Using Discrete Event Simulation to Model the Benefits of VMI and CPFR

by

Amanda Cha

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Management Sciences

Waterloo, Ontario, Canada, 2013

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Collaborative Planning, Forecasting, and Replenishment or CPFR, is a cooperative business methodology where supply chain members exchange demand information and develop a single shared forecast. CPFR promises to improve demand forecast accuracy, reduce inventory levels, and improve fill rates. Many organizations, including Wal-Mart, Michelin, and Heineken have successfully utilized CPFR to reduce their costs, lower their levels of inventory, and improve their fill rates. With advance notice of promotions or new product introductions, members of the supply chain can plan their own replenishment and manufacturing activities accordingly, and reduce their reliance on higher levels of safety stock.

Although there have been many successful CPFR pilot programs, few large scale implementations of CPFR can be found and some case studies have reported disappointing results. To determine when CPFR will deliver on its promises, a simulation study of a three-stage supply chain was devised. CPFR was compared to Vendor Managed Inventory (VMI), another popular information sharing supply chain methodology, and Independent Sourcing, where no information was shared and the supply chain members acted independently. A variety of demand patterns were tested, including steady demand and demand with promotions. The simulation was first tested using hypothetical data, then run with demand data provided by 3M, a large, conglomerate corporation.

The simulation results showed that when the supply chain members of VMI and CPFR had access to the same information, the two methodologies performed comparably. When promotions were not present, the information shared in CPFR was similar to the information shared in VMI and thus, there was no statistically significant difference between the performances of VMI and CPFR. When the supply chain members of CPFR were privy to information not shared in VMI, as was the case when promotions were present, CPFR had lower costs and inventory levels than

VMI. When promotions were planned by the retailer, their timing was only shared with the vendor in CPFR, and not with the vendor in VMI. To achieve the desired fill rates, the vendor in VMI held more inventory and therefore, incurred higher costs than CPFR.

While VMI and CPFR are easily differentiated in literature, in practice, VMI implementations can have aspects of CPFR, and vice versa. Our research has revealed that complete information sharing is of the utmost importance. When crucial information is withheld from supply chain partners, the ability of CPFR or VMI to reduce costs and inventory levels greatly diminishes. When working with incomplete information, supply chain members carry higher inventory levels to compensate for uncertainty.

Acknowledgements

I would like to thank my supervisor Professor Jim Bookbinder for his guidance and assistance throughout my Masters. Your expertise and immense knowledge were crucial to this thesis and I learnt a great deal from you. Thank you for your patience, your time and your encouragement.

I would like to also thank Paul Lambrinos from 3M for sharing his expertise and providing us with data for our research. Your valuable insights greatly improved the quality of our work. Thank you for taking the time to speak with us about your experiences in CPFR.

For his everlasting support, I would like to thank my wonderful husband Jason. Thank you for encouraging me to pursue a Masters degree and for being my personal cheerleader every step of the way. I would also like to thank my parents, Gil and Tracy for all they have given me. You both made my education your priority and I can't thank you enough for the sacrifices you both made so that I could pursue my passions. Stephen and Andrew, thank you for being my brothers, even though I know you didn't have a choice.

Finally, I would like to thank my friends for keeping me sane during my Masters. I could not have done it without you girls.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

1.1. Definition of Collaborative Planning, Forecasting and Replenishment

Collaborative Planning, Forecasting, and Replenishment (CPFR) is a cooperative business methodology where two or more supply chain members exchange information that can impact future demand and together, develop a single shared forecast. Supply chain firms trade promotional schedules and plans for new products and new stores to facilitate the creation of the joint forecast. By combining the intelligence of all participating firms, CPFR promises to improve demand forecast accuracy and as a result, reduce inventory levels, improve fill rates, and decrease lead times. Production schedules, replenishment policies, and other supply chain activities are then planned to meet the needs of the single forecast.

Developed by the Voluntary Inter-industry Commerce Standard Association (VICS) in 1998, CPFR is comprised of nine steps. The first step is to establish a collaborative relationship between the parties involved. Guidelines, rules, expectations, goals, and information sharing agreements are defined in this stage. Next, a joint business plan is created by exchanging and combining business plans and promotional schedules. The joint business plan involves formulating order minimums, lead times, and order intervals. Taking into account the information gathered by the joint business plan, each participating firm creates a sales forecast. The forecasts are compared and any differences, or "exceptions", are identified and resolved through discussion, and then agreement. This creates a single sales forecast that is used by all parties involved. This sales forecast is then utilized to create an order forecast by each participant. Once again, the individual order forecasts are compared. If exceptions are found, the supply chain members work together to solve them, which creates a shared order forecast. The firms use the order forecast to plan their supply chain activities so that upcoming orders can be filled. Manufacturers can plan their production schedules, vendors can determine what levels of inventory to hold, and both can plan their transportation needs. In the final step, the combined order forecast is turned into an order, which is placed by the retailer to the vendor or manufacturer (Harrington, 2003).

The reported benefits of CPFR are attractive to any firm wishing to improve the efficiency of their supply chain operations. Companies that have implemented CPFR pilot programs have experienced lower costs and higher service levels, which ultimately results in increased profit. Examples of these companies include Wal-Mart, Sara Lee, Heineken, Sears, and Michelin (Aviv, 2001; Steermann, 2003). Many organizations are turning to CPFR to improve forecast accuracy, as poor accuracy can result in either stockouts or excess inventory, each with its own cost consequence. Stockouts can result in lost sales, while excess inventory can lead to higher inventory carrying costs.

CPFR has also been said to alleviate the effects of the bullwhip effect (Disney et al., 2004). The bullwhip effect is a phenomenon where order information becomes distorted as it moves up the supply chain. The effect is caused by supply chain members misinterpreting information gleaned from an order (Lee et al., 1997). One of the proposed methods to combat the bullwhip effect is to share forecast information with upstream supply chain members, which is a part of the CPFR process. By sharing the demand forecast, all supply chain members can plan their own replenishment activities based on the end-customer demand forecast.

