost-effective service level.

The three papers described above represent the majority of the literature in their treatment of customer service level. In these models, the service level or cost constraints are either determined arbitrarily and *a priori*, or constraint parameters are left as variables for the decision-maker to determine. Even with using sensitivity analysis on these parameters, this approach is

unlikely to result in the most cost-effective service level. Additionally, none of the service level models described thus far consider the cost of stock-outs (when the specific product requested by a customer is not available at the requested time). This cost is equally as important as the cost of holding inventory in determining an agile strategy and customer service level.

Sonnet [2004] presents a methodology for estimating the cost of stock-out for semiconductor products. Customer surveys are utilized in order to estimate a customer's likelihood to postpone the sale, cancel the sale, or buy a different product from the same company when a stock-out occurs. This is the only work, to the author's knowledge, that quantifies the cost of stock-outs using actual market data.

With the exception of Sonnet, [2004], there has been no quantifiable basis for determining the cost of lost sales due to less than perfect customer service. This work builds upon Sonnet's research in order to find the most cost-effective inventory and customer service level for specific products by trading off the cost of holding inventory with stock-out costs.

This section has discussed literature that addresses customer service level, an important goal of supply chain agility. In the next section, literature that addresses another aspect of agility, the ability to detect and respond effectively to change, is presented.

2.3 Methods for Detecting/Reacting to Change

Achieving supply chain agility requires not only setting the right inventory levels initially, but also continually monitoring the market to detect and quickly react to changes in factors such as demand and competition. This section addresses another important component of agility, determining when and how to react to change, which can be extremely difficult in a

highly variable environment because noise can both disguise true demand shifts and falsely identify a demand shift.

Inventory levels must be evaluated on a regular basis to determine if the dynamic factors upon which they were based initially have changed significantly to merit new inventory targets. Due to long lead-times, it is important to react to change as quickly as possible. However, reacting frequently has two potential pitfalls: (1) a great deal of resources must be devoted to monitor and identify changes and (2) reacting too frequently can increase unnecessary changes of internal processes such as manufacturing, and the effects are amplified due to the bullwhip effect. Therefore, mild changes in inventory targets can result in extreme variability in the production echelon. Forrester [1969] describes potential consequences of the bullwhip effect as high and fluctuating inventory levels, longer lead-times, expensive under- or over-utilization of resources (i.e. overtime and extra machinery), extra set-up and changeover time, and even quality problems. Therefore, it is desirable to mitigate these effects by responding only when necessary, i.e., when true shifts in the underlying mean demand have occurred. The remainder of this section discusses a few papers that address methods for identifying and reacting effectively to true shifts while minimizing supply chain variability by filtering out noise whenever possible. The papers are grouped into system dynamics and control theory approaches and control chart approaches.

2.3.1. System Dynamics and Control Theory

There have been many applications of system dynamics and control theory, two closely related concepts, to demand and inventory management. This section defines both concepts and discusses several applications of each in the field of supply chain management.

System dynamics, originally presented by Jay Forrester, is a methodology for analyzing complex feedback systems such as a supply chain. System dynamics studies a system holistically as a feedback system rather than each component individually. The central concept of system dynamics is that all systems interact through causal relationships. Conservation of flow equations are typically used to represent interactions between each node of the system [Forrester, 1969].

Towill [1993] explains the concept of system dynamics and its application to supply chain management. He shows an example of smoothing a material flow in a supply chain by using all the information available in the market place rather than reacting to the distorted information passed on by the adjacent echelon. The bullwhip effect is minimized and the entire supply chain is more stable.

Lertpattarapong [2002] presents a model for detecting changes that occur in a supply chain, including demand shifts. The model incorporates neural networks and system dynamics, and involves the development of a causal loop diagram to explain the dynamic behavior of the supply chain under study. Seventeen independent variables are analyzed in order to determine the cause of oscillations taking place in finished goods inventory and required capacity. By examining these variables, the author finds the causes of the oscillations.

Optimal control theory is “a branch of mathematics developed to find optimal ways to control a dynamic system” [Sethi and Thompson, 2000]. Control theory has been applied to the control of inventory systems with the goal of reducing demand amplification due to the bullwhip effect. A few of the control theory applications to inventory management are described below.

Braun et al. [2003] apply Model Predictive Control (MPC), a methodology for inventory control that utilizes quadratic programming algorithms and control theory, to a semiconductor supply chain. The MPC controller takes current inventory levels and forecasts as inputs and

outputs factory production. It is programmed to smooth reactions to perceived demand shifts to minimize the bullwhip effect. The degree of smoothing can be controlled via a smoothing constant. The methodology is first applied to a single-product, two-node problem controlled by a predictive controller using anticipation. Inaccurate forecasting and shifts in demand are used to test the performance of the controller, which proves to be robust despite these conditions. Insights gained from this investigation are then applied to the design of a centralized MPC controller for a four-node problem, and many other extensions to this work follow [Wang et al., 2003; Wang et al., 2004, Rivera et al., 2005]. While the approach is promising, it is complicated, and an extensive knowledge of higher mathematics, programming, and model predictive control are needed to implement it, as well as expertise on the supply chain under study. Additionally, these papers do not describe the impact of the smoothing factor on the ability to react to true shifts in demand.

Although many of these works show promise for mitigating the effects of variability in the supply chain, the practicality of applying optimal control theory by the majority of planning personnel is questionable. A methodology that is easier to understand and implement is desirable.

2.3.2. Control Charts

Control charts have been used extensively for statistical quality control to detect when a process is out of control, or behaving abnormally. A control chart is a graphical display of a characteristic of one or more variables. The chart consists of a center line representing the mean of the characteristic assuming the process is in control, and two horizontal lines called upper and lower control limits (UCL and LCL). These control limits are a function of the variability of the process, and typically are chosen so that if the process is in control, most of the sample points

will fall between them. Typical control limits are placed three standard deviations above and below the mean, but may differ based on the desired probability of a sampling statistic falling outside of the control limits if the process is in control, which depends on the consequences of making a Type I or Type II error [Mitra, 1998].

Many types of control charts are described in the literature. Control charts for attributes track characteristics that cannot be measured numerically and do not indicate the degree to which a process is non-conforming, while control charts for variables are used for tracking characteristics that are measurable on a numerical scale (such as demand and inventory levels). Montgomery [1985] refers to variable control charts as *leading indicators* of trouble that can alert users to a quality problem before it becomes too serious, i.e., too many points are out of control.

Control of the process average or mean quality level is usually with the control chart for means, or the X-bar chart, while control of the process range is usually tracked with an R-chart. The exponentially weighted moving average (EWMA) control chart is considered to be the best type of control chart for detecting small changes in the mean of a process quickly [Mitra, 1998]. EWMA control charts use a weighting constant, λ , which determines the amount of weight given to past and current information.

Roberts [1959] first describes the use of the exponentially weighted moving average (EWMA) for control purposes. Several other applications follow [Lucas and Saccucci, 1990], including a comparison of the EWMA to other control chart techniques [Roberts, 1966].

Despite the many applications of control charts to statistical quality control, the literature contains few applications of control charts to inventory management. Buzacott [1999] presents one such paper, a methodology for a periodic review inventory system where forecasts are

generated from a simple moving average of past demand. An ABC system is used, with A items replenished weekly, B items every two weeks, and C items every 8 weeks. The research shows that dynamic inventory targets are preferred to stationary targets in the presence of abnormal demand activity, i.e., when demand becomes out of control.

Takahashi and Nakamura [1999] apply an exponentially weighted moving average control chart to the problem of minimizing the variability of a Just-in-Time (JIT) manufacturing environment. The JIT controller is used to detect unstable demand, or deviations from normal process variability. When demand falls outside of the control limits, buffer inventory sizes are readjusted by an amount that is determined by simulating the system under conditions of stable demand. The main shortcoming of the paper is that the methodology is applied to a simple two-echelon supply chain with short lead-times and it is unclear whether buffer size could be adjusted before a stock-out occurs in a more complex supply chain. In addition to the complexity of the supply chain, this research is different than the prior work because this research involves making adjustments to inventory levels and production more frequently when a forecast is determined to be out of control until the process stabilizes rather than making a single adjustment to inventory buffers.

2.4 Summary of Relevant Literature and Contribution of this Study

There is a great deal of literature on supply chain agility, but agility is generally dealt with qualitatively and measured subjectively. Two major goals of supply chain agility are prevalent in the literature: improving customer service levels and reacting effectively to change. However, most inventory models that consider service level choose a target arbitrarily, or maximize service level with arbitrary cost constraints, an approach which will not generally lead

to finding the most cost-effective customer service and inventory levels. There is a great deal of control and systems theory literature on minimizing the variability in inventory systems, but it is extremely complicated and its practicality is questionable. Finally, control charts have been shown to be very effective for statistical quality control, and show potential for effective control of inventory systems, although there have been few applications in this area.

This research fills these gaps in the literature by addressing two major aspects of supply chain agility that are necessities in complex environments like the semiconductor industry – the ability to achieve a cost-effective customer service level by storing the right amount of inventory for specific products, and the ability to manage variability by reacting effectively to change. The expected contribution of this research is a methodology for determining ideal customer-service and inventory levels; as well a production and inventory method that results in the ability to react effectively, avoid chasing noise, and minimize supply chain variability.

In the next two chapters, the methodology for achieving the stated objectives is described.

CHAPTER 3

RELATIONSHIP BETWEEN INVENTORY, CUSTOMER SERVICE LEVEL, AND COST

This research focuses on modeling supply chain agility and its associated costs. This chapter describes quantifying the relationship between inventory, customer service level, and cost. Variables related to agility are identified, data are collected related to these variables, and mathematical relationships for customer service level are developed using regression modeling. An equation for the cost of providing a specified customer service level is also developed. The primary goal of this phase is to use the customer service level and cost equations to determine a cost-effective customer service level and the finished goods inventory required to achieve it. A secondary goal is to explore the impact of uncontrollable factors on the customer service level and inventory relationship. The inventory levels determined in this chapter are used as a starting point in the second phase, in which several policies for updating inventory targets are evaluated to determine which policies can effectively react to actual demand shifts without reacting to noise.

The remainder of Chapter 3 is devoted to describing the modeling methodology and the case study application.

3.1 Introduction to Case Study

Intel Corporation, the world's largest semiconductor manufacturer, is selected for the case study application of this research. Two of Intel's product families, similar in nature but with different sales volumes and amounts of variability in demand, are selected for the analysis. Two product families are selected so that the results can be compared for validation purposes as well

as to understand how fundamental differences between the two groups affect the ideal inventory levels. Both product groups contain products in the mature stage of the product life cycle with highly variable demand (with a coefficient of variation greater than 0.4). For both product groups, historical data is available for more than one year, and experts testify that the available data is reasonably reliable. The problems associated with managing supply and demand for these products, as described by the experts who manage them, include short lead-times requested by customers coupled with long production lead-times, frequent and last-minute order cancellations and changes, high costs of holding inventory, and difficulty of accurately forecasting demand.

The remaining sections in this chapter describe the approach for quantifying the relationship between inventory, customer service level, and cost in detail.

3.2 Regression Modeling

Logistic regression modeling is employed to develop an equation to represent on-time delivery performance to customers, one of the primary goals of agility. Thousands of historical data points are collected describing inventory levels, forecast accuracy, order lead-time, and variability of demand over a period of one year. Logistic regression models are built for each of the two product families using customer service level, a binary variable indicating whether an order is late or on time, as a response. Models are constructed using SAS Version 8.2. Goodness of the model fit is measured using logistic regression model diagnostics including Max Rescaled R-Square (a measure equivalent to R-Square in ordinary least squares regression), Akaike's Information Criterion (a measure of model error), and the Hosmer-Lemeshow lack of fit test (a measure of the overall model's ability to predict the response) [Hosmer and Lemeshow, 2000]. Two types of models are built for each product family, a "planning" model that includes only the

inventory variable as an independent variable, and “insight” models that include both inventory and one of three uncontrollable factors (forecast error, demand variability, and order lead-time) as independent variables. While the purpose of the planning model is to determine the relationship between customer service level and inventory, the purpose of the insight models is to determine the influence of three important factors on this relationship. Each class of models is discussed in more detail in the following sub-sections.

3.2.1. Defining the Customer Service Level and Inventory Relationship

The planning model is developed to understand the relationship between inventory levels and the customer service level for the historical performance of the product groups. The inventory level is typically correlated with the demand forecast. This can be seen from both the Economic Order Quantity formula and from common industry practice of targeting a specified number of weeks of anticipated demand in inventory when demand is non-stationary. The relationship between inventory and demand and demand and delivery performance (i.e. as the demand increases while inventory stays the same, delivery performance decreases) adversely affects the models ability to relate inventory to delivery performance. Therefore, instead of using inventory level directly, the inventory levels are scaled by the forecasted demand, creating the independent variable, “weeks of inventory.” Thus, the regression planning model had a single independent variable, “weeks of inventory” (*WOI*), with delivery performance as a response. *WOI* can be defined as follows:

$$WOI = \frac{\text{inventory on hand}}{(\text{Demand forecast for next 13 weeks} / 13)} \quad (1)$$

WOI is the inventory in units divided by the average weekly forecasted demand for the next quarter. While the target multiple of *WOI* generally remains fairly constant, the forecast in units changes regularly.

The regression equation is used to plot the expected customer service level versus *WOI* to determine the inventory required to achieve a desired service level. Confidence limits are also plotted by adding the standard error in the variables' coefficients determined during regression multiplied by 1.96 (the Z-score corresponding to $\alpha=.05$). Figure 1 shows the relationship between inventory and expected service level for Product Group 1, while Figure 2 shows a nearly identical relationship for Product Group 2. However, a slightly (less than one percent) higher service level is attained for Product Group 1 when storing the same amount of inventory as Product Group 2. This is expected because Product Group 2 has a higher average coefficient of variation of monthly demand for its products (0.61 compared to 0.48) during the time period under study.




Figure 1. Weeks of Inventory vs. Expected Service Level with 95 Percent Confidence (Product Group 1)

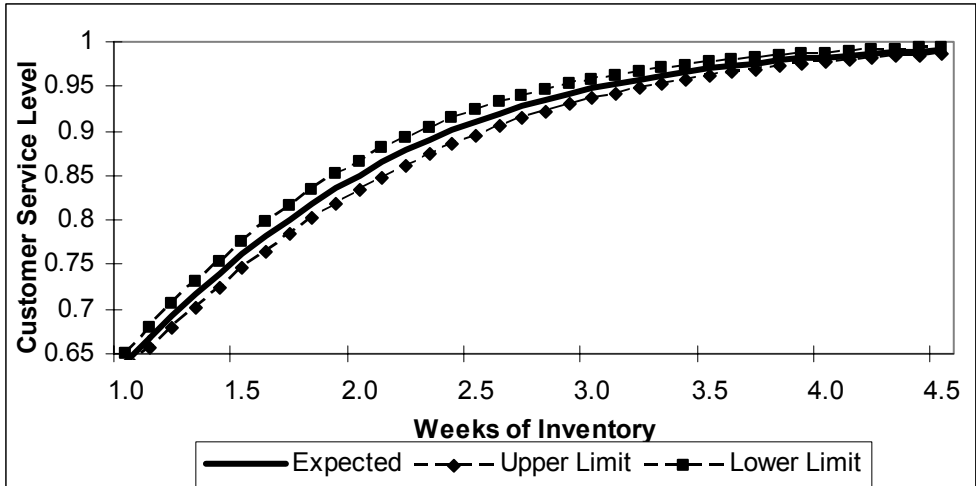


Figure 2. Weeks of Inventory vs. Expected Service Level with 95 Percent Confidence (Product Group 2)

3.2.2. Analysis of Factors Affecting the Service Level and Inventory Relationship

This section is devoted to the analysis of the effects of three additional factors on the relationships shown in the previous section. Order lead-time, forecast accuracy, and demand variability are paired with the “weeks of inventory” variable and customer service level as the response in three logistic regression insight models. These equations are developed in order to estimate the effect these variables have on the inventory and customer service level relationship. Although these factors cannot be controlled directly, this analysis provides insight into the appropriate reaction to changes in these factors. This allows organizations to prepare in advance of the change, and to develop strategies in advance for worst-case scenarios. By ensuring a plan is in place for such scenarios, organizations will be better prepared to react to change and remain agile when changes occur. Additionally, it may not be necessary to perform regression analysis for each of the products manufactured by the company. Rather, regression can be done on clusters of products and adjustments can be made for individual products based on the values of these three variables.