Despite many reports of reduced costs, inventory, and stockout rates, not all companies have found CPFR to be successful. Smaros (2007) gave several examples of pilot programs where CPFR did not deliver on its promises. Also, while many CPFR pilot programs exist, few large scale CPFR implementations can be found. Since CPFR can be time consuming and costly to set up, retailers and vendors are interested in determining if situations exist when implementing CPFR may not be beneficial, especially if an information sharing supply chain strategy, such as Vendor Managed Inventory (VMI), is already is in place.

1.2. Research Objectives and Approach

The objective of this thesis was to determine when the benefits of CPFR can be realized. Does CPFR always lead decreased inventory levels and reduced costs? Is CPFR recommended if a VMI program is already in place? What advantages does CPFR provide over VMI, if any? We attempted to answer these questions by using discrete event simulation to model the three-stage supply chain of a single, non-perishable, consumer item. The supply chain was composed of a single retailer, a single vendor, and a single manufacturer. A CPFR implementation between the retailer and vendor was compared to a VMI program between the retailer and vendor and to Independent Sourcing (IS), where no information was shared amongst the supply chain members. We programmed the simulation model using Crystal Ball, a simulation software from Oracle that utilizes Microsoft Excel.

The retailer in our simulation faced weekly stochastic demand with two types of patterns to determine how different types of demand could impact the success or failure of CPFR. In the first demand pattern, no promotions were present and demand was steady. In the second, promotions initiated by the retailer were present. During promotions, demand spiked depending on the attractiveness of the sale and the aggressiveness of the marketing campaign. The simulation model was first programmed using a hypothetical set of data, created to be representative of empirical sales data. Different levels of demand variability and target fill rates were tested to examine if and how they affected the performances of IS, VMI, and CPFR. Cost and inventory results were gathered from the simulation at the system-wide level, at the retailer level, and at the vendor level to compare each methodology's effectiveness. The simulation model was then run using empirical sales data provided by 3M, a large, conglomerate corporation. Sales data for items in four product groups, stationary, home care, first aid, and hardware, was provided by 3M. By running the simulation with a variety of demand patterns and cost and lead time parameters, we investigated how IS, VMI, and CPFR operated under different conditions and if each methodology's effectiveness would be affected.

1.3. Thesis Outline

Chapter two presents a literature review of CPFR starting with details about CPFR's predecessors and their strengths and weaknesses. We discuss the first CPFR pilot program undertaken by Wal-Mart and its supplier Warner Lambert, and its results. This is followed by a review of other CPFR implementations and pilot programs, both successful and unsuccessful, and details of either their accomplishments, or their failures. Finally we look at existing research on CPFR and examine the approaches various authors took to validate the claims of CPFR. We focus on comparisons to other supply chain methodologies and insights into what conditions are necessary for CPFR to succeed.

The third chapter of this thesis outlines how our simulation model was created using Crystal Ball and Microsoft Excel. Our definitions of IS, VMI, and CPFR are explained since various definitions for each methodology have been found. We describe how the demand and costs parameters were calculated, outline the characteristics of end-customer demand, and list the assumptions made. We also discuss the methods used to verify and validate the simulation models.

The results from the simulation described in chapter three are outlined in chapter four. Averages and confidence intervals of cost, inventory, and stockout results for the various cases tested are presented here. First, the results from the simulations using hypothetical data are detailed, then the results from each of the cases tested in simulations using empirical data.

Analysis and discussion of the results can be found in chapter five. We examine the implications of the results from simulations without promotions and the simulations with promotions. Results from both simulations using hypothetical data and simulations using empirical data are discussed. Finally we give our conclusions and recommendations which follow from our results and discussion.

Chapter 2: Literature Review

2.1. Predecessors to CPFR

2.1.1. Efficient Consumer Response (ECR)

The Efficient Consumer Response (ECR) Movement Group was developed by 14 trade association sponsors in 1992. Their purpose was to encourage the integration of supply chain members by building relationships and trust, with the aim of achieving better supply chain results. Four core principles, efficient promotions, efficient replenishment, efficient store assortment, and efficient product introductions, were defined by ECR as methods to optimize supply chain performance. By sharing strategic information with trading partners, ECR states that supply chain inefficiencies will be reduced, leading to lower inventory levels, quicker response times, and decreased costs (Barratt & Oliveira, 2001).

2.1.2. Vendor Managed Inventory (VMI) and Continuous Replenishment (CR)

Vendor Managed Inventory (VMI) is an information sharing supply chain methodology that can be used to implement the principles of ECR. Developed in the mid-1980s, it gives to the vendor, the responsibility of managing the retailer's inventory. The vendor determines what level of inventory the retailer should carry, and replenishes it as needed. Similar to VMI is Continuous Replenishment (CR), where the vendor is also given access to the retailer's point-of-sales data (POS). With this data, the vendor is expected to create a forecast which will help them determine an inventory policy on behalf of the retailer. By controlling replenishment, the vendor reduces demand uncertainty since they no longer solely rely on order information to make decisions. Companies that have implemented VMI have found that this results in a smoother demand signal and the need for safety stock is reduced. Vendors are also able to reduce transportation costs, as they no longer need to ship products when orders are placed. Instead, vendors can plan their replenishment to utilize full truckload shipments rather than more expensive, less-than-truckload shipments (Waller et al., 1999).

Despite these reported benefits, both VMI and CR have weaknesses. Both are poor at handling promotion since access to demand information is limited, and vendors are expected to interpret the inventory data they have access to without assistance from the retailer. In VMI, vendors only gain access to information about stock levels in the retailer's distribution center. Vendors gain additional access to information in CR, but POS data is historical, and does not assist in predicting when future promotions will occur. To prevent stockouts from occurring during promotions, vendors participating in VMI and CR found that they were forced to carry large amounts of inventory (Hill, 1999; Barratt & Oliveira, 2001).

CPFR promises to address the issues experienced with VMI and CR. Since sales forecasts are jointly created and shared, vendors are no longer surprised by upcoming promotions, and the risk of misinterpreting order information is eliminated. Events other than promotions that can influence sales are also shared, such as weather forecasts, new product introductions, assortment changes, and marketing campaigns. Any discrepancies between the retailer's and vendor's forecasts are discussed and resolved (Aichlmayr, 2000).

2.2. Successful CPFR Pilot Programs

In 1995, Wal-Mart was the first company to implement a CPFR-like pilot program with supplier Warner-Lambert, IT companies SAP and Manugistics, and consulting firm Benchmarking Partners. The group called this predecessor to CPFR, Collaborative Forecasting and Replenishment (CFAR), and their goal was to lower the inventory levels of Warner-Lambert's Listerine mouthwash products. Wal-Mart and Warner-Lambert shared information about their upcoming promotional activities and compared their sales and order forecasts to find and resolve any discrepancies on a weekly basis. Prior to the pilot program, Warner-Lambert was not given advance notice of Wal-Mart's planned promotions, which were capable of creating considerable increases to customer demand. To ensure they had enough stock for any potential promotions, Warner-Lambert had resorted to keeping a substantial amount of safety stock. With CFAR, Warner-Lambert knew the promotional schedule in advance, and was able to adjust their production schedule to manufacture according to customer demand, reducing their reliance on safety stock. Wal-Mart experienced a 25% drop in their inventory levels and in-stock averages grew from 85% to 98% which resulted in an $8.5 million increase in sales (Barratt & Oliveira, 2001; Seifert, 2003).

Since Wal-Mart's first successful CPFR implementation, many successful pilot programs have taken place and a variety of benefits have been reported. The majority of those programs have taken place in consumer product industries and some of their results are outlined below.

Dutch brewer Heineken was one of the earliest CPFR users, starting their program in 1996, with their North American distributors. By collaborating with their distributors on sales estimates, Heineken cut forecast error by 15%. The program also resulted in a 50% reduction in order-cycle time (Hill & Mathur, 1999; Aviv, 2001).

Wal-Mart also implemented CFAR with supplier Sara Lee Branded Apparel for its Hanes underwear products in 1998. The six month program involved 50 SKUs and 2500 Wal-Mart stores. It resulted in a 14% reduction of store-level inventory, a 2.7% improvement of in-stock levels, and a 30% rise in retail turns. While sales for the underwear category increased 35% over the six month period, sales for Hanes products increased 45% (Ireland & Bruce, 2000; Aviv, 2001).

American supermarket chain Wegmans Food Markets partnered with supplier Nabisco for a six month CPFR trial in 1998. The trial was split into two phases with the first focused on Nabisco's Planters Nuts products. The trial had positive outcomes; sales for Planters Nuts grew by 53.9% while category sales only grew by 16.3%. Days-on-hand inventory also decreased 18% and service level to stores rose 4% to 97% over the course of the trial. The second phase involved Nabisco's Milkbone products. Again the outcome was positive; while sales for the category were 7% higher, sales for Milkbone products increased 8% (Ireland & Bruce, 2000).

Retailer Ace Hardware and home and office product manufacturer Manco started a CPFR pilot program in 2000 to replace their existing VMI system. The VMI system was not meeting the expectations of the firms and both hoped to improve the speed and agility of the supply chain with CPFR. The two firms shared their sales goals and promotional and seasonal merchandise plans to jointly create a single sales forecast. Not only did Ace Hardware and Manco collaborate on the sales and order forecast, but also on product assortment and space planning for the stores. The pilot program was a success; sales increased by 9% and forecast accuracy rose by 10%. By analysing the order flow with Ace Hardware, the partners were able to attain better shipping economies, cutting distribution costs by 28% and freight costs by 18% (Seifert, 2003).

Motorola began their CPFR pilot program in 2001 to improve the performance of their mobile phone handset supply chain. The firm worked with their retailers to eliminate forecast error, reduce inventory, and improve on-time delivery. Motorola experienced high levels of forecast error, which lead to stockouts and lost revenue before implementing CPFR. Forecasting demand was difficult, as the company produced over 120 models globally and new models were continuously being introduced. The mobile phone industry also experienced short life-cycles, with the average life cycle lasting just over one year. While Motorola was well informed about their customers, their visibility of the retailer's market was limited. The company was only aware of the shipments they sent to their retailer's distribution centers. By implementing the CPFR pilot program with their retailers, Motorola was able to gain visibility of the end consumer and as a result, improve their forecast accuracy. Safety stock levels decreased and stockout rates dropped by two-thirds. Higher forecast accuracy allowed Motorola to better plan their production schedule and transportation, which led to a 50% drop in transportation costs (Cederlund et al., 2007).

American retailer Sears and French tire manufacturer Michelin started their CPFR program in 2001. Previous attempts by the two firms to boost fill rates had caused inventory levels at Sears' distribution centers and Michelin's warehouses to rise. Despite higher levels of inventory, fill rates were still falling short of acceptable levels. Through CPFR, the two companies hoped to reduce the occurrences of stockouts at Sears' retail stores while lowering the levels of inventory in their warehouses and distribution centers. The firms also hoped that CPFR would expand demand visibility throughout the supply chain. Using the VICS CPFR method, Sears and Michelin implemented the program on 80 SKUs, with positive results. Sears and Michelin experienced a 25% fall in their inventory levels and in-store, in-stock levels rose by 4.3%. Sears also saw a 10.7% improvement to their DC-to-retail store fill rate. Sears and Michelin found that CPFR was especially beneficial during product transitions. Greater demand visibility had allowed the companies to better plan promotions and markdowns that helped sell older products, with the purpose of making room for newer products. The CPFR program was eventually expanded to include all of Sears' Auto Centers and National Tire and Battery locations and to 220 of Michelin's SKUs (Steermann, 2003).

American boating supplies retailer West Marine partnered with 150 of their suppliers in 2002, with the goal to recover the coordination of their supply chain planning and replenishment activities. After their acquisition of competitor E&B Marine in 1996, West Marine suffered from inaccurate demand forecasts, which led to poor in-stock levels, causing an 8% drop in sales in just one year. By working with their suppliers, West Marine was able to achieve 85% forecast accuracy and their in-store, in-stock rates grew to 96% (Denend, 2005).