Figure 3 shows a graph of weeks of inventory and the requested order lead-time (the time from when an order is placed until it must be received by the customer to be on-time) required to achieve a 95 percent service level. This graph is based on a logistic regression model with two independent variables: weeks of inventory and the square of order lead-time. The 95 percent confidence limits are determined based on the standard error associated with each parameter estimate.




Figure 3. Relationship between order lead-time and inventory required to achieve a 95 percent service level with 95 percent confidence limits.

Figure 3 shows the inventory scaled by the forecasted demand (shown on the primary x-axis) required to achieve a 95 percent service level if order lead-time is equal to the quantity on the y-axis. Confidence levels based on an alpha of .05 are shown on the order lead-time. Based on the figure, if ten weeks of finished inventory are stored, the organization can expect to provide a service time to the customer of about five days. If only one week of finished inventory is held, an average of 30 to 40 days is needed to complete a customer's order. The figure assumes that the value of order lead-time is the same for every order. However, in practice order lead-time varies. However, the information in the graph can provide a general indication of the

value/cost of increasing/decreasing order lead-time, and this variable's effect on the inventory and service level relationship.

Figure 4 shows a graph of the inventory required to achieve a 95 percent service level versus forecast error, as measured by:

$$\text{Forecast Error} = \left| \frac{\text{Forecasted Demand} - \text{Actual Demand}}{\text{Actual Demand}} \right| \quad (2)$$

Forecast error is measured based on a forecast made one month in advance of requested delivery dates of orders.



Figure 4. Relationship between forecast error and inventory required to achieve a 95 percent service level with 95 percent confidence limits.

Based on Figure 4, weeks of inventory could be decreased from 5 to 3.6 (28 percent) while achieving the same service level if forecast error could be reduced from 30 percent to 20 percent. Presumably, if improvements could be made in forecast accuracy for this particular product, forecast accuracy could be improved for other products as well and the potential savings could be enormous. With this information, companies can determine the amount of resources to dedicate to developing better forecasting techniques.

Finally, Figure 5 shows the relationship between the coefficient of variation of demand (measured over a six-month period) and inventory required to achieve a 95 percent service level. The model contains weeks of inventory and the log of coefficient of variation of demand as its two independent variables.



Figure 5. Relationship between coefficient of variation of demand and inventory required to achieve a 95 percent service level with 95 percent confidence limits.

The amount of inventory required to achieve the same service level increases as the variability of demand increases because safety stock increases with uncertainty. Therefore, it is desirable to reduce demand variability as much as possible. Although it is very difficult to have any effect on demand variability, it may be possible to influence demand based on price incentives or long-term contracts.

In summary, these three variables that have a significant impact on supply chain agility and the relationship between inventory and customer service level have been analyzed. The graphs in this section provide insight into the effects of these factors so that when changes occur, the impact is known in advance. Organizations can easily understand the impact on required inventory levels if one of these factors were to suddenly change and be prepared for the change.

Although the effects of these uncontrollable variables have been shown to be significant, they may have been underestimated due to the technique used. For example, forecast error is not known in advance and, therefore, when the forecast is highly accurate inventory cannot be reduced as much as it could be if the forecast error was known in advance. This relationship is not reflected in the data and therefore is not considered in the analysis and figures. Therefore, it can be concluded that regression analysis cannot determine the full effect of the independent variables on the inventory required to achieve a given customer-service level.

In order to ensure the validity of the regression models presented in this chapter, validation of model fit is performed by partitioning the data into training (70 percent) and validation (30 percent) data sets. The model is initially fit using the training data set, and later fit using the validation data set for comparison. Table 3 summarizes the validation results.

Table 3. Model Validation Results

Product Group	Variables	R-sq	Validation R-sq
1	Weeks of Inventory	0.544	0.482
1	Weeks of Inventory & Order Lead-time	0.645	0.603
1	Weeks of Inventory & Forecast Error	0.571	0.527
1	Weeks of Inventory & Demand Variability	0.596	0.526
2	Weeks of Inventory	0.512	0.477
2	Weeks of Inventory & Order Lead-time	0.653	0.607
2	Weeks of Inventory & Forecast Error	0.568	0.522
2	Weeks of Inventory & Demand Variability	0.579	0.492

The validation scores are reasonably close to the initial R-squared values to conclude that over fitting of the training data has not significantly impacted the models and to proceed.

3.3 Customer Service Level Cost Equation

The next step in finding a cost-effective customer service level is the development of an equation quantifying the costs associated with achieving a given service level. These costs include inventory holding costs and the cost of lost sales due to stock-outs. In the supply chain

under study, the yearly values for the holding cost for finished goods inventory are estimated as 26 percentage of the variable cost to produce these products, which includes obsolescence (the decrease in the value of products from the time they are manufactured until sold), opportunity cost, and scrap [Bridge, 2004]. Therefore, the inventory holding cost (*IHC*) in dollars for a specific inventory level in units is as follows:

$$IHC = 0.26 * \text{Variable Cost of Production} * \text{Inventory [units]} \quad (3)$$

Inventory holding costs can also be defined based on weeks of inventory (*WOI*) as:

$$IHC = 0.26 * \text{Variable Cost of Production} * \text{WOI} * \text{WOI Target Multiplier} \quad (4)$$

Data such as variable cost and demand cannot be disclosed due to confidentiality issues, but is readily available, so the inventory holding cost for each service level can be easily calculated.

The cost of loss of business due to less than perfect delivery performance is estimated based on the survey-based method proposed by Sonnet [2004]. Customers of the specific products under study are asked to indicate their likelihood to wait for products when a stock-out occurs (the specific product desired is not available at the date requested). Results of the survey indicate that if the desired product is unavailable at the requested time, customers purchase an alternative product from the same company or wait for the desired product approximately 80 percent of the time, and buy from the competition the remaining 20 percent of the time. The time value of money is not considered, and it is assumed that delivery performance for a particular product does not affect customers' willingness to buy other products from the same company.

The cost to achieve a specific customer service level, $C(SL)$, is as follows:

$$C(SL) = \frac{\text{Inventory} * \text{Inventory Holding Cost} + \text{Expected Lost Sales} * \text{Profit Margin}}{\text{Sales}} \quad (5)$$

Using Sonnet's findings and (4), the cost of achieving a customer service level can also be written as:

$$C(SL) = WOI * WOI \text{ Target} * 0.26 * \text{Variable Cost of Production Multiplier} + 0.2 * (1-SL) * \text{Average Weekly Demand} * \text{Profit Margin} \quad (6)$$

Some of the components in the service level cost equation are estimates. Therefore, each component is scrutinized under sensitivity analysis. The effect of the profit margin (product cost as a percent of revenue) is shown in Figure 6.




Figure 6. Sensitivity Analysis on the Effect of Product Margin on Ideal Service Level

In order to find the ideal customer service level, it is not necessary to know the exact values of revenue and the cost to produce a product, but the ratio of these two quantities must be known. Although the results are sensitive to this ratio, the cost and revenue of products will be known with near certainty (bulk discounts may be given so revenue may vary slightly and exact costs may be difficult to predict), so this graph should be used mainly to determine the inventory level for products with various profit margins rather than for sensitivity analysis. It is important to note that due to the sensitivity of the results to this ratio, it can be expected that some products with high margins may have very different ideal inventory targets than those with lower margins. It must also be noted that inventory targets must be reevaluated frequently because margins, as well as other factors inventory targets are based on, change quickly for products with short life cycles such as those manufactured in the semiconductor industry.

3.4 Determination of Minimum Cost Service Level

After equations for customer service level and its cost are obtained, a simple optimization is performed and the required inventory as a function of forecasted demand to achieve the minimum cost service level is estimated. Plots are generated for both product groups containing the expected customer service level for a given finished-goods inventory level with upper and lower confidence limits, along with the total cost of achieving the service level based on Equation (6). Actual costs are not shown due to confidentiality reasons.




Figure 7. Cost and Service Level versus Weeks of Inventory (Product Group 1)

In Figures 7 and 8, the minimum cost service level is found by locating the minimum range on the cost curve, projecting upward to the service level curves, and finding the corresponding service level on the primary y -axis. Figure 7 shows that the cost is minimized by storing 5.1 weeks of finished inventory, which will result in a service level of 95.68 percent plus or minus 0.44 percent. However, the range from 4.7 to 5.6 weeks of inventory results in a cost

that is within one percent of the minimum cost and a customer service level of 94.57 to 96.78 percent.




Figure 8. Cost and Service Level versus Weeks of Inventory (Product Group 2)

Figure 8 shows that for Product Group 2, 5.25 weeks of inventory resulting in a service level of 95.56 percent is ideal. However, a range of 5.1 to 5.7 weeks of finished inventory with a service level range of 94.73 to 97.04 percent deviates less than one percent from the minimum cost. The reason that more inventory is required by Product Group 2 to achieve virtually the same customer service level as Product Group 1 is likely due to the fact that Product Group 2 has higher variability of demand. It can be concluded based on the cost curve in both figures that it is more desirable to exceed inventory targets than to fall short of them because the cost of stock-out is greater than the cost of holding excess inventory.

3.5 Conclusions

To summarize the chapter, a methodology is presented for determining the relationship between inventory, customer service level, and cost. Logistic regression analysis is performed on

historical inventory and delivery performance data from Intel Corporation to formalize the relationship between inventory and customer service level. A cost of achieving a given customer service level is obtained via financial analysis of inventory holding costs and stock-out costs. This information is used to identify a minimum cost customer service level and the inventory required to support it for specific products. Example results from two product groups manufactured by the case organization are presented. The impact of three uncontrollable factors (forecast error, demand variability, and order lead-time) on the relationship between inventory and customer service level is evaluated in order to understand the necessary adjustments to target inventory and services levels should one or more of these factors change.

The inventory levels determined in this chapter are based on a demand forecast that changes over time. A method for determining how often to forecast and when to change these inventory targets and production based on new information is necessary because an organization must react effectively to change in order to be agile. In Chapter 4, frequency and event-based policies for production and inventory are examined via a Monte Carlo Simulation.

CHAPTER 4

INVENTORY CONTROL POLICIES

The methodology for determining inventory targets developed in Chapter 3 requires determining inventory levels based on a multiple of forecasted demand. Since the demand forecast changes over time and inventory targets are based on the demand forecast, the inventory targets must change over time as well. The frequency with which inventory targets are updated determines how quickly organizations can react to demand changes, but also impacts the variability of the entire supply chain. More frequent updating may mean more production variability in the supply chain. However, updating less frequently may mean being slow to react to change (i.e., less agile) and may lead to stock-outs or excessive inventory and scrap. This chapter focuses on determining the frequency at which to forecast demand and the policy for identifying when to react by updating inventory targets and/or production. This is done by comparing the results from a Monte Carlo simulation for different inventory policies. The goals are to determine a policy that results in a suitable trade-off between minimizing cost and variability, as well as to determine a policy that is robust, meaning it performs reasonably well for all products.

In this chapter, several control chart-based planning policies are proposed and described. A description of the simulation model used to evaluate the various planning policies for the supply chain under study is provided. Section 4.3 describes the method for generating forecasts used by the simulation model and assumptions used. Results of the simulation runs are presented in the form of statistical hypothesis tests to compare the various policies in terms of cost and variability. Finally, a summary is presented.

4.1 Inventory Planning Policies

Several policies for production and inventory control are described in this section. The policies are evaluated by the supply chain model that is presented in the next section. The first set of policies to be evaluated is the frequency-based planning policies. For these policies the forecasts are updated weekly (7 days), biweekly (14 days), monthly (30 days), and quarterly (91 days). These planning policies involve forecasting customer demand at a given frequency and subsequently updating inventory targets and production levels each time a forecast is made. The current practice at the case company is to update forecasts and production schedules monthly for these two product lines.

Additionally, several hybrid policies that incorporate frequency and event-based planning are also evaluated. These policies call for forecasting and updating production levels more frequently when a forecast becomes out of control, i.e. when the forecast falls outside of the control limits. Due to its ability to quickly detect changes to the mean of the quantity of interest (in this case, demand), the exponentially weighted moving average (EWMA) control chart is used as a tool for determining when an out of control demand event has occurred. The EWMA is calculated as follows:

$$EWMA_t = \lambda \bar{X}_t + (1 - \lambda)^2 EWMA_{t-1} \quad (7)$$

The weighting factor, λ , determines the emphasis placed on more recent versus less recent observations. For each control chart-based policy, λ is varied at 0.1 and 0.3 because the typical range of values is from 0.1 to 0.3 [Mitra, 1998]. Additionally, the control limits are set at 1, 1.5, 2, and 3 multiples of the standard deviation from the mean. When the forecast becomes out of

control, forecasting is done more often and for shorter durations, thus the production levels are updated more frequently until the forecast returns to an in control state. For example, one such policy is to update the forecast monthly when the forecast is in control, and weekly when out of control. This policy is paired with the two settings for λ and the four settings for the width of the control limits. The resulting 28 policies that are evaluated by the simulation model are shown in Table 4.

Table 4. Planning Policies Evaluated by Supply Chain Model

Policy	Basis	In Control Frequency	Out of Control Frequency	Width of Forecast Control Limit	Control Chart λ
1	Frequency	Quarterly	N/A	N/A	N/A
2	Frequency	Monthly	N/A	N/A	N/A
3	Frequency	Biweekly	N/A	N/A	N/A
4	Frequency	Weekly	N/A	N/A	N/A
5	EWMA Control Chart	Monthly	Biweekly	1σ	0.1
6	EWMA Control Chart	Monthly	Biweekly	1σ	0.3
7	EWMA Control Chart	Monthly	Biweekly	1.5σ	0.1
8	EWMA Control Chart	Monthly	Biweekly	1.5σ	0.3
9	EWMA Control Chart	Monthly	Biweekly	2σ	0.1
10	EWMA Control Chart	Monthly	Biweekly	2σ	0.3
11	EWMA Control Chart	Monthly	Biweekly	3σ	0.1
12	EWMA Control Chart	Monthly	Biweekly	3σ	0.3
13	EWMA Control Chart	Monthly	Weekly	1σ	0.1
14	EWMA Control Chart	Monthly	Weekly	1σ	0.3
15	EWMA Control Chart	Monthly	Weekly	1.5σ	0.1
16	EWMA Control Chart	Monthly	Weekly	1.5σ	0.3
17	EWMA Control Chart	Monthly	Weekly	2σ	0.1
18	EWMA Control Chart	Monthly	Weekly	2σ	0.3
19	EWMA Control Chart	Monthly	Weekly	3σ	0.1
20	EWMA Control Chart	Monthly	Weekly	3σ	0.3
21	EWMA Control Chart	Biweekly	Weekly	1σ	0.1
22	EWMA Control Chart	Biweekly	Weekly	1σ	0.3
23	EWMA Control Chart	Biweekly	Weekly	1.5σ	0.1
24	EWMA Control Chart	Biweekly	Weekly	1.5σ	0.3
25	EWMA Control Chart	Biweekly	Weekly	2σ	0.1
26	EWMA Control Chart	Biweekly	Weekly	2σ	0.3
27	EWMA Control Chart	Biweekly	Weekly	3σ	0.1
28	EWMA Control Chart	Biweekly	Weekly	3σ	0.3

The model is run for a period of one year (365 time periods of one day each) for 10,000 replications per planning policy for each of 18 products from two product groups studied. The

following statistics are collected for each planning policy: number of stock-outs, average inventory level, variability of production in the factory and assembly/test nodes (measured by the sum of the variance of the amount of material that begins production in these two nodes each day), and the cost of stock-outs and holding inventory (the total cost of achieving a customer service level discussed in Chapter 3).

The supply chain model and the generation of forecasts described later in this chapter are programmed using the PERL programming language. The complete code can be found in Appendix B. The results are described and policies are compared following the detailed description of the supply chain model and forecast generation method.

4.2 Monte Carlo Supply Chain Simulation Model

In order to evaluate the impact of various inventory control policies, a four-echelon Monte Carlo simulation model of the supply chain under study is developed. Figure 9 depicts a simplified version of the case company's supply chain.




Figure 9. Case Supply Chain Node Diagram.

The factory node is where wafers are produced with an average 10-week throughput time, which is assumed to be constant. Production levels are determined by demand forecasts and inventory levels. After manufacturing, inventory is stored in the semi-finished inventory warehouse, Node 2, until being sent to assembly/test for processing. The quantity of inventory sent to assembly/test is based on a forecast and the finished-inventory level. Assembly/test is modeled as a constant two-week processing time. After assembly/test, inventory is stored in finished inventory, Node 3, until requested by a customer. The variables included in the model are described below.