Chung and Leung (2005) studied a CPFR pilot program in a Chinese copper clad laminate company, and examined if CPFR could be successfully implemented by the manufacturing company. The copper clad laminate industry is very different than the grocery industry, where the majority of CPFR pilot programs occur. Unlike the grocery industry, sales are made directly to the customer and ERP software is not used. Copper clad laminate is used in the manufacturing of electronic products which have short life cycles. Having high levels of inventory will lead to obsolete scrap. The company, referred to as "MA", was suffering from poor order visibility and wanted to decrease out-of-stock occurrences and inventory levels. The pilot program was implemented with supplier "JA" and resulted in a fall in lead time from 5 weeks to 2 weeks. Inventory levels and out-of-stock occurrences also fell.

2.3. Journal Articles Examining the Benefits of CPFR

2.3.1. Examining CPFR Exclusively

Raghunathan (1999) was one of the first to perform an analysis of the benefits of CPFR, or CFAR as it was known at that time, with an analytical model . The research aimed to determine who benefited from implementing the methodology, manufacturers or retailers, using a modified inventory theoretic model (Min & Yu, 2008). In the model, Raghunathan assumed that demand was stationary, the manufacturer did not have capacity constraints, and prices at the retailer were constant. Three cases were examined, each containing one manufacturer and two independent retailers in a two-tier supply chain. In the first case, neither retailer participated in CFAR. Orders were placed to the manufacturer and if the orders could not be fully filled, shortages were equally divided between the two retailers. In the second case, one retailer participated in CFAR and provided the manufacturer with demand information. It was assumed that the retailer had better knowledge of customer demand than the manufacturer. In exchange for the demand information, the manufacturer promised the participating retailer that their orders would be filled. If any shortages in inventory occurred during a demand period, the participating retailer's order would be filled first. Only the non-participating retailer would experience a shortage during that time. In the third case, both retailers participated in CFAR and both provided the manufacturer with demand information. If any shortages in supply occurred, they were equally divided amongst the two participating suppliers. However, since the retailers were better able to predict customer demand, the manufacturer could use the supplied demand information to more accurately plan their production. Since demand uncertainty was reduced, if shortages did occur, they were of smaller magnitude than in the first case.

Raghunathan found that the manufacturer would benefit as long as at least one retailer participated in CFAR. Participating in CFAR also reduced costs for the retailers as their demand was better met. In the case where only one retailer participated in CFAR, the non-participating retailer faced negative consequences for not participating. If any shortages occurred, they were only experienced by the non-participating retailer.

Aviv (2001) examined the benefits of collaborative forecasting in CPFR by developing a mathematical framework for a two-tier supply chain. Three forecasting methods were compared in the supply chain consisting of a single retailer and single vendor. In the first method, called local forecasting, each supply chain member updated their own forecast as new demand information became available and incorporated that forecast into their own replenishment process. Local forecasting was a decentralized information structure where inventory levels and forecasts were only known locally. The second forecasting method involved centralizing the forecasts. Called collaborative forecasting, supply chain members shared and managed a single forecast. Both methods were compared to a benchmark, where forecasts were not integrated into the replenishment process. In his paper, Aviv assumed that the supplier and retailer were cooperative, that it, they set their inventory and order policies to optimize the overall cost of the supply chain rather than minimize their own costs. Aviv found that collaborative forecasting was more beneficial than local forecasting when the supplier and retailer had different forecasting abilities. If the supplier and retailer had the same forecasting abilities, the two forecasts created by local forecasting would not differ from the one created by collaborative forecasting. Collaborative forecasting with different forecasting abilities would create a forecast that combined the abilities of the supplier and retailer, thus creating a more accurate demand forecast.

Boone et al. (2002) used data from a Fortune 500 consumer products company to compare the performance of CPFR against a traditional reorder point system using a simulation model. The performance was judged based on two key performance measures, consumer service or fill rates, and supply chain cycle time. In the traditional reorder point system, the manufacturer did not know the needs of the retailer, while in CPFR, the manufacturer did. Boone et al. examined several test cases, varying forecast errors, service levels, transportation modes, safety stock levels, and demand levels. The simulation showed that CPFR would improve fill rates, lower inventory levels, and decrease cycle times. Supply chain cycle times fell because the product was spending less time as inventory, due to reduced inventory levels. The different test cases also shed light on when the benefits of CPFR would be greater. High forecast error and high demand levels resulted in greater CPFR benefits when compared to the traditional reorder point system.

Aviv (2007) furthered his previous research by examining collaborative forecasting in CPFR in a two-stage supply chain, with a manufacturer and a retailer. However, in this paper, supply chain members were not cooperative, that is, they did not set their inventory policies to maximize overall supply chain performance, but according to specified performance metrics. While the manufacturer was concerned with production smoothing, inventory levels, and production schedule stability, the retailer was concerned with inventory levels and service rates. Aviv found that the benefits of collaborative forecasting depended on the relative explanatory power of the supply chain member, which he defined as the ability of that member to anticipate demand. Collaborative forecasting benefits grew when the manufacturer had the highest relative explanatory power. This occurred because the retailer's demand information is somewhat already shared with the manufacturer through their orders. Supply side agility, the ability of the supplier to act upon new demand information, was also deemed to be important.

2.3.2. Comparing CPFR and VMI

Disney et al. (2004) compared the performance of a traditional supply chain, a VMI supply chain, an information sharing supply chain (EPOS), and a collaborative, CPFR-like, supply chain supply using a management game. Created in the 1950s by MIT, the Beer Game was played by student teams. Disney et al. simulated several different supply chain methodologies and ran the game for 25 time periods. The performance of each method was measured by inventory holding costs and the effect of the bullwhip effect in each supply chain. The results showed that the CPFR-like methodology was best at reducing the bullwhip effect in the supply chain, while VMI performed the worst in inventory holding costs and in bullwhip effect. EPOS also performed poorly, which led to the conclusion that while the students all had access to the same information, they were each interpreting it differently, leading to stockouts.

Aviv (2002) also compared a traditional supply chain to VMI and CPFR in a two-stage supply chain. Aviv used an auto-regressive time series model to determine under which circumstances each methodology would provide the greatest amount of benefits. In the traditional case, the retailer and supplier were cooperative, that is, they coordinated their policy parameters to minimize costs in the overall supply chain. While both the supplier and retailer observed market signals that enabled them to explain future demand, they did not observe the same signals and did not share their observations with each other. In VMI, the retailer's POS data was shared with the supplier. However, the POS data did not contain all information relevant to future demand, such as promotional schedules. The supplier made supply chain decisions based on partial data. In CPFR, the inventory was managed centrally, and the market signals observed by the retailer and supplier were shared with one another. Compared to the traditional supply chain and VMI, CPFR required more investment. The participating firms needed to commit organizational resources and had to be willing to share information with each other. Aviv found that the benefits of VMI and CPFR depended on the relative explanatory power of either the supplier or the retailer. When the supplier had large relative explanatory power, that is, the market signals observed directly by the supplier were more informative than those observed directly by the retailer, VMI provided greater benefits. CPFR would provide more benefits than VMI when the retailer had greater explanatory power than the supplier.

Sari (2008a) compared the benefits of CPFR, VMI and a traditional supply chain using a Monte Carlo simulation model. All three methodologies were examined using a four-stage supply chain under stationary and non-stationary demands. The four echelons were a manufacturing plant, a warehouse, a distributor, and a retailer. In the traditional supply chain, each echelon formed their policies to optimize their own costs, and the decisions were based only on order information. In VMI, the distributor had access to the retailer's real-time inventory levels and POS data. The remainder of the supply chain operated in the traditional supply chain method. In CPFR, inventory levels, POS data, promotional plans, and sales forecasts were shared with all four levels of the supply chain. Sari examined demand seasonality, lead times, customer demand uncertainty, and production capacity to determine how these factors affected the benefits of CPFR and VMI. The benefits were measured using the customer service level of the retailer and the total cost incurred by the supply chain. The simulation showed that the benefits of CPFR were always greater than those of VMI. The gap between the benefit levels were especially noticeable when demand variability was high. When lead times were short or when manufacturing capacity was tight, the gap between the benefits of CPFR and VMI was smaller. In these cases, Sari concluded that the additional benefits provided by CPFR did not justify the additional resources needed for CPFR over VMI.

2.4. Unsuccessful CPFR Implementations

Despite many successful CPFR pilot programs and research supporting the claims of CPFR, not all pilot programs have been successful and few large scale implementations can be found (Smaros, 2007). Smaros presented case studies of four collaborative forecasting pilot programs that took place in Europe and found that only three of the four were deemed successful. The pilot programs were run by a European grocer and four of its suppliers, referred to DairyCo, MeatCo, CandyCo, and ChemCo. The grocer's identity was also hidden and referred to as RetailCo in the paper.

RetailCo and DairyCo implemented a collaboration method similar to CPFR with the goal of improving store-level forecast accuracy. Though retail-level forecast accuracy improved, store-level forecast accuracy did not. RetailCo found the process arduous and determined that significant investment in both manpower and technology would be required to implement the program on a larger scale.

RetailCo and MeatCo wished to improve forecast accuracy by collaboratively forecasting promotions in their partnership. To facilitate this, historical promotional data, such as sales that occurred prior to, throughout, and following a promotion period, were analyzed. MeatCo and RetailCo did not find this project helpful in increasing forecast accuracy. RetailCo's managers were not as experienced in forecasting as the MeatCo managers and only a small amount of historical promotion data was made available to be examined.

RetailCo and CandyCo's collaborative program was centered around new product introductions. By giving CandyCo access to RetailCo's POS data, the duo hoped to improve CandyCo's sales forecast accuracy. CandyCo found that POS data was helpful, in that it allowed them to determine which products were at risk for stockout and which products were not selling as expected. Normally, CandyCo would not gain these insights until an order from RetailCo was placed. From the information gleaned from the POS data, CandyCo was able to update their production schedules and thus, improve forecast accuracy by 7% and service levels by 2.6%.

A similar project was undertaken by RetailCo and ChemCo. However, unlike CandyCo, ChemCo was a large multinational firm with long production lead-times. Despite having access to POS data for new products, ChemCo did not believe gaining sales information a few weeks earlier would make a difference to the forecast accuracy.

Smaros concluded from these case studies that in practice, the forecasting abilities of the retailer can severely hinder collaborative forecasting. In the cases with DairyCo and MeatCo, RetailCo's lack of forecasting experience hindered the success of the pilot programs. Smaros also pointed out that retailers and manufacturers have different planning horizons and aggregation levels which make combining forecasts difficult. Should the manufacturer have a long production interval, integrating up-to-date demand information may not be possible, as was the case with ChemCo.

2.5. Factors that Impact the Success or Failure of CPFR

2.5.1. Empirical Observations of the Key Success Factors of CPFR

Barratt and Oliveira (2001) surveyed supply chain practitioners that were knowledgeable in CPFR to determine what the inhibitors of success were, in implementing CPFR. The survey revealed that the chief inhibitor was when visibility into the supply chain was limited. This could occur when there was a lack of trust, minimal information sharing, poor communication of demand forecasts to supply chain members, and insufficient technology. The presence of any of these factors could result in an unsuccessful CPFR implementation.

Smaros (2007) concluded from the four case studies examined, that collaborative forecasting methods like CPFR do not provide benefits in every circumstance. The value of collaborative forecasting was diminished in the case of ChemCo, due to their extremely long lead times. RetailCo's limited forecasting abilities hindered their ability to create a more accurate demand forecast with MeatCo.

Fu et al. (2009) used the fuzzy analytic hierarchy process to rank the factors that impacted the success or failure of a CPFR program. Ten experts, all of whom had experience implementing CPFR, were surveyed in three areas, technology, organization, and environment, for the study. Fu et al. found that supply chain members must share the same goals for collaboration for CPFR to be successful. Coordinated communication was also revealed to be important. Various departments within each participating organization, marketing, finance, and procurement, must all work together to ensure CPFR's success. Similar to the results of Barratt and Oliveira's research, the study revealed that trust between supply chain partners was necessary for CPFR to succeed.