The following indices are used in the model:

i	product index (1..18)
t	time index, measured in days (1..365)
m	planning method (1..28)
f	frequency of forecasting (1..4, where 1:quarterly, 2: monthly, 3: biweekly, 4: weekly)

The following variables indicate production levels:

$Prod_{i,t,m}$	units of product i started in the fab during period t based on method m
$SItOFI_{i,t,m}$	units of product i shipped from semi-finished inventory to finished inventory through assembly/test during period t based on method m

The following are inventory variables:

$SI_{i,t,m}$	units of product i in semi-finished inventory at the beginning of period t based on method m
$FI_{i,t,m}$	units of product i in finished inventory at the beginning of period t based on method m

The following variables determine the inventory targets:

$FIT_{i,t,m}$	target units of finished inventory for product i in period t based on method m
$SIT_{i,t,m}$	target units of semi-finished inventory for product i in period t based on method m

l_f length of the period between forecasts corresponding to the forecasting frequency f
 $F_{i,t,m}$ daily demand forecast for product i in time t based on method m , which determines the frequency f (the forecast for the period corresponding frequency f divided by l_f)

The following variables are used to calculate units of product shipped by period:

$A_{i,t}$ actual demand for product i during period t
 $B_{i,t,m}$ back-orders for product i during period t based on method m
 $Ship_{i,t,m}$ units of product i shipped from the finished inventory warehouse to customers during period t based on method m

The following variables are measures of performance tracked to compare the planning scenarios:

$MD_{i,t,m}$ units of demand for product i in period t that is lost because customers are unwilling to wait based on method m
 $S_{i,t,m}$ number of stock-outs for product i in period t (total number of days when all demand was not met on time) based on method m

The following are input parameters determined in the previous analysis in Chapter 3:

SM_i target multiple of semi-finished inventory in relation to one week of forecasted demand (week of inventory) for product i
 FM_i target multiple of finished inventory in relation to one week of forecasted demand (week of inventory) for product i
 ρ percent of unsatisfied demand that becomes back-ordered in the next period (probability that a customer is willing to wait for an order)

The inventory targets are calculated based on demand forecasts for the next quarter and the target multiple of semi-finished or finished inventory in relation to one week of forecasted demand determined in Chapter 3 as follows:

$$SIT_{i,t,m} = \begin{cases} SM_i * 7 * average(F_{i,t,m} : F_{i,t+91,m}) & \forall i, m, t = 1 \dots 273 \\ SM_i * 7 * average(F_{i,274,m} : F_{i,365,m}) & \forall i, m, t = 274 \dots 365 \end{cases} \quad (8)$$

$$FIT_{i,t,m} = \begin{cases} FM_i * 7 * average(F_{i,t,m} : F_{i,t+91,m}) & \forall i, m, t = 1 \dots 273 \\ FM_i * 7 * average(F_{i,274,m} : F_{i,365,m}) & \forall i, m, t = 274 \dots 365 \end{cases} \quad (9)$$

Equation (8) states that the semi-finished inventory target is equal to the semi-finished inventory target multiple times the average daily forecasted demand over the next quarter (91 days) multiplied by seven. This value of average daily forecasted demand remains the same for the last 91 periods to avoid an out of range error. Equation (9) describes the same relationships for the finished inventory. It is important to note that the daily forecast does not change for every time t , this value will change only when a new forecast is made, as explained later in this section. The $F_{i,t,m}$ variable is the forecast based on method (planning policy) m , which determines the frequency of forecasting f . The relationship between the frequency of forecasting f to method m is shown in Table 4.

The following flow balance equations describe the relationships between each stage in the supply chain:

$$SI_{i,t,m} = SI_{i,t-1,m} + \text{Prodi,t-70,m} - SItoFI_{i,t-1,m} \quad \forall i, m, t = 70 \dots 365 \quad (10)$$

$$FI_{i,t,m} = FI_{i,t-1,m} + SItoFI_{i,t-14,m} - \text{Ship}_{i,t-1,m} \quad \forall i, m, t = 14 \dots 365 \quad (11)$$

Equation (10) states that the semi-finished inventory at the beginning of time t equals the semi-finished inventory from the previous time period plus the production from the fab started 70 periods (the average production lead-time) prior minus the semi-finished inventory that was sent to the finished inventory warehouse through assembly/test in the previous period. Equation (11) states that the finished inventory at the beginning of time period t equals finished inventory in period $t-1$ plus inventory that was sent to the finished inventory warehouse through assembly/test 14 periods (the average assembly/test lead-time) prior to period t minus the finished inventory that was shipped to the customer in period $t-1$.

The following logic is used to calculate shipments to customers and back-orders:

$$Ship_{i,t,m} = \begin{cases} A_{i,t} + B_{i,t-1,m} & \text{if } FI_{i,t,m} \geq A_{i,t} + B_{i,t-1,m} \\ FI_{i,t,m} & \text{Otherwise} \end{cases} \quad \forall i, m, t > 1 \quad (12)$$

$$B_{i,t,m} = \begin{cases} 0 & \text{if } FI_{i,t,m} \geq A_{i,t} + B_{i,t-1,m} \\ \rho B_{i,t-1,m} - FI_{i,t,m} + A_{i,t} & \text{Otherwise} \end{cases} \quad \forall i, m, t > 1 \quad (13)$$

Equation (12) states that if the actual demand plus any backorders from previous periods is less than or equal to the quantity currently stored in finished inventory, shipments to the customer will equal the actual demand plus any back-orders. Otherwise, the entire quantity of inventory in the finished inventory warehouse is shipped to the customer. Equation (13) states that if the finished inventory is greater than or equal to the actual demand for the current period plus any cumulative back-orders, the backorder quantity is set to zero. Otherwise, the current backorder quantity equals the previous backorder quantity multiplied by the percent of customers that are willing to wait, plus the difference between finished inventory and actual demand.

To describe production in the fab or factory node, $Prod_{i,t,m}$, and the assembly/test node, $SIToFI_{i,t,m}$, additional variables must be defined to ensure that production and inventory levels are updated when new forecasts are released. The variables are used in the time index in place of t to relate the time t to the start of the period when a new forecast is made.

$p_{f,t}$ time corresponding to the start of the period between forecasts (day when new forecast is made) corresponding to frequency f and time t , plus 70 days (the production lead-time), used to calculate $Prod_{i,t,m}$

$r_{f,t}$ time corresponding to the start of the period between forecasts (day when new forecast is made) corresponding to frequency f and time t , plus 14 days (the assembly/test lead-time), used to calculate $SIToFI_{i,t,m}$

$s_{f,t}$ time corresponding to the start of the period between forecasts (day when new forecast is made) corresponding to frequency f and time t , used to calculate $SIT_{i,t,m}$ and $FIT_{i,t,m}$

The equations for $p_{f,t}$, $r_{f,t}$, and $s_{f,t}$ can be written as:

$$p_{f,t} = t - (t + 70) \% l_f \quad (14)$$

$$r_{f,t} = t - (t + 14) \% l_f \quad (15)$$

$$s_{f,t} = t - t \% l_f \quad (16)$$

Equation (14) states that $p_{f,t}$, the start of a period between forecasts corresponding to frequency f and time t , equals the current time t minus the remainder of t plus 70 divided by l_f , the length of the period corresponding to frequency f . The calculation for $s_{f,t}$ is similar except that 14 (the assembly test lead-time) is added instead of 70. These calculations are necessary so that updates are made based only upon available forecast information. For example, with a monthly forecasting frequency, during times $t=32$ through $t=60$, production is based on the forecast and inventory targets determined in time $t=31$, the first day of the new month.

$Prod_{i,t,m}$ and $SItOFI_{i,t,m}$ are calculated as follows:

$$Prod_{i,t,m} = \begin{cases} \text{average}(F_{i,p_{f,t}} : F_{i,p_{f,t}+l_f,m}) + SI_{i,t,m} - SIT_{i,s_{f,t},m} & \forall i, t = 1 \dots 295 - l_f \\ \text{average}(F_{i,296-l_f,m} : F_{i,365,m}) + SI_{i,t,m} - SIT_{i,s_{f,t},m} & \forall i, t = 296 - l_f \dots 365 \end{cases} \quad (17)$$

$$SItOFI_{i,t,m} = \begin{cases} \text{average}(F_{i,s_{f,t}} : F_{i,s_{f,t}+l_f,m}) + FI_{i,t,m} - FIT_{i,s_{f,t},m} & \forall i, t = 1 \dots 350 - l_f \\ \text{average}(F_{i,351-l_f,m} : F_{i,365,m}) + FI_{i,t,m} - FIT_{i,s_{f,t},m} & \forall i, t = 351 - l_f \dots 365 \end{cases} \quad (18)$$

Equation (17) states that the production is calculated as the average forecast from 70 periods from the current period (the average production lead time) through 70 plus l_f periods from the current period, plus each day an adjustment is made for the difference between the actual semi-finished inventory level and target semi-finished inventory level corresponding to $s_{f,t}$, the first day of the forecast period. Equation (18) states that the amount of inventory sent from the semi-

finished to finished inventory warehouse is equal to the average forecast from 14 periods from the current period (the expected assembly test processing time) to 14 plus 1f periods from the current period, plus an adjustment is made for the difference between target and actual finished inventory levels. Two different formulas are used for each variable in order to avoid an out of range error.

The following are additional constraints:

$$SItOFI_{i,t,m} \leq SI_{i,t,m} \quad \forall i,t,m \quad (19)$$

$$\begin{aligned} Prod_{i,t,m}, SItOFI_{i,t,m}, SI_{i,t,m}, FI_{i,t,m}, SM_i, FM_i, SIT_{i,t,m}, \\ FIT_{i,t,m}, MD_{i,t,m}, A_{i,t}, B_{i,t,m}, Ship_{i,t,m} \geq 0 \quad \forall i,t,m \quad (20) \end{aligned}$$

Equation (19) states that the amount sent from semi-finished to finished inventory through assembly/test cannot exceed the amount of inventory in the semi-finished inventory warehouse at any given time. Equation (20) is a non-negativity constraint for feasibility, and to ensure that no inventory is sent back to a prior node in the supply chain.

The forecasts used in the calculations vary based on the planning policy used because different forecasting frequencies are used for different policies. The generation of forecasts used in the model is described in Section 4.3.

4.3 Generation of Forecast Data

The forecasts used in the model are calculated based on historical forecast error and actual demand data. The process for generating the forecasts and assumptions used are described in the remainder of this section.

Since forecasts are currently made on a monthly basis, historical forecast data are not available for other time periods. For the simulation models, forecasts are generated based on a combination of actual daily demand and a randomly generated forecast error based on the historical mean and standard deviation of monthly forecast error. The forecast error is approximated by the standard deviation of the historical monthly forecasts. Actual demand data are available for one year (2004), and forecasts are generated over the same period for each timeframe of interest (weekly, biweekly, monthly, and quarterly).

The following additional variables must be defined before describing the generation of forecasts used in the model:

$A_{i,t,f}$	actual demand for product i aggregated for the period between forecasts corresponding to time t and frequency f
$\sigma_{i,f,q,t}$	the standard deviation of actual demand for product i by frequency f during the quarter q that corresponds to time t
$MAD_{i,t,f}$	actual demand for product i in time t for the period corresponding to frequency f , modified to reflect bias in forecasts for individual products
FEM	a forecast error multiplier that is varied to test different amounts of forecast error
$F_{i,t,f}$	daily demand forecast for product i in time t based on frequency f (the forecast for the period corresponding frequency f divided by l_f)

First, actual demand for each product is aggregated into the time periods corresponding to the four forecasting frequencies (quarterly, monthly, biweekly, and weekly) and stored as the variable $A_{i,t,f}$, calculated as follows:

$$A_{i,t,f} = \sum_{t=s_{f,t}}^{t=s_{f,t}+l_f} A_{i,t} \quad (21)$$

Next, data analysis is done to compare the distributions of historical monthly and forecasted actual demand using 2004 data for the 18 products involved in the study. Actual and forecasted demand are plotted and visually examined, and both appear to approximately follow a normal distribution. Based on a non-directional Wilcoxon Signed-Ranks test of differences in variance [Stamatis, 2003] between monthly demand and monthly forecasts, the Z -score of -0.75 is not enough evidence to conclude that the variances of actual and forecasted demand are unequal at the 95% significance level. Therefore, the demand and forecasts are assumed to be equally variable at the monthly level. Since no forecast data are available for the other time periods, it must be assumed that the distribution of forecasted and actual demand is approximately the same for these periods as well. For this reason, the standard deviation of forecasts is approximated with the standard deviation of actual demand for the time period of interest when randomly generating forecast data.

For the products and time periods of interest, forecasts are negatively biased by an average of one percent, meaning a slight under forecasting occurred. Since there is not evidence to conclude this difference is statistically different from zero based on a one-tailed t -test with an alpha of 0.1 and 17 degrees of freedom, it is assumed the expected value of forecast error averaged over all products and periods is zero. However, the forecasts for individual products averaged over all twelve months are biased by an average of 14 percent with a standard deviation of ten percent. To account for this bias, the actual demand for product i in time t for frequency f is substituted with modified actual demand, $MAD_{i,t,f}$. Via a simple simulation, it is determined that a normally-distributed random distribution with a mean of 1 and standard deviation 0.18 results in an average difference from 1 of 14 percent with a standard deviation of 10 percent,

approximately the same distribution as the overall forecast bias by product. Therefore, $MAD_{i,t,f}$ is defined as follows:

$$MAD_{i,t,f} = N \sim \left(\frac{A_{i,t,f}}{l_f}, \frac{A_{i,t,f}}{l_f} * 0.18 * FEM \right) \quad (22)$$

The same random number is used to calculate $MAD_{i,t,f}$ for all times t and frequencies f within a simulation run, but the number varies for each product and each simulation run. The forecast error multiplier, FEM , is set to 1 initially, and it is varied to test various amounts of forecast error.

Forecasts for a given product are generated based on a random number with a mean equal to $MAD_{i,t,f}$ and standard deviation of $\sigma_{i,f,q,t}$. Forecasts are based on the standard deviation for the quarter rather than for the year because demand contains seasonality, and the variability of demand changes with seasons as well as the mean. The forecast, $F_{i,t,f}$, is calculated as follows:

$$F_{i,t,f} = N \sim (MAD_{i,t,f}, \sigma_{i,f,q,t}) \quad (23)$$

It is determined that the resulting forecasts deviate for the actual demand by approximately 20% based on Equation (2) and a forecast error multiplier (FEM) of 1. The FEM is also varied to test each policy with 10% and 30% forecast error as well, in order to determine the stability of the policies with different levels of forecast error. These forecasts are the basis for planning inventory targets in the supply chain model. The length of the time period covered by the forecast is determined by the planning policy as discussed above.

Results

The output of the simulation runs is summarized and compared with the goal of finding a policy that results in the best trade-off between the objectives of minimizing overall supply chain costs (due to inventory and stock-outs) and minimizing supply chain production variability. The first goal is measured based on the cost of achieving a customer service level equation (including the cost of holding inventory and cost of lost sales due to stock-outs). Supply chain variability is measured by summing the variances of daily production in the factory and assembly/test echelons for the 365 periods over which the model is run.

The simulation is run for 10,000 replications with forecast error at the 10, 20, and 30 percent levels. Results are obtained for each of these scenarios to determine the stability and robustness of policies. First, general results are presented in order to provide an understanding of how each policy performs in terms of the two goals and to determine where further comparison is necessary. Next, several policies are compared in pairs in order to determine if statistical differences exist between their resulting costs and variances and to examine the trade-offs.

4.3.1. Overall Cost and Variance Results by Policy

Before presenting results for the control chart-based policies, cost and variance results are shown graphically for each frequency-based policy in Figure 10.




Figure 10. Cost and Variance Results for Frequency-Based Policies

Figure 10 shows that cost tends to increase as the frequency of updating decreases. A correlation exists between production variance and the frequency of updating, but it is not as strong as the relationship between cost and frequency of updating. For example, the quarterly policy results in a higher production variance than the biweekly and monthly policies due to the drastic changes in production that occur at the start of most quarters to compensate for large deviations from inventory targets based on the new forecast. Therefore, it is clear that the quarterly updating policy offers no advantage over any of the other frequency-based policies.

The sum and variance of supply chain costs as well as the sum of production variance over all 18 products and all three forecast errors combined resulting from each policy are shown in Table 5. Policies are listed in ascending order based on the sum of their costs, and production variances are also ranked in ascending order. All costs are computed based on equation (3) and an assumed 50% profit margin. Costs are summed over the 18 products and averaged over the 10,000 replications. The variable cost of production of each product is divided out of the cost for confidentiality reasons. All variances shown in the tables are divided by 1,000.

Table 5. Summarized Results for All Products, Forecast Errors

Frequency	λ	σ	Sum of Cost	Sum of Cost Variance	Sum of Prod. Var.	Cost Rank	Variance Rank
Biweekly to Weekly	0.1	1	287,847	103,292	35,485	1	28
Monthly to Weekly	0.3	1	290,127	104,028	34,527	2	20
Monthly to Weekly	0.1	1	290,186	103,239	35,050	3	27
Weekly	N/A	N/A	290,998	103,248	34,892	4	23
Biweekly to Weekly	0.1	1.5	291,365	109,527	35,017	5	26
Monthly to Weekly	0.1	1.5	291,802	107,542	34,991	6	25
Monthly to Biweekly	0.3	1	291,944	107,863	33,999	7	14
Monthly to Weekly	0.3	1.5	292,540	106,875	34,206	8	15
Monthly to Weekly	0.1	2	292,786	111,156	34,938	9	24
Monthly to Biweekly	0.1	1	292,865	109,156	34,211	10	16
Biweekly to Weekly	0.3	1	293,292	111,506	33,613	11	9
Monthly to Biweekly	0.3	1.5	293,497	108,683	33,935	12	12
Monthly to Biweekly	0.1	1.5	293,582	111,196	34,439	13	19
Biweekly to Weekly	0.1	2	293,674	114,924	34,422	14	18
Monthly to Weekly	0.3	2	293,890	109,922	33,951	15	13
Monthly to Biweekly	0.1	2	294,002	113,254	34,571	16	22
Monthly to Biweekly	0.3	2	294,549	110,472	33,803	17	10
Monthly to Weekly	0.1	3	294,961	112,844	34,548	18	21
Monthly to Biweekly	0.1	3	295,489	113,337	34,389	19	17
Biweekly to Weekly	0.1	3	295,574	116,284	33,927	20	11
Monthly to Weekly	0.3	3	296,131	112,984	33,602	21	8
Monthly to Biweekly	0.3	3	296,372	112,953	33,581	22	7
Biweekly to Weekly	0.3	1.5	296,740	116,476	33,328	23	2
Biweekly to Weekly	0.3	2	296,781	118,178	33,432	24	4
Biweekly to Weekly	0.3	3	297,080	117,275	33,370	25	3
Biweekly	N/A	N/A	297,185	117,158	33,276	26	1
Monthly	N/A	N/A	298,589	116,891	33,536	27	6
Quarterly	N/A	N/A	312,642	146,376	33,437	28	5

An Analysis of Variance (ANOVA) is performed on the control chart-based policies in order to determine which of the factors have a significant effect on the cost. The factors included in the study are the λ parameter, the size of the control limits, and the frequency of updating when in control and out of control. Results of the ANOVA indicate that the λ parameter and the size of the control limits have a significant impact on cost at the 5% significance level, while the frequency of updating does not. All three factors have a significant impact of variance.

Next, a Tukey's multiple comparison procedure is used to test the difference between individual levels of each factor. Tukey's procedure is applied because it protects the overall α , or experimentwise error, for the entire procedure rather than an individual comparison [Mendenhall and Sincich, 1995]. At the 5% significant level, results of this procedure indicate a significant difference in cost between a sigma multiplier of one compared to all three other levels (1.5, two, and three). A sigma multiplier of 1.5 is significantly different from a sigma multiplier of three but not one of two. A sigma multiplier of two does not result in a significantly different cost than a sigma multiplier of three at the 5% significance level. At the 10% significance level, all levels of size of control limits result in a significantly different cost except the two- and three-sigma control limits. A λ of 0.1 results in a significantly lower cost than a λ of 0.3 at the 5% significance level.

When testing for differences between variances, the Tukey's procedure shows the difference between variance is not significant for any level of sigma multiplier but the difference between a λ of 0.1 and a λ of 0.3 is significant at the 5% significance level. At the 10% significance level, a significant difference in variance exists between a sigma multiplier of one and a sigma multiplier of two or three, and between a sigma multiplier of 1.5 and three. No significant difference exists between a sigma multiplier of one and 1.5; between 1.5 and two, and between two and three at the 10% significance level.

The effects of these three factors are shown graphically in Figures 11-16.




Figure 11. Effect of the λ Parameter on Cost




Figure 12. Effect of the Sigma Multiplier Parameter on Cost




Figure 13. Effect of In Control and Out of Control Updating Frequency on Cost




Figure 14. Effect of the λ Parameter on Variance




Figure 15. Effect of the Sigma Multiplier Parameter on Variance




Figure 16. Effect of In Control and Out of Control Updating Frequency on Variance

Based on Table 5 and Figures 11-16, policies with smaller control limits and smaller λ values generally result in a lower total cost but higher variability than larger control limits. Table 6 helps to explain these effects by indicating the percent of out-of-control points detected by the control charts for each level of the two variable parameters, each level of forecast error, and each in control updating frequency: monthly and biweekly.

Table 6. Effect of λ Parameter and Size of Control Limits on Percent of Out of Control Points Detected by Control Chart

Lambda	Sigma	Forecast Error = 10%		Forecast Error = 20%		Forecast Error = 30%	
		Monthly	Biweekly	Monthly	Biweekly	Monthly	Biweekly
0.1	1	51%	43%	53%	44%	54%	44%
0.1	1.5	32%	19%	33%	19%	34%	20%
0.1	2	28%	17%	29%	17%	30%	17%
0.1	3	20%	5%	20%	5%	21%	5%
0.3	1	34%	23%	35%	23%	38%	23%
0.3	1.5	17%	7%	19%	7%	21%	7%
0.3	2	15%	6%	16%	6%	17%	6%
0.3	3	8%	2%	8%	2%	9%	2%

The table shows that lower values of the λ parameter and smaller control limits reduce costs by detecting more out of control points, and therefore switch to updating more frequently (biweekly or weekly) from less frequently (monthly or biweekly) more often. This, in turn, results in a higher variance.

Table 7 shows the rankings for cost and variance for each level of forecast error applied in the model. Policies are sorted ascending by their cost rank for the 20% forecast error scenario, and rankings in the lowest quartile for their columns are shown in bold. While the cost rankings are fairly stable, the benefits of the more frequent updating seem to increase with the level of forecast error. While the monthly policy ranked 15 out of 28 policies for the 10% forecast error scenario, its rank is 27 out of 28 for the 20% and 30% forecast error scenarios. By contrast, the weekly policy improves its cost rank from seventh with 10% forecast error to fourth with 20% and 30% forecast error.

Table 7. Rankings of Cost and Variance by Policy and Level of Forecast Error

Frequency	λ	σ	Forecast Error = 10%		Forecast Error = 20%		Forecast Error = 30%	
			Cost Rank	Var. Rank	Cost Rank	Var. Rank	Cost Rank	Var. Rank
Biweekly to Weekly	0.1	1	6	28	1	28	1	28
Monthly to Weekly	0.3	1	1	22	2	21	5	17
Monthly to Weekly	0.1	1	2	26	3	27	3	26
Weekly	N/A	N/A	7	24	4	23	4	23
Biweekly to Weekly	0.1	1.5	20	27	5	26	2	24
Monthly to Biweekly	0.3	1	5	14	6	14	8	27
Monthly to Weekly	0.1	1.5	3	25	7	25	9	11
Monthly to Weekly	0.3	1.5	4	15	8	16	11	16
Monthly to Biweekly	0.1	1	8	16	9	15	10	25
Monthly to Weekly	0.1	2	9	23	10	24	12	15
Monthly to Biweekly	0.1	1.5	14	18	11	19	7	3
Biweekly to Weekly	0.3	1	25	9	12	9	15	12
Monthly to Biweekly	0.3	1.5	10	13	13	13	16	19
Biweekly to Weekly	0.1	2	24	19	14	18	6	18
Monthly to Weekly	0.3	2	11	12	15	12	13	13
Monthly to Biweekly	0.1	2	12	21	16	22	14	21
Monthly to Biweekly	0.3	2	13	10	17	10	17	10
Monthly to Weekly	0.1	3	19	20	18	20	19	22
Monthly to Biweekly	0.1	3	18	17	19	17	21	20
Biweekly to Weekly	0.1	3	21	11	20	11	18	14
Monthly to Weekly	0.3	3	16	8	21	8	22	9
Monthly to Biweekly	0.3	3	17	7	22	7	24	8
Biweekly to Weekly	0.3	1.5	27	2	23	2	20	2
Biweekly to Weekly	0.3	2	26	5	24	5	23	5
Biweekly to Weekly	0.3	3	22	4	25	3	25	4
Biweekly	N/A	N/A	23	1	26	1	26	1
Monthly	N/A	N/A	15	6	27	6	27	6
Quarterly	N/A	N/A	28	3	28	4	28	7

Next, the average customer service level (over 10,000 replications and 18 products) resulting from each policy is shown for each level of forecast error. Service levels that fall in the upper quartile for the column are shown in bold.

Table 8. Customer Service Level Resulting from Each Policy by Level of Forecast Error

			Average Customer Service Level		
Frequency	λ	σ	10% Forecast Error	20% Forecast Error	30% Forecast Error
Biweekly to Weekly	0.1	1	92.77%	90.36%	87.11%
Biweekly to Weekly	0.1	1.5	92.47%	90.17%	86.60%
Monthly to Weekly	0.1	1	92.57%	90.14%	86.19%
Monthly to Weekly	0.1	1.5	92.68%	90.11%	86.53%
Monthly to Weekly	0.1	2	92.60%	90.07%	86.69%
Monthly to Weekly	0.3	1	92.52%	90.04%	86.10%
Biweekly to Weekly	0.1	2	92.43%	90.02%	86.13%
Weekly	N/A	N/A	92.45%	90.00%	86.34%
Monthly to Biweekly	0.1	2	92.61%	89.93%	86.27%
Monthly to Weekly	0.1	3	92.38%	89.92%	86.16%
Monthly to Weekly	0.3	1.5	92.41%	89.92%	86.15%
Monthly to Biweekly	0.1	1	92.53%	89.91%	85.43%
Monthly to Biweekly	0.1	1.5	92.47%	89.90%	85.49%
Monthly to Biweekly	0.3	1	92.44%	89.88%	85.65%
Monthly to Biweekly	0.1	3	92.43%	89.85%	86.02%
Biweekly to Weekly	0.1	3	92.34%	89.84%	85.83%
Monthly to Biweekly	0.3	1.5	92.50%	89.84%	85.81%
Monthly to Weekly	0.3	2	92.32%	89.79%	85.72%
Monthly to Biweekly	0.3	2	92.37%	89.74%	85.78%
Biweekly	N/A	N/A	92.37%	89.69%	84.73%
Biweekly to Weekly	0.3	3	92.34%	89.68%	84.77%
Biweekly to Weekly	0.3	1	92.25%	89.66%	85.38%
Biweekly to Weekly	0.3	2	92.33%	89.66%	84.97%
Biweekly to Weekly	0.3	1.5	92.27%	89.59%	85.30%
Monthly to Weekly	0.3	3	92.26%	89.59%	84.99%
Monthly to Biweekly	0.3	3	92.25%	89.57%	84.98%
Monthly	N/A	N/A	92.25%	89.48%	83.47%
Quarterly	N/A	N/A	91.24%	88.51%	82.52%

Although the customer service level may not be accurate due to the simplistic nature of the model, the relative customer service level of each policy is of interest.

Table 8 shows a similar relationship between customer service level and forecast error to the relationship between cost and forecast error shown in Table 7. The service level of the policies that involve more frequent updating (i.e., weekly) improves relative to the other policies as forecast error increases. While the difference between the service levels achieved with the various policies is minimal for the 10% forecast error scenario, it increases exponentially as

forecast error increases. Because the customer service level has a strong inverse relationship with the overall cost, it is obvious that lost sales is a major driver of cost.

Next, the products are segregated by their level of demand variability. Nine of the 18 products are classified as highly variable, meaning the coefficient of variation of their monthly demand is greater than 0.8, and the remaining nine products are classified as moderately variable.

In Table 9, the relative rankings for the costs and variances resulting from each policy are shown grouped by all products as well as by moderately and highly variable products. These two rankings are summed to provide insight into how well the policies perform overall. The difference between cost and variance rankings for the moderately and highly variable product classifications are also shown along with the sum of these differences in order to give an indication of the stability of policies across moderately and highly variable products. Rankings that fall within the lowest quartile for the column are shown in bold.

The table shows that policies can result in drastically different costs and variances when applied to highly variable products compared to moderately variable products. In particular, the weekly policy, which results in the lowest cost for the highly variable products, ranks 23 out of 28 in terms of cost for the moderately variable products. The variability rankings are much more stable, indicating that demand variability is an important factor to consider when identifying the policy that results in the best tradeoff between cost and variance minimization.

Table 9. Comparison of Cost and Variance Rankings by Product Variability Classification

Planning Policy			All Products			Mod. Variable		Highly Variable		Ranking Deltas		
Frequency	λ	σ	Cost	Var.	Sum	Cost	Var.	Cost	Var.	Cost	Var.	Sum
Biweekly to Weekly	0.1	1	1	28	29	1	28	3	28	2	0	2
Monthly to Weekly	0.3	1	2	20	22	7	18	4	23	3	5	8
Monthly to Weekly	0.1	1	3	27	30	18	25	2	26	16	1	17
Weekly	N/A	N/A	4	23	27	23	19	1	27	22	8	30
Biweekly to Weekly	0.1	1.5	5	26	31	6	23	8	25	2	2	4
Monthly to Weekly	0.1	1.5	6	25	31	15	26	6	24	9	2	11
Monthly to Biweekly	0.3	1	7	14	21	13	11	7	15	6	4	10
Monthly to Weekly	0.3	1.5	8	15	23	3	15	17	19	14	4	18
Monthly to Weekly	0.1	2	9	24	33	9	27	12	22	3	5	8
Monthly to Biweekly	0.1	1	10	16	26	20	16	5	18	15	2	17
Biweekly to Weekly	0.3	1	11	9	20	16	2	10	13	6	11	17
Monthly to Biweekly	0.3	1.5	12	12	24	5	12	20	12	15	0	15
Monthly to Biweekly	0.1	1.5	13	19	32	19	20	9	20	10	0	10
Biweekly to Weekly	0.1	2	14	18	32	17	17	11	21	6	4	10
Monthly to Weekly	0.3	2	15	13	28	2	14	21	11	19	3	22
Monthly to Biweekly	0.1	2	16	22	38	14	24	18	17	4	7	11
Monthly to Biweekly	0.3	2	17	10	27	4	10	23	9	19	1	20
Monthly to Weekly	0.1	3	18	21	39	8	22	22	16	14	6	20
Monthly to Biweekly	0.1	3	19	17	36	11	21	24	14	13	7	20
Biweekly to Weekly	0.1	3	20	11	31	22	13	15	10	7	3	10
Monthly to Weekly	0.3	3	21	8	29	10	9	25	8	15	1	16
Monthly to Biweekly	0.3	3	22	7	29	12	8	26	5	14	3	17
Biweekly to Weekly	0.3	1.5	23	2	25	24	1	19	7	5	6	11
Biweekly to Weekly	0.3	2	24	4	28	25	4	16	4	9	0	9
Biweekly to Weekly	0.3	3	25	3	28	26	6	14	2	12	4	16
Biweekly	N/A	N/A	26	1	27	27	3	13	1	14	2	16
Monthly	N/A	N/A	27	6	33	21	7	27	6	6	1	7
Quarterly	N/A	N/A	28	5	33	28	5	28	3	0	2	2

Additional ANOVA is performed on the products segregated by level of demand variability. For the highly variable products, all three factors have a significant impact on both cost and variance at the 5% significance level. For the moderately variable products, the λ parameter and the size of the control limits have a significant impact on cost and variance at the 5% significance level, while the frequency of updating does not have a significant effect on either measure.

Based on Tukey's multiple comparison procedure for the individual product groupings, results for the highly variable products indicate a significant difference in cost between all four levels of sigma multipliers except between 1.5 and two at the 5% significant level. At the 10%

significance level, all levels of size of control limits result in a significantly different cost. A λ of 0.1 results in a significantly lower cost than a λ of 0.3 at the 5% significance level. For the moderately variable products, a difference exists only between a sigma multiplier of one and three, and 1.5 and three at the 5% significant level; while the difference between and sigma multipliers of one and two is also significant with an overall α of 10%.