Danese (2010) examined the collaborative planning programs implemented by a variety of companies across Europe. Data was collected by conducting interviews, making direct observations, and examining documentation at ten companies. Danese was especially interested in determining why each firm chose the collaboration method they did. The various methods could be categorized based on their level of integration and multiplexity. Low integration was characterized by simple data exchange, whereas high integration involved synchronizing information and joint planning. Multiplexity was defined as the number of business areas that were involved with the collaboration program.

Danese called the lowest level of collaboration, communication. At this level, firms had low levels of integration, where only data was exchanged. According to Danese, VMI and CR fit into this category. The next level was called limited collaboration, and it featured higher levels of integration, but lower levels of multiplexity. In this category, data was exchanged and firms worked together to create a joint order forecast. At the highest level of collaboration, called full collaboration, there were high levels of both integration and multiplexity. The firms exchanged data and worked together to create many plans, including sales, order, and promotional plans.

In examining the characteristics of the ten firms, Danese discovered that there were several common factors that could explain why each firm had chosen the method of collaboration it had. These factors included collaboration goals, demand elasticity, and product diversity.

Danese defined two core goals of collaboration, efficiency and responsiveness. When the goal of the collaboration was to improve efficiency, firms tended towards the lower levels of collaboration. However, when responsiveness was the goal, higher levels of collaboration were implemented. Danese concluded that higher levels of integration were necessary to facilitate higher levels of responsiveness.

Danese also found a link between the demand elasticity of a product and the level of collaboration. Firms that sold products with low demand elasticity, that is, during a promotion, the average increase in the sales volume was under 40%, lower levels of collaboration were utilized. When demand elasticity was high, that is, the products experienced sales volume increases greater than 200% during promotions, businesses were more likely to participate in higher levels of collaboration.

Product diversity was also found to affect the level of collaboration employed by firms. Product diversity was defined by companies either selling the same product or different products. For example, a tool manufacturer that produces hammers and a hardware retailer will experience low product diversity, since both companies sell the same product, the hammer. A tool manufacturer that makes hammers and its packaging supplier experience high product diversity, as one sells the hammer and the other sells plastic packaging. Danese found that collaborating firms with low product diversity, implemented higher levels of collaboration. Danese hypothesised that this was due to the inability of firms with high product diversity to collaborate on promotional and sales forecasts.

2.5.2. Research into the Key Success Factors of CPFR

Sari (2008b) examined how data errors in inventory systems could impact the benefits of CPFR. Using the Monte Carlo simulation model from earlier work (2008a), Sari introduced inaccuracy in the three inventory models by using an error factor in the inventory levels. At the retailer level, inventory and sales numbers were adjusted by multiplying the true values with random values. Inventory and shipment quantities were altered at the distributor and warehouse level, while inventory, production, and shipment quantities were changed at the manufacturer. The error levels, the manufacturing capacity, and the lead times were all varied to determine how these factors would affect the outcome. Sari found that inaccurate inventory records had a significant negative effect on the performance of a supply chain, especially if the members collaborated. Inaccurate records had a greater impact on CPFR than VMI, due to the lesser amount of safety stock used in CPFR. Sari also found that short lead times and low demand uncertainty increased the negative impact of information inaccuracy. The benefits of CPFR and VMI could drop to insignificant levels under these circumstances.

Chapter 3: Model Description

3.1. Discrete Event Simulations

In this chapter, we will discuss our approach to modeling CPFR. We will outline how we defined VMI and CPFR and the steps we took to create a simulation model to evaluate their performances. To compare the abilities of VMI and CPFR, discrete event simulation models were programmed in Oracle's Crystal Ball software. Crystal Ball, a popular simulation software, was selected due to its easy integration into Microsoft Excel. Along with VMI and CPFR, Independent Sourcing (IS) was also programmed to serve as a base case. A three-tier supply chain consisting of a retailer, a vendor, and a manufacturer was modeled on a weekly basis in the simulations.

While stockouts at the retailer level resulted in lost sales, stockouts at the vendor level were backordered. The vendor placed orders with the manufacturer, who had an unlimited supply and was able to fulfill any order the vendor placed. This assumption was made to "end" the supply chain. That is, rather than have the manufacturer experience stockouts and need to wait for supplies from its own supplier, who in turn could also experience stockouts, and so on, and so forth, the supply chain ended with the manufacturer who we assumed could fulfill the vendor's orders with an adequate lead time. Despite having unlimited supply at the manufacturer, the vendor could still experience stockouts if its forecasts were incorrect.

3.2. Model Conceptualization

3.2.1. Definition of Independent Sourcing (IS)

In IS, the retailer and vendor do not share any information. The vendor can only base its own replenishment decisions on orders it receives from the retailer. These orders are placed intermittently, which makes it more difficult for the vendor to forecast. The vendor is also unable to distinguish between orders placed for regular demand and orders placed for an upcoming promotion. This may lead the vendor to believe that demand is increasing when in reality, demand will return to its previous characteristics after the promotion period has ended.

Figure 3-1: Flow of Information through the Supply Chain in IS

3.2.2. Definition of Vendor Managed Inventory (VMI)

In VMI, the vendor places orders on behalf of the retailer and in exchange, the vendor is given access to the retailer's POS and inventory data. Since the vendor controls the retailer's replenishment schedule, the vendor can determine its own replenishment schedule accordingly. The retailer does not share all relevant demand information with the vendor however. While the vendor is aware that promotions occur, their exact timing is not revealed by the retailer.

Figure 3-2: Flow of Information through the Supply Chain in VMI

In VMI, it is often assumed that the vendor can order from itself at a lower cost than the retailer can, however, the exact differences between these two costs are not standard. The cost savings will differ from vendor to vendor. Since an exact difference in order costs cannot be determined, the same order cost will be used, since at most, the vendor's order costs will be equal to the retailer's order cost. Since the vendor orders on the retailer's behalf, the vendor incurs this cost. However, it is assumed that any cost savings realized by the VMI system will be shared by vendor and the retailer. This sharing of benefits will offset the increase in costs that the vendor in VMI incurs. How the cost savings are shared is outside the scope of this research.