For the highly variable products, Tukey's procedure shows the difference between variance is significant for all levels of sigma multipliers except between 1.5 and two, and the difference between a λ of 0.1 and a λ of 0.3 is significant at the 5% significance level. At the 10% significance level, all differences between factor levels are significant. For the moderately variable products, the difference between variance is not significant for any level of sigma multipliers but the difference between a λ of 0.1 and a λ of 0.3 is significant at the 5% significance level. At the 10% significance level, a significant difference in variance exists only between a sigma multiplier of three and a sigma multiplier of one and 1.5.

Several policies are selected for further analysis based on their rankings as well as their consistency, measured by the delta between rankings for the moderately and highly variable product groupings. These policies include the monthly-to-weekly control chart-based policy with $\lambda=0.3$ and $\sigma=1$, which results in the second lowest cost, a moderate variance, and is consistent; the biweekly-to-weekly control chart-based policy with $\lambda=0.3$ and $\sigma=1$, because the sums of its ranks for cost and variance is the lowest; and the biweekly-to-weekly control chart-based policy with $\lambda=0.3$ and $\sigma=1.5$, because it has the second-lowest overall variance, a moderate cost, and is consistent. Next, more detailed statistical tests are performed to compare these policies amongst each other and to several frequency-based policies.

4.3.2. Paired Comparison of Policies

In this subsection, the results of tests to determine if a statistically significant difference exists between the policies in terms of cost and supply chain variability are shown and discussed. The costs and variability are summed for all 18 products, and hypothesis tests are performed to determine if a statistically significant difference between means or variances exists between each pair of policies for each level of forecast error as well as each grouping of products (all products, moderately variable products, and highly variable products).

The difference between the sum of the total cost achieved with various policies is evaluated by a large-sample test of hypothesis about the difference between two means for independent samples, while differences in variances are tested using an F-test for independent samples. A 5% significance level ($\alpha = 0.05$) is used for each statistical test. Additionally, each of the 18 products is evaluated individually. The statistical tests comparing policies in pairs are shown in the following tables. P-values that indicate a statistical difference at the 95% confidence level are shown in bold. All variances are divided by 1,000.

Because the control chart-based policies alternate between updating monthly when forecasts are in control and weekly or biweekly when forecasts are out of control, a comparison of the weekly and monthly frequency-based planning policies is shown first in Table 10.

Table 10. Comparison of Weekly vs. Monthly Planning Policies

Product Grouping	FE	Monthly			Weekly			p-value (cost)	p-value (var)
		Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.	Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.		
All	10%	95,866	31,867	11,081	94,478	30,271	11,522	<.01	0.04
All	20%	99,181	37,224	11,152	96,734	33,565	11,596	<.01	0.04
All	30%	103,542	47,800	11,303	99,786	39,411	11,774	<.01	0.04
Mod. Var.	10%	52,112	12,240	7,579	53,108	15,689	7,766	<.01	0.22
Mod. Var.	20%	54,493	16,260	7,633	54,878	18,168	7,796	<.01	0.30
Mod. Var.	30%	57,694	24,361	7,744	57,067	22,970	7,935	<.01	0.22
High Var.	10%	43,755	19,626	3,502	41,370	14,583	3,756	<.01	<.01
High Var.	20%	44,688	20,964	3,519	41,856	15,397	3,800	<.01	<.01
High Var.	30%	45,848	23,440	3,559	42,719	16,442	3,839	<.01	<.01

The weekly policy results in a significantly lower cost than the monthly policy overall and for the highly variable products with 95% confidence. However, the monthly policy has a significantly lower cost for the moderately variable products, except in the case of high forecast error. While the monthly policy exhibits a statistically lower variance for the overall and highly variable product groupings, no statistical difference in variance exists for the moderately variable products. Individually, 11 of 18 products (two moderately variable and all nine highly variable) achieved a statistically lower cost with the weekly policy, five of 18 products (all moderately variable) achieved a statistically lower cost with the monthly policy, and two moderately variable products showed no statistical difference in cost between the two policies. Fourteen out of 18 products have a significantly lower variance with the monthly policy, while the variance of the remaining four products is not significantly different.

Next, the monthly frequency-based policy is compared to the monthly-to-weekly control chart-based policy with $\lambda = 0.3$ and $\sigma = 1$, the policy with the second-lowest overall cost of all policies.

Table 11. Comparison of Monthly vs. Monthly to Weekly Control Chart with $\lambda=0.3$, $\sigma=1$
Planning Policies

Product Grouping	FE	Monthly			Monthly->Weekly CC $\lambda=0.3$ $\sigma=1$			p-value (cost)	p-value (var)
		Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.	Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.		
All	10%	95,866	31,867	11,081	93,975	28,794	11,427	<.01	0.14
All	20%	99,181	37,224	11,152	96,464	33,330	11,508	<.01	0.12
All	30%	103,542	47,800	11,303	99,688	41,905	11,592	<.01	0.20
Mod. Var.	10%	52,112	12,240	7,579	52,009	12,889	7,744	0.04	0.28
Mod. Var.	20%	54,493	16,260	7,633	53,764	15,988	7,787	<.01	0.32
Mod. Var.	30%	57,694	24,361	7,744	56,004	22,062	7,840	<.01	0.54
High Var.	10%	43,755	19,626	3,502	41,967	15,905	3,683	<.01	<.01
High Var.	20%	44,688	20,964	3,519	42,700	17,342	3,722	<.01	<.01
High Var.	30%	45,848	23,440	3,559	43,684	19,842	3,752	<.01	<.01

It can be concluded based on Table 11 that the monthly-to-weekly control chart-based policy with $\lambda = 0.3$ and $\sigma = 1$ results in a statistically lower cost than the monthly frequency-based policy for all product groupings and forecast scenarios with 95% confidence. On an individual product basis, 14 out of 18 (seven moderately variable and seven highly variable) products achieve a statistically lower cost with the control chart-based policy at the 95% confidence level, while three of 18 (all moderately variable) achieve a statistically lower cost with the monthly policy, and one highly variable products shows no statistical difference in cost. Only for the highly variable products groupings show a statistically lower variance resulting from the monthly policy for these scenarios. Individually, nine out of 18 products (five highly variable and four moderately variable) result in a lower variance with the monthly policy, while the remaining nine show no statistical difference.

Results for the control chart-based policy resemble those of the monthly policy for the moderately variable products, while the results for the highly variable products are similar to those of the weekly policy. The difference between results for the two product groupings is due

to the fact that for highly variable products, the difference between monthly and weekly forecasts is greater than for moderately variable products.

In Table 12, the monthly-to-weekly control chart-based policy with $\lambda=0.3$ is compared to the monthly-to-weekly control chart-based policy with $\lambda=0.1$ in order to determine the effect of the λ parameter, the weight applied to more recent compared to less recent observations in the calculation of the exponentially weighted moving average. A higher λ means more emphasis is given to the most recent point, while a lower λ places more emphasis on past data.

Table 12. Comparison of Monthly-to-Weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits vs. Monthly-to-Weekly Control Chart with $\lambda=0.1$ and 1σ Control Limits

Product Grouping	FE	Monthly->Weekly CC $\lambda=0.3$ $\sigma=1$			Monthly->Weekly CC $\lambda=0.1$ $\sigma=1$			p-value (cost)	p-value (var)
		Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.	Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.		
All	10%	93,975	28,794	11,427	94,024	28,988	11,560	0.52	0.54
All	20%	96,464	33,330	11,508	96,544	33,381	11,679	0.33	0.46
All	30%	99,688	41,905	11,592	99,618	40,869	11,811	0.44	0.34
Mod. Var.	10%	52,009	12,889	7,744	52,476	14,352	7,810	<.01	0.68
Mod. Var.	20%	53,764	15,988	7,787	54,233	17,105	7,893	<.01	0.48
Mod. Var.	30%	56,004	22,062	7,840	56,371	22,458	7,969	<.01	0.40
High Var.	10%	41,967	15,905	3,683	41,549	14,636	3,750	<.01	0.36
High Var.	20%	42,700	17,342	3,722	42,310	16,277	3,785	<.01	0.40
High Var.	30%	43,684	19,842	3,752	43,246	18,412	3,842	<.01	0.24

The policy with $\lambda=0.1$ results in a statistically lower cost than the policy with $\lambda=0.3$ for the moderately variable products, while the reverse is true for the highly variable products. Individually, seven out of 18 products (five moderately variable and two highly variable) have a statistically lower cost when applying the policy with $\lambda=0.1$, while five out of 18 products (one moderately variable and four highly variable) have a statistically lower cost when applying the policy with $\lambda=0.3$. No statistical difference in variance can be proven for any of the product groupings, but individually, five of 18 products have a statistically lower variance with the policy with $\lambda=0.3$, while the remaining 13 do not differ statistically in variance.

Table 13. Comparison of Monthly vs. Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits Planning Policies

Product Grouping	FE	Monthly			Biweekly->Weekly CC $\lambda=0.3$ $\sigma=1$			p-value (cost)	p-value (var)
		Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.	Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.		
All	10%	95,866	31,867	11,081	95,383	32,368	11,163	<.01	0.72
All	20%	99,181	37,224	11,152	97,482	36,042	11,203	<.01	0.80
All	30%	103,542	47,800	11,303	100,427	43,096	11,248	<.01	0.80
Mod. Var.	10%	52,112	12,240	7,579	52,678	15,567	7,576	<.01	0.99
Mod. Var.	20%	54,493	16,260	7,633	54,177	18,314	7,611	<.01	0.88
Mod. Var.	30%	57,694	24,361	7,744	56,078	22,900	7,636	<.01	0.48
High Var.	10%	43,755	19,626	3,502	42,706	16,801	3,587	<.01	0.24
High Var.	20%	44,688	20,964	3,519	43,305	17,728	3,591	<.01	0.30
High Var.	30%	45,848	23,440	3,559	44,350	20,196	3,611	<.01	0.44

The biweekly-to-weekly control chart-based policy with $\lambda = 0.3$ and $\sigma = 1$ results in a statistically lower cost than the monthly policy for the all product groupings with 95% confidence, with no statistical difference in variance. On an individual product basis, 12 out of 18 products (six moderately variable and six highly variable) have a statistically lower cost with the control chart-based policy, while five (three highly variable and two moderately variable) have a statistically lower cost with the monthly policy, and one show no statistical difference in cost. Seven products (all highly variable) have a statistically lower variance with the control chart-based policy, while seven (five moderately variable and two highly variable) have a lower variance with the monthly policy and four are not statistically different at the 95% confidence level. Similar results occur when comparing the monthly policy to the biweekly to weekly control chart-based policy with $\lambda = 0.3$ and $\sigma = 1.5$; cost is improved in all but the low forecast error scenario for moderately variable products and no significant difference in variance is observed.

A comparison of the biweekly-to-weekly control chart-based policies with $\lambda = 0.3$ and $\sigma = 1$ versus $\sigma = 1.5$ is shown in Table 14 to evaluate the effect of the width of control limits.

Table 14. Comparison of Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits vs. Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1.5σ Control Limits

		Biweekly->Weekly CC $\lambda=0.3$ $\sigma=1$			Biweekly->Weekly CC $\lambda=0.3$ $\sigma=1.5$				
Product Grouping	FE	Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.	Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.	p-value (cost)	p-value (var)
All	10%	95,383	32,368	11,163	96,274	33,293	11,004	<.01	0.48
All	20%	97,482	36,042	11,203	98,569	37,942	11,106	<.01	0.64
All	30%	100,427	43,096	11,248	101,896	45,241	11,218	<.01	0.88
Mod. Var.	10%	52,678	15,567	7,576	53,202	16,106	7,512	<.01	0.64
Mod. Var.	20%	54,177	18,314	7,611	54,855	19,717	7,582	<.01	0.84
Mod. Var.	30%	56,078	22,900	7,636	57,200	25,293	7,653	<.01	0.92
High Var.	10%	42,706	16,801	3,587	43,072	17,187	3,493	<.01	0.18
High Var.	20%	43,305	17,728	3,591	43,715	18,225	3,524	<.01	0.34
High Var.	30%	44,350	20,196	3,611	44,696	19,948	3,565	<.01	0.50

The policy with $\sigma = 1$ performs significantly better in terms of cost with no significant difference in variability. Individually, seven out of 18 products (two moderately variable and five highly variable) have a significantly lower cost with the policy with $\sigma = 1$, while four out of 18 products (three moderately variable and one highly variable) have a significantly lower cost with the policy with $\sigma = 1.5$, and the remaining seven products show no statistical difference in cost. 16 out of 18 products are not statistically different in terms of variance, while one product results in a statistically lower variance with each policy.

Table 15, a comparison of the biweekly frequency-based vs. biweekly control chart-based policy with $\lambda = 0.3$ and $\sigma = 1$, shows a significantly lower cost is achieved with the control chart-based policy for all scenarios except the moderately variable product grouping with 10% forecast with 95% confidence, and individually in ten out of 18 products (with one product achieving a lower cost with the biweekly policy and seven showing no statistical difference in cost). Variance is significantly lower with the control chart-based policy only for the highly variable product grouping with 10% forecast error, and individually for seven out of 18 products, with four products significantly lower in variance for the biweekly policy and seven showing no significant difference.

Table 15. Comparison of Biweekly vs. Biweekly-to-weekly Control Chart with $\lambda=0.3$ and 1σ Control Limits

Product Grouping	FE	Biweekly			Biweekly->Weekly CC $\lambda=0.3$ $\sigma=1$			p-value (cost)	p-value (var)
		Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.	Sum of Cost	Sum of Cost Var.	Sum of Prod. Var.		
All	10%	95,993	32,935	10,990	95,383	32,368	11,163	<.01	0.42
All	20%	98,736	37,735	11,076	97,482	36,042	11,203	<.01	0.58
All	30%	102,457	46,488	11,209	100,427	43,096	11,248	<.01	0.88
Mod. Var.	10%	53,341	16,092	7,547	52,678	15,567	7,576	<.01	0.84
Mod. Var.	20%	55,126	19,531	7,604	54,177	18,314	7,611	<.01	0.96
Mod. Var.	30%	57,773	26,408	7,717	56,078	22,900	7,636	<.01	0.58
High Var.	10%	42,652	16,843	3,444	42,706	16,801	3,587	0.35	0.04
High Var.	20%	43,610	18,203	3,472	43,305	17,728	3,591	<.01	0.10
High Var.	30%	44,684	20,080	3,492	44,350	20,196	3,611	<.01	0.10

4.3.3. Summary of Results

In this chapter, a novel control chart-based methodology for determining when to update inventory targets and production has been presented. The cost and production variance resulting from control chart-based policies with varying levels of the important parameters, as well as several frequency-based policies, have been compared under different levels of forecast error and demand variability. Several control chart-based policies have been shown to improve the cost without a significant increase in variability compared to the case organization's current planning method. However, for each policy there is a trade-off between minimizing cost and minimizing variability; no single policy is best at achieving both objectives.

The results presented in this chapter are specific to Intel's products, their characteristics, and supply chain. However, much of the findings can be generalized so that other companies and industries many apply them. In the remainder of this section, results are summarized for Intel and general findings are presented as well.

4.3.3.1 General Results

It is clear that moderately and highly variable products should be treated differently in order to minimize cost and variability. Highly variable products generally need tighter control, meaning more frequent forecasting and updating of production levels, than moderately variable products. However, tighter control of processes generally results in higher variability, so it is important to be selective about applying tighter control only when it is warranted. Policies that involve updating frequently generally do not improve cost for moderately variable products, an indication that more frequent planning is not necessary for many products. Control chart-based policies are robust because by applying the same control policy to moderately variable and highly variable products, the highly variable products will be updated more frequently than the moderately variable products, and as the underlying factors that determine how often updates should occur (demand variability and forecast error) change, the policy will automatically adjust the updating frequency. By applying control charts to indicate when tighter control of processes is necessary, resources are not invested unnecessarily in forecasting and variability is not increased by chasing noise in demand that could be managed with safety stock.

The results presented in this chapter show that forecast error is an important factor that must be considered when determining the ideal control policy. Forecast error increases the need for tight control of processes as the improvement in cost resulting from forecasting more frequently is magnified by forecast error. While the production variability also increases with forecast error, the relative rankings of policies with respect variance are much more stable than their rankings with respect to cost, so the trade-off between cost and variability minimization is impacted by the amount of forecast error. Therefore, although general relationships between

control policies and cost and variance can be understood by examining the Intel-specific results, the methodology presented in this dissertation should be followed and individual results obtained for other organizations to better understand these relationships for their specific products and characteristics.