3.2.3. Definition of Collaborative Planning, Forecasting and Replenishment (CPFR)

In CPFR, the retailer and vendor together create one forecast and one replenishment plan. The retailer also shares with the vendor its promotional plan, POS data, and inventory data. Because the retailer shares its promotional calendar, the vendor is aware, well in advance, of the exact timing of each promotional period and thus, able to order additional stock in anticipation. In the case without promotions, the information shared with the vendor in CPFR does not differ from the information shared with the vendor in VMI. The advantage of CPFR over VMI comes when the retailer is able to share information that the vendor cannot deduce from either the inventory or POS data. In our case, this unique information comes from prior knowledge of the *promotions*. Another advantage comes from collaborating on the replenishment plan. The vendor is permitted to make small adjustments to the retailer's orders if it is unable to completely fill an order because it is short a small amount. This prevents backorders from occurring and the penalties that are associated with it.

Figure 3-3: Flow of Information through the Supply Chain in CPFR

As discussed in Section 2.3.2 and Section 2.5.2, in the work done by Sari (2008a; 2008b), it is assumed that the parameters of the end-customer demand can be predicted when the retailer and vendor collaborate on the forecast in CPFR. That is, through collaboration, the retailer and vendor produce a more accurate forecast than the ones created by each party independently. Our work differs from Sari in that we do not make this assumption. In the case with without promotions, the vendor in CPFR is given access to the same information that the vendor in VMI is given access to. In this case, CPFR does not produce a more accurate forecast than VMI.

3.3. Simulating with Hypothetical Data

Two sets of simulations were created, the first, using a hypothetical, but representative set of data and the second, using an empirical set of sales data from 3M. The hypothetical data was developed to reflect empirical data. Weekly demand was generated from a normal distribution with a mean of 250 which, after examining the empirical data from 3M, was deemed reasonable. Varying levels of demand variability were examined by using three different coefficients of variation (CV), a low, medium, and high, as outlined in Table 3-1. Any values generated less than zero were discarded.

Table 3-1: End-Customer Demand Standard Deviations for Simulations using Hypothetical Data

	Demand Variability Level		
	Low	Medium	High
Coefficient of Variation	0.2	0.4	0.6
Regular Demand Standard Deviation	50	100	150
Promotional Demand Standard Deviation	100	200	300

The lead time between the retailer and the vendor was one week and the lead time between the vendor and the manufacturer was two weeks. Order costs for the retailer and vendor were $100 and $75 respectively, and holding costs were $0.05 per item, per week for the retailer and $0.01 per item, per week for the vendor. These cost parameters were selected as they gave reasonable reordering times. The retailer would order approximately once every four weeks while the vendor would place an order on average, every seven weeks.

Two scenarios, one with retailer-determined promotions and one without any promotions, were created to compare the performances of VMI and CPFR. In the scenario without promotions, weekly demand was generated from the same distribution throughout the year. End-customer demand was forecasted by the retailer using exponential smoothing since it is an effective method to use when demand is steady (Silver et al., 1998). Silver et al. suggested that a smoothing constant between 0.01 and 0.30 should be utilized. Several α values less than 0.3 were tested, and α=0.1 was found to produce the lowest cost for the retailer for all three different levels of demand variability. A new forecast was created at the beginning of each quarter using the sales data from the previous quarters. The forecast was then used to calculate the retailer's economic order quantity during that quarter. In the case of CPFR, the forecast was created several weeks before the beginning of each quarter to accommodate the vendor's time constraints. The vendor ordered in advance of the retailer and therefore required the retailer's demand forecast a few weeks before the start of each quarter.

The retailer's reorder point was determined using the method devised by Silver et al. (2009). Since consumer demand was modeled on a weekly basis, a true continuous-review method could not be implemented. Silver et al. recognized that the inventory would reach a continuous-review, reorder point at a random time between review periods and denoted the time between that reorder point and the next review period as τ. Silver et al. could then liken the problem to a continuous-review model with a lead time of $L + \tau$, where L was the original lead time. We selected two different customer fill rates to be tested, 95% and 99%. The fill rate, or P_2 service level, refers to the percent of end customer demand that is fulfilled from inventory on hand. Using the method devised by Silver et al. (2009), we easily calculated reorder points that would attain both fill rates. The reorder points for both 95% and 99% fill rates are listed in Table 3-2.

Table 3-2: Reorder Points for the Retailer in Simulations using Hypothetical Data

Target Fill Rate (%)	Demand Variability Level		
	Low	Medium	High
95	355	389	442
99	458	529	630

Since there were no costs associated with a stockout at the retailer, it was necessary to specify target fill rates. In doing so, similarly performing systems would be compared. IS will be compared to VMI where both retailers have achieved a target fill rate of 95%, rather than to a VMI system that has achieved a 98% fill rate at the retailer.

The vendor's replenishment strategy varied based on the supply chain strategy used. In IS, periodic order quantity (POQ) was used to determine when orders should be placed since the demand the vendor faced was intermittent and POQ could provide better cost performance than other lot sizing methods (Silver et al., 1998). The calculated POQ values can be found in Table 3-3.

Table 3-3: Periodic Order Quantities for the Vendor in Simulations using Hypothetical Data

POQ in Weeks	Demand Variability Level		
	Low	Med	High
Without Promotions	7	7	7
With Promotions	7	7	6

Orders were placed with the manufacturer by the vendor in IS once every seven days except in the case with promotions and high demand variability, where an order was placed every six days. However, if the vendor already had sufficient inventory, the order was delayed until stock levels dropped below a reasonable level. When calculating the number of weeks each replenishment would cover, the results were rounded down to determine how frequently orders should be placed. For example, in the case with low demand variability and no promotions, the POQ was 7.748 weeks, which was rounded down to 7 weeks. This resulted in orders being placed slightly more often than necessary. Delaying orders placed to the manufacturer until inventory levels dropped below a reasonable amount prevented inventory levels from growing. The level of inventory that would delay an order being placed, which we will refer to as the postponement amount, depended on the targeted fill rate and demand variability. The postponement amounts were selected to ensure no stockouts at the vendor level would occur. By eliminating stockouts at the vendor, stockouts at the retailer would only occur due to errors in the retailer's forecast. This also eliminated the need for a backorder cost. Since these costs can vary from company to company, it is difficult to determine a representative backorder cost. Instead, as was the case for retailer stockouts, the vendors in the different systems achieved the same rate of stockouts.