4.3.3.2 Intel Specific Results

The choice of the appropriate control policy for Intel is dependent upon the decision-makers' desire to minimize variability compared to cost, as well as the amount of resources that can be invested in forecasting. In order to significantly improve cost without statistically increasing variability from the current monthly-based policy, the biweekly frequency-based policy or the biweekly-to-weekly control chart-based policies with $\lambda=0.1$ and $\sigma=1$ or 1.5 are all good options because while cost is significantly decreased compared to the monthly policy currently in place at the case organization, variability is not significantly increased. If cost minimization is the most important goal and more variability can be tolerated, the monthly-to-weekly control chart policy with $\lambda=0.3$ and $\sigma=1$ is a good option, because it results in the second-lowest overall cost of the 28 policies, it performs well for both moderately and highly variable products, and it produces moderate variability compared to the other low cost policies.

If a simpler policy is desired that requires fewer calculations than the control chart-based policies, one possible strategy is to update the highly variable products weekly and update the moderately variable products monthly or biweekly. However, demand variability of individual products may change over time and therefore, this strategy would require constant monitoring of demand variability to quickly detect changes. By contrast, control chart-based policies are more stable and robust, meaning the same policy can be applied to all products without monitoring of

the data to determine if changes have occurred. The biweekly policy is robust (performs well for both highly variable and moderately variable products) and results in the lowest overall variance (although this variance is not significantly lower than that of the monthly policy) and the lowest cost other than the weekly policy. Therefore, if a single frequency-based is to be selected for all products, the biweekly updating policy is recommended.

Control chart-based planning policies have been shown to provide a robust technique that adjusts appropriately for the situation and for all products regardless of their level of variability. The degree to which minimization of cost is prioritized compared to minimization of variability can be controlled through the manipulation of the λ and σ parameters. By applying a control chart-based policy rather than a strictly frequency-based policy, the risks of scrapping products due to obsolescence and excessive loss of sales due to stock-outs is reduced without increasing the amount of production variability and resources required for forecasting and planning more than necessary.

CHAPTER 5

SUMMARY AND CONCLUSIONS

5.1 Summary and Contribution

There is a considerable need to improve supply chain agility and to determine how to satisfy customers in the most cost-effective manner, which is evidenced by the focus on agility in the recent supply chain literature as a way for companies to increase their customer service levels and maintain market share. Two components of agility are addressed in this research: the ability to achieve a high customer service level and the ability to react effectively to change, by presenting a methodology for identifying a cost effective inventory and customer service level as well as a policy for production and inventory control that increases the ability to respond effectively to change while minimizing the likelihood of reacting to noise. This research provides insight into the ability of inventory planning policies to be agile, i.e. to react quickly to true shifts in demand without overreacting to false shifts and causing excess variability in the supply chain. The basis for the research is a case study of Intel Corporation, the world's largest semiconductor manufacturer.

The major contribution of this research is a methodology for production and inventory control in supply chain systems with non-stationary, stochastic demand. This methodology is based on a novel application of control charts to forecasting to detect when tighter control of production and inventory is desirable, which reduces the risk of excess inventory and stock-outs but does not unnecessarily increase variability and the resources required for planning.

The most unique aspect of the methodology is its ability to quantify the relationship between inventory, customer service level, cost, and variability. The methodology presented in

this study also provides the ability to determine a cost-effective customer service level that includes the cost of stock-outs, as well as the impact of forecast error, demand variability, and order lead-time on this relationship. This research contrasts most of the inventory modeling literature, which typically assumes demand is stationary and selects customer service level goals arbitrarily. Additionally, this research contributes a quantitative base in the area of supply chain agility, a research area that thus far has been primarily qualitative.

5.2 Future Research

There are several promising areas for future work based on the research in this dissertation. First, the approach developed to attain an agile supply chain was tested on a case study in the semiconductor industry. It would be interesting to apply the method to another industry and compare the results and insights gained. Additionally, there are some specific issues that arose in the development of the modeling approach that could provide additional avenues for deeper analysis as discussed in more detail below.

This study identifies cost-effective finished inventory targets for specific products while assuming the semi-finished inventory target is constant. One such future development could address the impact of semi-finished inventory levels on customer service level, and determine the most cost-effective ratio of finished to semi-finished goods inventory.

In the study, forecast error was assumed to be the same for all time periods because no actual forecast data exists for time periods other than monthly. Forecast error is typically higher for shorter time periods (weekly) than more aggregated time periods (quarterly). This may be offset completely or to some degree because forecasting for shorter time periods is done closer to

the time period of interest when more information is known about demand. Another direction for future research could be to study the effect of time period on forecast accuracy, and the results could be reflected in the comparison of policies.

This research addresses the cost of achieving a customer service level based on inventory holding costs and the immediate loss of revenue due to a stock-out. Future research could address an additional cost associated with stock-outs, the cost of lost future orders due to a loss of customer goodwill, possibly by correlating customer service level to customer satisfaction surveys. The addition of this cost would be likely to drive the ideal customer service and inventory levels higher.

Finally, another potential direction for future research involves quantifying the cost associated with variability, and in particular, the cost benefits achieved by reduction of throughput time. This would result in the ability to select the overall minimum cost policy rather than evaluating the tradeoff between cost and variability.

APPENDIX – CODE FOR SUPPLY CHAIN MODEL


```

#!/usr/bin/perl -w
use strict;
use Statistics::Descriptive;
use Math::Random;

#Read in input parameters
open(INFILE, "< $ARGV[0]") or die "Can not open the input file $ARGV[0].\n";

my $ADImult = 3;    #ADI inventory target multiple of forecasted demand
my $configs = 16;
my @cost;
my $costAvgCW = 0.26;
my $costAvgADI = 0.13;
my $costMissDem = 0.2;
my $costout = 1;
my $stepsout = 0;
my $CWMult = 5;    #CW inventory target multiple of forecasted demand
my $errorAvg = 1;
my $errorStd = 0.2;
my $ICL = $ARGV[1]; #flag to determine if inventory control level applied
my @M = qw(3 2);  #weighting factor for control chart
my $output = 0;
my @overallSD;    #overall standard deviation
my $Pwait = 0.8;  #% customers willing to wait
my @R = qw(0.3 0.1); #weighting factor for control chart
my $simTime = 359;
my $trials = 5000;

my %dailyAct;    #Hash of daily parameters defined in the input file
my %weeklyAct;   #Hash of weekly parameters defined in the input file
my %sweeklyAct;  #Hash of sweekly parameters defined in the input file
my %monthlyAct;  #Hash of monthly parameters defined in the input file
my %quarterlyAct; #Hash of quarterly parameters defined in the input file

if($output > 0)
{
    open(OUTFILE, "> output_ $ARGV[2]") or die "Can not open the output_ $ARGV[2] file.\n";

    #Header labels for the outfile
    print OUTFILE "TRIAL\tICL\tPRODUCT\tPERIOD\tPROD STDEV\tADI2CW
STDEV\tCW AVG\t";
    print OUTFILE "ADI AVG\tStock Out\tMissed Demand\tBack Orders\tService Level\t";
    print OUTFILE "Overall SD\tCost\n";
}

```

```

if($costout > 0)
{
  open(COSTFILE, "> cost_$ARGV[2]") or die "Can not open the cost_$ARGV[2] file.\n";
}

if($stepsout > 0)
{
  open(STEPSFILE, "> output_steps.tsv") or die "Can not open the output_steps file.\n";
  print STEPSFILE "Product\tConfig\tQTRtmp\tADIT\tADI\tADILCL\tADIUCL\tADIsum\t";
  print STEPSFILE "ADI2CW\tCWT\tCW\tCWLCL\tCWUCL\tCWsum\t";
  print STEPSFILE "prod\tship\tMD\tBO\tBOC\tSO\tActual\n";
}

```

#Read in first line of input file to get product names

```

$_ = <INFILE>;
chomp;
my @products = split /\t/;

```

##Parse actual data from a file##

```

while (<INFILE>)
{
  chomp;
  #Skip the line if it is blank
  if ($_ !~ /^s*$/)
  {
    #Split the input file based on tabs
    my @line = split /\t/;

    #reading in the actual demand by product
    for my $i (5..19)
    {
      my $day = $line[0];
      my $WW = $line[1];
      my $month = $line[2];
      my $quarter = $line[3];
      my $sweek = $line[4];

      #summing the actual demand by time period
      push @{$dailyAct{$products[$i]}}, $line[$i];
      $weeklyAct{$products[$i]}{$WW} += $line[$i];
      $sweeklyAct{$products[$i]}{$sweek} += $line[$i];
      $monthlyAct{$products[$i]}{$month} += $line[$i];
    }
  }
}

```

```

        $quarterlyAct{$products[$i]} {$quarter} += $line[$i];
    }
}
}
close INFILE;

###Generate forecast data###
my $count;
my $error;
my @norm;
my @rand;
my @stdev;

for my $trial (0..($trials - 1))
{
    print "TRIAL: $trial\n";

    #Consider each product separately
    for my $prod (@products[5..19])
    {
        my @weeklyFC;    #Weekly forecasts
        my @sweeklyFC;  #Semiweekly forecasts
        my @monthlyFC;  #Monthly forecasts
        my @quarterlyFC; #Quarterly forecasts

        #Variables for control charts
        my @monthStdevByQtr;
        my @sweekStdevByQtr;
        my @monthlyFC_CC;
        my @sweeklyFC_CC;

        ##Quarterly generation##
        my $stat = Statistics::Descriptive::Sparse->new();
        @stdev = ();
        @rand = ();
        $count = 0;

        #Iterate over each quarter for the product
        for my $quarter (keys %{$quarterlyAct{$prod}})
        {
            $stat->add_data($quarterlyAct{$prod} {$quarter});
        }

        #Calculate quarterly standard deviation

```

```

push @stdev, $stat->standard_deviation();

#Add forecast for each quarter w/ actual quarterly mean and quarterly
#stdev over the year
for my $quarter (sort keys %{$quarterlyAct{$prod}})
{
    #Generate one random forecast value for each quarter
    $error = random_normal(1, $errorAvg, $errorStd);
    @norm = random_normal(1, ($quarterlyAct{$prod}{$quarter}) * $error, $stdev[0]);

    #Only include the forecast if it is > 0; otherwise, use 0
    for my $normVal (@norm)
    {
        ($normVal > 0 ? push @rand, $normVal : push @rand, 0);
    }
}

#Divide each quarter forecast by 91 and add 91 copies to array
for my $quarterlyVal (@rand)
{
    my $dailyVal = $quarterlyVal / 91;
    push @quarterlyFC, ($dailyVal) x 91;
}

##Monthly generation##
$stat = Statistics::Descriptive::Sparse->new();
@stdev = ();
@rand = ();
$count = 0;

#Iterate over each month for the product
for my $month (sort keys %{$monthlyAct{$prod}})
{
    $count += 1;
    $stat->add_data($monthlyAct{$prod}{$month});

    #Calculate standard deviation every 3 months (quarter)
    if($count % 3 == 0)
    {
        push @stdev, $stat->standard_deviation();
        $stat = Statistics::Descriptive::Sparse->new();
    }
}

```

```

#Keep a copy of monthly stdev by quarter for control chart limit calc
@monthStdevByQtr = @stdev;

#Add forecast for each month w/ quarterly mean divided by 3 and monthly
#stdev over quarter
for my $quarter (sort keys %{$quarterlyAct{$prod}})
{
  #Generate three random forecast values for each quarter
  $error = random_normal(1, $errorAvg, $errorStd);
  @norm = random_normal(3, ($quarterlyAct{$prod} {$quarter}) * $error / 3, shift
@stdev);

  #Only include the forecast if it is > 0; otherwise, use 0
  for my $normVal (@norm)
  {
    ($normVal > 0 ? push @rand, $normVal : push @rand, 0);
  }
}

#Keep a copy of monthly forecasts for control chart calculations
@monthlyFC_CC = @rand;

#Divide each monthly forecast by 30 and add 30 copies to array
for my $monthlyVal (@rand)
{
  my $dailyVal = $monthlyVal / 30;
  push @monthlyFC, ($dailyVal) x 30;
}

###Biweekly generation###
$stat = Statistics::Descriptive::Sparse->new();
@stdev = ();
@rand = ();
$count = 0;
my $goal = 6; #semiweeks per quarter

#Iterate over each biweek for the product
for my $sweek (sort keys %{$sweeklyAct{$prod}})
{
  $count += 1;
  $stat->add_data($sweeklyAct{$prod} {$sweek});

  #Calculate standard deviation every 6 or 7 semiweeks (quarter)
  if($count == $goal)

```

```

{
  $count = 0;

  #Biweek boundaries are split at Q1/Q2 and Q3/Q4 boundaries, so
  #alternate 6 and 7 biweeks per quarter
  if($goal == 6)
  {
    $goal = 7;
  }
  else
  {
    $goal = 6;
  }

  push @stdev, $stat->standard_deviation();
  $stat = Statistics::Descriptive::Sparse->new();
}
}

#Keep a copy of montly stdev by quarter for control chart limit calc
@weekStdevByQtr = @stdev;

#Add forecast for each semiweek using quarterly mean divided by 6.5
# and sweekly stdev over quarter
$goal = 6;
for my $quarter (sort keys %{$quarterlyAct{$prod}})
{
  #Generate {goal} random forecast values for each quarter
  $error = random_normal(1, $errorAvg, $errorStd);
  @norm = random_normal($goal, ($quarterlyAct{$prod}{$quarter}) * $error / 6.5, shift
@stdev);

  #Only include the forecast if it is > 0; otherwise, use 0
  for my $normVal (@norm)
  {
    ($normVal > 0 ? push @rand, $normVal : push @rand, 0);
  }

  #Semiweek boundaries are split at Q1/Q2 and Q3/Q4 boundaries, so
  #alternate 6 and 7 semiweeks per quarter
  if($goal == 6)
  {
    $goal = 7;
  }
}

```

```

else
  {
    $goal = 6;
  }
}

#Keep a copy of biweekly forecasts for control chart calculations
@sweeklyFC_CC = @rand;

#Divide each biweekly forecast by 14 and add 14 copies to array
for my $sweeklyVal (@rand)
{
  my $dailyVal = $sweeklyVal / 14;
  push @sweeklyFC, ($dailyVal) x 14;
}

###Weekly generation###
$stat = Statistics::Descriptive::Sparse->new();
@rand = ();
@stdev = ();
$count = 0;

#Iterate over each month for the product
for my $week (sort keys %{$sweeklyAct{$sprod}})
{
  $count += 1;
  $stat->add_data($sweeklyAct{$sprod} {$week});

  #Calculate standard deviation every 13 weeks (quarter)
  if($count % 13 == 0)
  {
    push @stdev, $stat->standard_deviation();
    $stat = Statistics::Descriptive::Sparse->new();
  }
}

#Add forecast for each week using quarterly mean divided by 13 and
#weekly stdev over quarter
for my $quarter (sort keys %{$squarterlyAct{$sprod}})
{
  #Generate 13 random forecast values for each quarter
  $error = random_normal(1, $errorAvg, $errorStd);
  @norm = random_normal(13, ($squarterlyAct{$sprod} {$quarter}) * $error / 13, shift
@stdev);

```

```

#Only include the forecast if it is > 0; otherwise, use 0
for my $normVal (@norm)
{
    ($normVal > 0 ? push @rand, $normVal : push @rand, 0);
}
}

#Divide each monthly forecast by 30 and add 7 copies to array
for my $weeklyVal (@rand)
{
    my $dailyVal = $weeklyVal / 7;
    push @weeklyFC, ($dailyVal) x 7;
}

##Calculate control charts on forecast##
my @swFC_WK_3_3 = @swweeklyFC; #Biweekly--> weekly, R=.3, 3 sigma control limits
my @swFC_WK_1_3 = @swweeklyFC; #Biweekly--> weekly, R=.1, 3 sigma control limits
my @swFC_WK_3_2 = @swweeklyFC;
my @swFC_WK_1_2 = @swweeklyFC;
my @mFC_SW_3_3 = @monthlyFC; #Monthly--> biweekly, R=.3, 3 sigma control limits
my @mFC_SW_1_3 = @monthlyFC;
my @mFC_SW_3_2 = @monthlyFC;
my @mFC_SW_1_2 = @monthlyFC;
my @mFC_WK_3_3 = @monthlyFC; #Monthly--> weekly, R=.3, 3 sigma control limits
my @mFC_WK_1_3 = @monthlyFC;
my @mFC_WK_3_2 = @monthlyFC;
my @mFC_WK_1_2 = @monthlyFC;
my @movingAvg = ($monthlyFC_CC[0]) x 2;
my $qtr = 0; #index of quarter in moving average is being calculated

#Calculate control chart variations
for(my $month = 1; $month < $#monthlyFC_CC; $month++)
{
    #Increment index of stdev by quarter
    ($month % 4 == 0 ? $qtr++ : 0);

    #Calculate moving average and limits
    $movingAvg[0] = ($R[0] * $monthlyFC_CC[$month]) + ((1 - $R[0]) * $movingAvg[0]);
    my $LCL_3_3 = $movingAvg[0] - ($M[0] * $monthStdevByQtr[$qtr]);
    my $UCL_3_3 = $movingAvg[0] + ($M[0] * $monthStdevByQtr[$qtr]);
    my $LCL_3_2 = $movingAvg[0] - ($M[1] * $monthStdevByQtr[$qtr]);
    my $UCL_3_2 = $movingAvg[0] + ($M[1] * $monthStdevByQtr[$qtr]);
}

```



```

$movingAvg[1] = ($R[1] * $monthlyFC_CC[$month]) + ((1 - $R[1]) * $movingAvg[1]);
my $LCL_1_3 = $movingAvg[1] - ($M[0] * $monthStdevByQtr[$qtr]);
my $UCL_1_3 = $movingAvg[1] + ($M[0] * $monthStdevByQtr[$qtr]);
my $LCL_1_2 = $movingAvg[1] - ($M[1] * $monthStdevByQtr[$qtr]);
my $UCL_1_2 = $movingAvg[1] + ($M[1] * $monthStdevByQtr[$qtr]);

#Determine if forecast is outside limits
if($monthlyFC_CC[$month] > $UCL_3_3 or $monthlyFC_CC[$month] < $LCL_3_3)
{
  #print "Replacing month: $i for 0.3 3\n";
  #When forecast is OOC, replace monthly forecast with semiweekly
  for my $day (($month * 30) .. ($month * 30 + 30))
  {
    $mFC_SW_3_3[$day] = $weeklyFC[$day];
    $mFC_WK_3_3[$day] = $weeklyFC[$day];
  }
}

if($monthlyFC_CC[$month] > $UCL_3_2 or $monthlyFC_CC[$month] < $LCL_3_2)
{
  #print "Replacing month: $i for 0.3 2\n";
  #Replace monthly forecast with semiweekly
  for my $day (($month * 30) .. (($month * 30) + 30))
  {
    $mFC_SW_3_2[$day] = $weeklyFC[$day];
    $mFC_WK_3_2[$day] = $weeklyFC[$day];
  }
}

if($monthlyFC_CC[$month] > $UCL_1_3 or $monthlyFC_CC[$month] < $LCL_1_3)
{
  #print "Replacing month: $i for 0.1 3\n";
  #When forecast is OOC, replace monthly forecast with semiweekly
  for my $day (($month * 30) .. ($month * 30 + 30))
  {
    $mFC_SW_1_3[$day] = $weeklyFC[$day];
    $mFC_WK_1_3[$day] = $weeklyFC[$day];
  }
}

if($monthlyFC_CC[$month] > $UCL_1_2 or $monthlyFC_CC[$month] < $LCL_1_2)
{
  #print "Replacing month: $i for 0.1 2\n";
  #Replace monthly forecast with semiweekly
  for my $day (($month * 30) .. (($month * 30) + 30))

```

```

    {
        $mFC_SW_1_2[$day] = $sweeklyFC[$day];
        $mFC_WK_1_2[$day] = $weeklyFC[$day];
    }
}

@movingAvg = ($sweeklyFC_CC[0]) x 2;
$qtr = 0;
$goal = 6;
$count = 0;
#Calculate control chart variations
for(my $sweek = 1; $sweek < $#sweeklyFC_CC; $sweek++)
{
    $count++;
    #Increment index of stdev by quarter
    if($count == $goal)
    {
        $count = 0;
        $qtr++;
        if($goal == 6)
        {
            $goal = 7;
        }
        else
        {
            $goal = 6;
        }
    }
}

#Calculate moving average and limits
$movingAvg[0] = ($R[0] * $sweeklyFC_CC[$sweek]) + ((1 - $R[0]) * $movingAvg[0]);
my $LCL_3_3 = $movingAvg[0] - ($M[0] * $sweekStdevByQtr[$qtr]);
my $UCL_3_3 = $movingAvg[0] + ($M[0] * $sweekStdevByQtr[$qtr]);
my $LCL_3_2 = $movingAvg[0] - ($M[1] * $sweekStdevByQtr[$qtr]);
my $UCL_3_2 = $movingAvg[0] + ($M[1] * $sweekStdevByQtr[$qtr]);

$movingAvg[1] = ($R[1] * $sweeklyFC_CC[$sweek]) + ((1 - $R[1]) * $movingAvg[1]);
my $LCL_1_3 = $movingAvg[1] - ($M[0] * $sweekStdevByQtr[$qtr]);
my $UCL_1_3 = $movingAvg[1] + ($M[0] * $sweekStdevByQtr[$qtr]);
my $LCL_1_2 = $movingAvg[1] - ($M[1] * $sweekStdevByQtr[$qtr]);
my $UCL_1_2 = $movingAvg[1] + ($M[1] * $sweekStdevByQtr[$qtr]);

#Determine if forecast is outside limits

```

```

if($sweeklyFC_CC[$sweek] > $UCL_3_3 or $sweeklyFC_CC[$sweek] < $LCL_3_3)
{
  #print "Replacing sweek: $i for 0.3 3\n";
  #When forecast is OOC, replace sweekly forecast with semiweekly
  for my $day (($sweek * 14) .. ($sweek * 14 + 14))
  {
    $swFC_WK_3_3[$day] = $weeklyFC[$day];
  }
}

if($sweeklyFC_CC[$sweek] > $UCL_3_2 or $sweeklyFC_CC[$sweek] < $LCL_3_2)
{
  #print "Replacing sweek: $i for 0.3 2\n";
  #Replace sweekly forecast with semiweekly
  for my $day (($sweek * 14) .. (($sweek * 14) + 14))
  {
    $swFC_WK_3_2[$day] = $weeklyFC[$day];
  }
}

if($sweeklyFC_CC[$sweek] > $UCL_1_3 or $sweeklyFC_CC[$sweek] < $LCL_1_3)
{
  #print "Replacing sweek: $i for 0.1 3\n";
  #When forecast is OOC, replace sweekly forecast with semiweekly
  for my $day (($sweek * 14) .. ($sweek * 14 + 14))
  {
    $swFC_WK_1_3[$day] = $weeklyFC[$day];
  }
}

if($sweeklyFC_CC[$sweek] > $UCL_1_2 or $sweeklyFC_CC[$sweek] < $LCL_1_2)
{
  #print "Replacing sweek: $i for 0.1 2\n";
  #Replace sweekly forecast with semiweekly
  for my $day (($sweek * 14) .. (($sweek * 14) + 14))
  {
    $swFC_WK_1_2[$day] = $weeklyFC[$day];
  }
}
}

## Start simulation##
my @ADI = (0)x$configs;    #ADI
my @ADILCL = (0)x$configs; #ADI lower control limit

```

```

my @ADICnt = (0)x$configs; #count of ADI used to calculate average
my @ADISum = (0)x$configs; #sum of ADI used to calculate average
my @ADIT = (0)x$configs; #ADI Target
my @ADIUCL = (0)x$configs; #ADI upper control limit
my @ADI2CW; #ADI to CW
my @ADI2CWSD = (0)x$configs; #ADI2CW standard deviation
my @BO = (0)x$configs; #backorders
my @BOC = (0)x$configs; #cumulative backorders
my @CW = (0)x$configs; #CW
my @CWcnt = (0)x$configs; #count of CW used to calculate average
my @CWLCL = (0)x$configs; #CW lower control limit
my @CWsum = (0)x$configs; #sum of CW used to calculate average
my @CWT = (0)x$configs; #CW Target
my @CWUCL = (0)x$configs; #CW upper control limit
my @MD = (0)x$configs; #missed demand
my @prod; #production
my @prodSD = (0)x$configs; #production standard deviation
my @ship = (0)x$configs; #shipments
my @SO = (0)x$configs; #stock-outs
my @SL = (0)x$configs; #service level

```

#Find initial counts for averaging and then just add/sub a single val

#for subsequent calculations

```
my @QTRtmp = (0)x$configs;
```

for my \$day (0..90) #sums the forecast for the Qtr by frequency

```

{
    $QTRtmp[0] += $quarterlyFC[$day];
    $QTRtmp[1] += $monthlyFC[$day];
    $QTRtmp[2] += $weeklyFC[$day];
    $QTRtmp[3] += $weeklyFC[$day];

    $QTRtmp[4] += $mFC_SW_3_3[$day];
    $QTRtmp[5] += $mFC_SW_1_3[$day];
    $QTRtmp[6] += $mFC_SW_3_2[$day];
    $QTRtmp[7] += $mFC_SW_1_2[$day];

    $QTRtmp[8] += $mFC_WK_3_3[$day];
    $QTRtmp[9] += $mFC_WK_1_3[$day];
    $QTRtmp[10] += $mFC_WK_3_2[$day];
    $QTRtmp[11] += $mFC_WK_1_2[$day];

    $QTRtmp[12] += $swFC_WK_3_3[$day];
    $QTRtmp[13] += $swFC_WK_1_3[$day];
    $QTRtmp[14] += $swFC_WK_3_2[$day];

```

```
$QTRtmp[15] += $swFC_WK_1_2[$day];  
}
```

```
my @prodtmp = (0)x$configs;
```

```
for my $day (69..97)
```

```
{
```

```
  $prodtmp[0] += $quarterlyFC[$day];
```

```
  $prodtmp[1] += $monthlyFC[$day];
```

```
  $prodtmp[2] += $weeklyFC[$day];
```

```
  $prodtmp[3] += $weeklyFC[$day];
```

```
  $prodtmp[4] += $mFC_SW_3_3[$day];
```

```
  $prodtmp[5] += $mFC_SW_1_3[$day];
```

```
  $prodtmp[6] += $mFC_SW_3_2[$day];
```

```
  $prodtmp[7] += $mFC_SW_1_2[$day];
```

```
  $prodtmp[8] += $mFC_WK_3_3[$day];
```

```
  $prodtmp[9] += $mFC_WK_1_3[$day];
```

```
  $prodtmp[10] += $mFC_WK_3_2[$day];
```

```
  $prodtmp[11] += $mFC_WK_1_2[$day];
```

```
  $prodtmp[12] += $swFC_WK_3_3[$day];
```

```
  $prodtmp[13] += $swFC_WK_1_3[$day];
```

```
  $prodtmp[14] += $swFC_WK_3_2[$day];
```

```
  $prodtmp[15] += $swFC_WK_1_2[$day];
```

```
}
```

```
my @ADI2CWtmp = (0)x$configs;
```

```
for my $day (9..17)
```

```
{
```

```
  $ADI2CWtmp[0] += $quarterlyFC[$day];
```

```
  $ADI2CWtmp[1] += $monthlyFC[$day];
```

```
  $ADI2CWtmp[2] += $weeklyFC[$day];
```

```
  $ADI2CWtmp[3] += $weeklyFC[$day];
```

```
  $ADI2CWtmp[4] += $mFC_SW_3_3[$day];
```

```
  $ADI2CWtmp[5] += $mFC_SW_1_3[$day];
```

```
  $ADI2CWtmp[6] += $mFC_SW_3_2[$day];
```

```
  $ADI2CWtmp[7] += $mFC_SW_1_2[$day];
```

```
  $ADI2CWtmp[8] += $mFC_WK_3_3[$day];
```

```
  $ADI2CWtmp[9] += $mFC_WK_1_3[$day];
```

```
  $ADI2CWtmp[10] += $mFC_WK_3_2[$day];
```

```
  $ADI2CWtmp[11] += $mFC_WK_1_2[$day];
```

```

$ADI2CWtmp[12] += $swFC_WK_3_3[$day];
$ADI2CWtmp[13] += $swFC_WK_1_3[$day];
$ADI2CWtmp[14] += $swFC_WK_3_2[$day];
$ADI2CWtmp[15] += $swFC_WK_1_2[$day];
}

#Simulate for 12 x 30 = 360 days
for my $time (0..$simTime)
{
  #Iterate through all configurations
  for my $i (0..($configs - 1))
  {
    my $QTRavg= $QTRtmp[$i] / 91; #avg forecast for Qtr

    #Calculate ADIT and CWT based on the average of one quarter
    $ADIT[$i] = $ADImult * $QTRavg;
    $CWT[$i] = $CWmult * $QTRavg;

    #Calculate the ICL based on the average of the first 91 forecasts
    $ADILCL[$i] = ($ADImult - $W) * $QTRavg;
    $ADIUCL[$i] = ($ADImult + $W) * $QTRavg;
    $CWLCL[$i] = ($CWmult - $W) * $QTRavg;
    $CWUCL[$i] = ($CWmult + $W) * $QTRavg;

    #For time < 14, CW equals CWT
    if($time < 14)
    {
      $CW[$i] = $CWT[$i];
    }
    else
    {
      $CW[$i] = $CW[$i] - $ship[$i] + $ADI2CW[$i][$time - 14];
    }

    $CWsum[$i] += $CW[$i];
    $CWcnt[$i] += 1;

    #For time < 70, ADI equals ADIT
    if($time < 70)
    {
      $ADI[$i] = $ADIT[$i];
    }
    else

```

```

{
  $ADI[$i] = $ADI[$i] - $ADI2CW[$i][$time - 1] + $prod[$i][$time - 70];
}

```

```

$ADIsun[$i] += $ADI[$i];
$ADICnt[$i] += 1;

```

```

#Calculate PROD as average of 69..97 + ADI - ADIT
if(0 == $ICL or $ADI[$i] < $ADILCL[$i] or $ADI[$i] > $ADIUCL[$i])
{
  $prod[$i][$time] = ($prodtmp[$i] / 29) + $ADIT[$i] - $ADI[$i];
}
else
{
  $prod[$i][$time] = $prodtmp[$i] / 29;
}

```

```

if($prod[$i][$time] < 0)
{
  $prod[$i][$time] = 0;
}

```

```

#Calculate ADI2CW as (average of 9..17) + CW - CWT
if(0 == $ICL or $CW[$i] < $CWLCL[$i] or $CW[$i] > $CWUCL[$i])
{
  $ADI2CW[$i][$time] = ($ADI2CWtmp[$i] / 9) + $CWT[$i] - $CW[$i];
}
else #Do not make adjustment if in control
{
  $ADI2CW[$i][$time] = $ADI2CWtmp[$i] / 9;
}

```

```

if($ADI2CW[$i][$time] > $ADI[$i]) #can only send what's in ADI to CW
{
  $ADI2CW[$i][$time] = $ADI[$i];
}

```

```

if($ADI2CW[$i][$time] < 0)
{
  $ADI2CW[$i][$time] = 0;
}

```

```

#Calculate shipments, backorders, and stock-outs
my $actual = ${$dailyAct{$prod}}[$time];

```

```

if($CW[$i] >= $actual + $BO[$i])
{
    $ship[$i] = $actual + $BO[$i];
    $BO[$i] = 0;
}
else
{
    $ship[$i] = $CW[$i];
    $MD[$i] += (1 - $Pwait) * $BO[$i];
    $BO[$i] = $Pwait * $BO[$i] - $CW[$i] + $actual;
    $BOC[$i] += $BO[$i];
}

#Stock-out
if($CW[$i] == 0)
{
    $SO[$i] += 1;
}
}

#Update the sum of the ICL's for the next iteration except when
#time > 269 when the average will be kept the same
if($time < 270)
{
    $QTRtmp[0] += ($quarterlyFC[$time + 90] - $quarterlyFC[$time]);
    $QTRtmp[1] += ($monthlyFC[$time + 90] - $monthlyFC[$time]);
    $QTRtmp[2] += ($weeklyFC[$time + 90] - $weeklyFC[$time]);
    $QTRtmp[3] += ($weeklyFC[$time + 90] - $weeklyFC[$time]);

    $QTRtmp[4] += ($mFC_SW_3_3[$time + 90] - $mFC_SW_3_3[$time]);
    $QTRtmp[5] += ($mFC_SW_1_3[$time + 90] - $mFC_SW_1_3[$time]);
    $QTRtmp[6] += ($mFC_SW_3_2[$time + 90] - $mFC_SW_3_2[$time]);
    $QTRtmp[7] += ($mFC_SW_1_2[$time + 90] - $mFC_SW_1_2[$time]);

    $QTRtmp[8] += ($mFC_WK_3_3[$time + 90] - $mFC_WK_3_3[$time]);
    $QTRtmp[9] += ($mFC_WK_1_3[$time + 90] - $mFC_WK_1_3[$time]);
    $QTRtmp[10] += ($mFC_WK_3_2[$time + 90] - $mFC_WK_3_2[$time]);
    $QTRtmp[11] += ($mFC_WK_1_2[$time + 90] - $mFC_WK_1_2[$time]);

    $QTRtmp[12] += ($swFC_WK_3_3[$time + 90] - $swFC_WK_3_3[$time]);
    $QTRtmp[13] += ($swFC_WK_1_3[$time + 90] - $swFC_WK_1_3[$time]);
    $QTRtmp[14] += ($swFC_WK_3_2[$time + 90] - $swFC_WK_3_2[$time]);
    $QTRtmp[15] += ($swFC_WK_1_2[$time + 90] - $swFC_WK_1_2[$time]);
}

```



```
}
```

```
#Update the sum of the PROD's for the next iteration except when
```

```
#time > 262 when the average will be kept the same
```

```
if($time < 263)
```

```
{
```

```
  $prodtmp[0] += ($quarterlyFC[$time + 97] - $quarterlyFC[$time + 69]);
```

```
  $prodtmp[1] += ($monthlyFC[$time + 97] - $monthlyFC[$time + 69]);
```

```
  $prodtmp[2] += ($weeklyFC[$time + 97] - $weeklyFC[$time + 69]);
```

```
  $prodtmp[3] += ($weeklyFC[$time + 97] - $weeklyFC[$time + 69]);
```

```
  $prodtmp[4] += ($mFC_SW_3_3[$time + 97] - $mFC_SW_3_3[$time + 69]);
```

```
  $prodtmp[5] += ($mFC_SW_1_3[$time + 97] - $mFC_SW_1_3[$time + 69]);
```

```
  $prodtmp[6] += ($mFC_SW_3_2[$time + 97] - $mFC_SW_3_2[$time + 69]);
```

```
  $prodtmp[7] += ($mFC_SW_1_2[$time + 97] - $mFC_SW_1_2[$time + 69]);
```

```
  $prodtmp[8] += ($mFC_WK_3_3[$time + 97] - $mFC_WK_3_3[$time + 69]);
```

```
  $prodtmp[9] += ($mFC_WK_1_3[$time + 97] - $mFC_WK_1_3[$time + 69]);
```

```
  $prodtmp[10] += ($mFC_WK_3_2[$time + 97] - $mFC_WK_3_2[$time + 69]);
```

```
  $prodtmp[11] += ($mFC_WK_1_2[$time + 97] - $mFC_WK_1_2[$time + 69]);
```

```
  $prodtmp[12] += ($swFC_WK_3_3[$time + 97] - $swFC_WK_3_3[$time + 69]);
```

```
  $prodtmp[13] += ($swFC_WK_1_3[$time + 97] - $swFC_WK_1_3[$time + 69]);
```

```
  $prodtmp[14] += ($swFC_WK_3_2[$time + 97] - $swFC_WK_3_2[$time + 69]);
```

```
  $prodtmp[15] += ($swFC_WK_1_2[$time + 97] - $swFC_WK_1_2[$time + 69]);
```

```
}
```

```
#Update the sum of the ADI2CW's for the next iteration except when
```

```
#time > 342 when the average will be kept the same
```

```
if($time < 343)
```

```
{
```

```
  $ADI2CWtmp[0] += ($quarterlyFC[$time + 17] - $quarterlyFC[$time + 9]);
```

```
  $ADI2CWtmp[1] += ($monthlyFC[$time + 17] - $monthlyFC[$time + 9]);
```

```
  $ADI2CWtmp[2] += ($weeklyFC[$time + 17] - $weeklyFC[$time + 9]);
```

```
  $ADI2CWtmp[3] += ($weeklyFC[$time + 17] - $weeklyFC[$time + 9]);
```

```
  $ADI2CWtmp[4] += ($mFC_SW_3_3[$time + 17] - $mFC_SW_3_3[$time + 9]);
```

```
  $ADI2CWtmp[5] += ($mFC_SW_1_3[$time + 17] - $mFC_SW_1_3[$time + 9]);
```

```
  $ADI2CWtmp[6] += ($mFC_SW_3_2[$time + 17] - $mFC_SW_3_2[$time + 9]);
```

```
  $ADI2CWtmp[7] += ($mFC_SW_1_2[$time + 17] - $mFC_SW_1_2[$time + 9]);
```

```
  $ADI2CWtmp[8] += ($mFC_WK_3_3[$time + 17] - $mFC_WK_3_3[$time + 9]);
```

```
  $ADI2CWtmp[9] += ($mFC_WK_1_3[$time + 17] - $mFC_WK_1_3[$time + 9]);
```

```
  $ADI2CWtmp[10] += ($mFC_WK_3_2[$time + 17] - $mFC_WK_3_2[$time + 9]);
```

```

$ADI2CWtmp[11] += ($mFC_WK_1_2[$time + 17] - $mFC_WK_1_2[$time + 9]);

$ADI2CWtmp[12] += ($swFC_WK_3_3[$time + 17] - $swFC_WK_3_3[$time + 9]);
$ADI2CWtmp[13] += ($swFC_WK_1_3[$time + 17] - $swFC_WK_1_3[$time + 9]);
$ADI2CWtmp[14] += ($swFC_WK_3_2[$time + 17] - $swFC_WK_3_2[$time + 9]);
$ADI2CWtmp[15] += ($swFC_WK_1_2[$time + 17] - $swFC_WK_1_2[$time + 9]);
}

#Calculate statistics
for my $i (0..($configs - 1))
{
    $SL[$i] = ($simTime - $SO[$i]) / $simTime;

    #Calculate stdev of production and ADI2CW
    $stat = Statistics::Descriptive::Sparse->new();
    $stat->add_data(@{$prod[$i]});
    $prodSD[$i] = $stat->standard_deviation();

    $stat = Statistics::Descriptive::Sparse->new();
    $stat->add_data(@{$ADI2CW[$i]});
    $ADI2CWSD[$i] = $stat->standard_deviation();

    #Calculate average inventory levels
    $ADIsun[$i] /= $ADICnt[$i];
    $CWsum[$i] /= $CWent[$i];
    $overallSD[$i]{$prod}{$Strial} = sqrt(($prodSD[$i] * $prodSD[$i]) +
($ADI2CWSD[$i] * $ADI2CWSD[$i]));
    $cost[$i]{$prod}{$Strial} = ($costAvgCW * $CWsum[$i]) + ($costAvgADI *
$ADIsun[$i]) + ($costMissDem * $MD[$i]);
}
}

```

LIST OF REFERENCES

- Bowersox, D., and Closs, D. (1996). *Logistical Management: The Integrated Supply Chain Process*. McGraw-Hill/Irwin.
- Braun, M., Rivera, D. Flores, M., Carlyle, M., and Kempf, K., (2003). "A Model Predictive Control Framework for Robust Management of Multi-Product, Multi-Echelon Demand Networks," *Annual Reviews in Control*, Vol.27, No. 2, pp. 229-245.
- Bridge, R. Personal Communication. (2004, October 10). Intel Corporation, CPLG Finance.
- Buzacott, J. (1999). "Dynamic inventory Targets Revisited," *Journal of the Operational Research Society*, Vol. 50, pp. 697-703.
- Christopher, M. and Towill, D. (2001). "An Integrated Model for the Design of Agile Supply Chains," *International Journal of Physical Distribution and Logistics Management*, Vol. 13, No. 4 (2001), pp. 235-246.
- Christopher, M. and Towill, D. (2002). "The Supply Chain Strategy Conundrum: To be Lean or Agile or to be Lean and Agile", *International Journal of Logistics Management: Research & Application*, Vol. 5, No. 3, 2002.
- Dong, M. (2001). "Process Modeling, Performance Analysis and Configuration Simulation in Integrated Supply Chain Network Design," Doctoral Dissertation. Virginia Polytechnic Institute and State University.
- Dove, R. (1995). "Measuring agility: The toll of turmoil," *Journal of Applied Manufacturing Systems*, Vol. 7, No. 2, pp. 12-14.
- Dove, R. (1999). "Knowledge Management, Responsible Ability, and the Agile Enterprise". *Journal of Knowledge Management*. pp. 1-16.
- Dubelaar, Chow, and Larson (2001). "Relationships Between Retail, Sales, and Service in a Retail Chain Store Operation" *International Journal of Physical Distribution & Logistics Management*, Vol. 31, No. 2, pp. 96-108.
- Ettl, M., Feigin, G., Lin, G., and Yao, D. (2000, March/April), "A Supply Network Model with Base-Stock Control and Service Requirements", *Operations Research*, Vol. 28. No 2., pp. 216-232.
- Forrester, J. (1969). *Industrial Dynamics*, MIT Press, Cambridge.

- Giachettia, R., Martinez, L., Saenza, A., and Chena, C. (2003). *Analysis of the structural measures of flexibility and agility using a measurement theoretical framework*. International Journal of Production Economics. Vol. 86. pp. 47–62.
- Goldman, S., Nagel, R., and Preiss, K. (1994). *Agile Competitors and Virtual Organizations: Strategies for Enriching the Customer*. Van Nostrand Reinhold: New York, New York.
- Goranson, H. (2000). *The Agile Virtual Enterprise, Cases, Metrics, and Tools*. Quorum Books, Westport, CT.
- Hosmer, D. and Lemeshow, S. (2000). *Applied Logistic Regression*. 2nd Edition. Wiley and Sons.
- Inderfurth, K. and Minner, S. (1998). “Safety Stocks in Multi-stage Inventory Systems Under Different Service Measures”, *European Journal of Operational Research*, Vol. 106, pp. 57-73.
- Jain, A., and Dubes, R. (1988). “Algorithms for Clustering Data,” Englewood Cliffs, NJ: Prentice-Hall.
- Katayama, H., and Bennett, D. (1999). “Agility, adaptability and leanness: A comparison of concepts and a study of practice,” *International Journal of Production Economics*, Vol. 60–61. pp. 43–51.
- Kok, A. De and Verrijdt, J. (1995). “Distribution planning for a divergent N-echelon network without intermediate stocks under service restrictions.” *International Journal of Production Economics*, Vol. 38, pp. 25-43.
- Lagodimos, A. (1992). “Multi-echelon service models for inventory systems under different rationing policies”, *International Journal of Production Research*, Vol. 30, pp. 939-958.
- Lee, H., Padmanabhan, V, and Whang, S. (1997, April). “Information Distortion in a Supply Chain: The Bullwhip Effect.” *Management Science*. Vol. 43, No. 4.
- Lertpattarapong C. (2002). “Applying system dynamics approach to the supply chain management problem,” Master’s Thesis. Massachusetts Institute of Technology.
- Liu, L., Lui, X., & Yao, D. (2004). Analysis and Optimization of a Multistage Inventory-Queue System. *Management Science*, Vol. 50, No. 3, pp. 365-380.
- Lucas, J. and Saccucci, M., (1990). “Exponentially weighted moving average control schemes: properties and enhancements,” *Technometrics*, Vol. 32, No. 1., pp. 1-12.
- Mahoney, R. and Plossl, G. (1997). *High Mix Low Volume Manufacturing*. Prentice Hall.

- Maltz, A., Grenoble, W., Rogers, D., Baseman, R., Gray, W., and Katircioglu, K. (2000). Lessons from the Semiconductor Industry. *Supply Chain Management Review*. November/December, pp. 42-52.
- Mason-Jones, R. and Towill, D. (1999). "Total Cycle Time Compression and the Agile Supply Chain," *International Journal of Production Economics*, Vol. 62, No. 1, pp. 61-73.
- Mendenhall, W. and Sincich, T. (1995). *Statistics for Engineering and the Sciences*. Fourth Edition. Prentice Hall.
- Mitra, A. (1998). *Fundamentals of Quality Control and Improvement*. Second Edition. Prentice Hall.
- Montgomery, D. (1985). *Introduction to Statistical Quality Control*. New York: John Wiley.
- Naylor, J., Naim, M., and Berry, D. (1999). "Leagility: Integrating the Lean and Agile Manufacturing Paradigms into the Total Supply Chain," *International Journal of Production Economics*, Vol. 62, No. 1, pp. 107-118.
- Ortega, M. and Lin, L. (2004). "Control theory applications to the production-inventory problem: a review," *International Journal of Production Research*, Vol. 42. No. 11. pp. 2303-2322.
- Power, D. and Sohal, S. (2001). "Critical success factors in agile supply chain management: An empirical study," *International Journal of Physical Distribution & Logistics Management*, Vol. 31, No. 4, pp. 247-265.
- Ramasesh, R., Kulkarni, S., and Jayakumar, M. (2001). "Agility in manufacturing systems: an exploratory modeling framework and simulation," *Integrated Manufacturing Systems*, Vol. 12 No. 7, pp. 534-548.
- Rivera, D., Mittelman, H., Sarjoughian, H., and Kempf, K. (2005). "A Novel Model Predictive Control Algorithm for Supply Chain Management in Semiconductor Manufacturing," 2005 NSF DMII Grantees' Conference, Scottsdale, Arizona.
- Roberts, S. (1959). "Control Chart Tests Based on Geometrically Weighted Moving Averages," *Technometrics*, Vol. 1, pp. 239-250.
- Roberts, S. (1966). "A Comparison of Some Control Chart Procedures," *Technometrics*, Vol. 20, pp.85-93.
- Sabri, E. and Beamon, B. (2000). "A Multi-Objective Approach to Simultaneous Strategic and Operational Planning in Supply Chain Design," *Omega*, Vol. 28, No. 5, pp. 581-598.

- Sarkis, J., Presley, A., and Liles, D. (1995). "Building a Business Case for Agility". *4th Agility Forum Conference*, Atlanta, GA.
- Schwarz, L., Deuermeyer, B., and Badinelli, R. (1985). "Fill-rate optimization in a one-warehouse N-identical retailer distribution system," *Management Science*, Vol. 31, pp. 488–498.
- Sethi, S., and Thompson, G. (2000). *Optimal control theory: applications to management science and economics*. Kluwer Academic Publishers.
- Shaw, A, McFarlane, D, Chang, Y, and Noury, P. (2002). "Measuring Response Capabilities in the Order Fulfillment Process," *Proceedings 9th International Conference European Operations Management Association (EurOMA)*, Copenhagen, Denmark.
- Slack, N. (1983). "Flexibility as a Manufacturing Objective", *International Journal of Operations and Production*, Vol. 3, No. 3, pp. 4-13.
- Stamatis, D. (2003). *Six Sigma and Beyond: Statistics and Probability*, Vol. III, St. Lucie Press.
- Sonnet, C. (2004). "Cost of Stock-outs in the Microprocessor Business and its Impact in Determining the Optimal Service Level," Master's Thesis. Massachusetts Institute of Technology.
- Swafford, P., Gosh, S., and Murthy, N. (2001). "Development of Flexibility and Agility Constructs for Supply Chain Agility Model," *Proceedings of the Decision Sciences Institute*, 2001.
- Takahashi, K. and Nakamura, N. (1999). "Reacting JIT Ordering Systems to the Unstable Changes in Demand," *International Journal of Production Research*, Vol. 37, No. 10, pp. 2293-2313.
- Towill, D., (1993). "System Dynamics-background, methodology and applications", *Computing and Control Engineering Journal*, Vol. 4, No. 6, pp. 261-268
- Towill, D. and McCullen, P. (1999). "The impact of an agile manufacturing programme on supply chain dynamics," *International Journal of Logistics Management*, Vol. 10, No. 1, pp. 83-96.
- Wang, W., Rivera, B., and Kempf, K. (2003). "Centralized Model Predictive Control Strategies for Inventory Management in Semiconductor Supply Chains," *Proceedings of the American Control Conference*, Denver, CO.

Wang, W., Rivera, D. and Kempf, K. (2004). "A Novel Model Predictive Control Strategy for Tactical Decision-Making in Semiconductor Manufacturing Supply Chain Management," paper 421e, presented at the 2004 AIChE Annual Meeting, Austin, TX.

Weber, M. (2002). "Measuring supply chain agility in the virtual organization" *International Journal of Physical Distribution & Logistics Management*, Vol. 32, No. 7, pp. 577-590.