The vendor also forecasted using exponential smoothing and updated its forecast after an order was received from the retailer. The vendor used $\alpha=0.1$ for all three CVs modeled, as it produced the lowest cost of the α values tested.

In VMI and CPFR, the vendor was able to make its own replenishment decisions based on the levels of inventory at the retailer. The vendor could anticipate when orders would be placed and if inventory levels were insufficient, place an order with the manufacturer. Since demand was steady and orders were controlled by the vendor in VMI, the vendor could determine a constant order amount for the retailer and order multiples of this amount from the manufacturer. In CPFR, the vendor could determine how much the retailer would order by consulting the agreed upon demand forecast and replenishment plan. The vendor could also make small adjustments to incoming orders from the retailer if the original orders would result in stockouts. If the vendor was short less than 10% of the retailer's order, the vendor could adjust the order to the amount of inventory on hand and avoid a stockout from occurring.

In the scenario with promotions, a two-week promotional period would take place every quarter with demand remaining steady during non-promotion weeks. Weekly demand for promotional periods was generated from a normal distribution with a mean of 500. The coefficient of variation remained constant, which resulted in standard deviations of 100, 200, and 300 for low, medium, and high levels of demand variability respectively. The timing of the promotional periods was randomly varied from one quarter to the next, by generating a number from a discrete uniform U(1,13) distribution and in IS, was only known to the retailer. The timing of each promotion was determined using a random number generator to ensure that the vendors in IS and VMI could not predict their occurrence. Both the retailer and vendor knew the exact timing of each promotion in CPFR. In VMI, the vendor was only informed that a two-week promotion period would take place each quarter, but when exactly during the quarter the promotion would take place was not shared. To ensure there would be sufficient inventory during the promotional period, the retailer would inform the vendor of the upcoming promotion a week prior to its start. This would give the vendor enough time to ship additional inventory to the retailer provided the vendor had it available. Therefore, it was a part of the vendor's policy to order additional inventory for promotions at the beginning of every quarter.

The hypothetical data simulation was run as a steady state simulation over 676 weeks or 13 years. The first 156 weeks or 3 years were deleted to remove the effects of the transient stage. To determine the transient stage, the vendor's inventory levels were graphed over the first 400 weeks of a simulation run as shown in Figure 3-4. The warm-up period was selected conservatively to ensure that the effects of initial transient in all the different simulations were not included in the results. Vendor inventory levels were plotted for all simulations to ensure that the warm-up period selected would work in all situations. Simulations with promotions were found to take longer to warm-up than simulations without promotions.

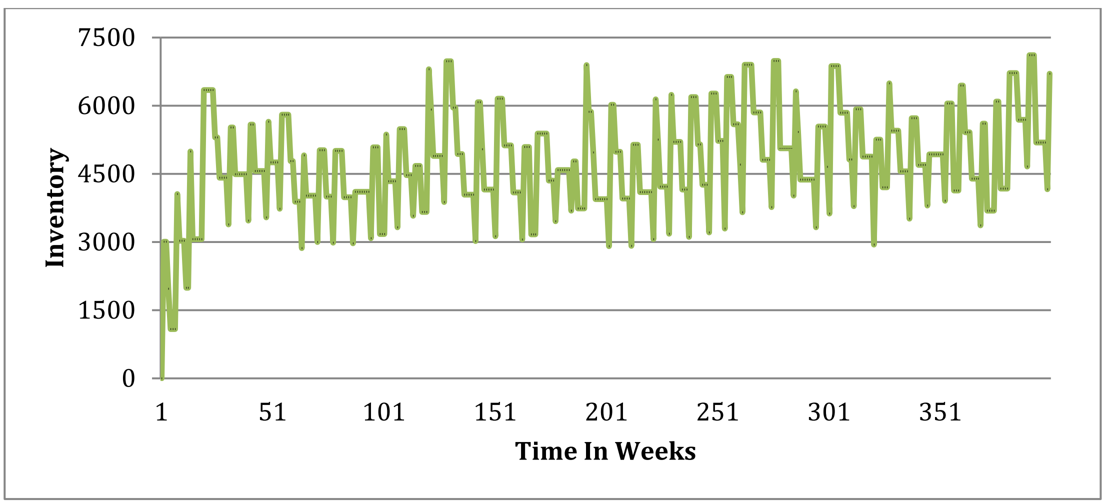

Figure 3-4: Inventory Levels at the Vendor in Simulations using Hypothetical Data

The transient stage appeared to end around the 110th week so observations up to the 156th week were discarded. In the steady state, inventory levels at the vendor appear to fluctuate approximately between 3000 and 6750 units.

While the method of batch means could have been used in the simulations using hypothetical data, it could not have been used in the case when there was a single, yearly promotion, a situation that arose in the simulations using empirical data. For this reason, the replication deletion approach was used. Law (2007) also states that using this approach gives good statistical performance and is easy to understand. 2000 replications were run for each of the 12 cases. Running the simulation with 3000 replications resulted in an insignificant change to the results.

To reduce variance, common random numbers were used for the end-customer demand and promotional weeks. According to Law (2007), when common random numbers are utilized, the differences observed in the simulation results are due to the differences in the systems, and not due to differences in the random numbers generated. Therefore, when assessing two or more systems, they are evaluated under comparable conditions.

3.3.1. Model Verification and Validation

As suggested by Law (2007) to assist in debugging, the simulation model was written in subprograms, starting with programming the retailer in IS with no promotions. IS with steady demand was modeled first, as it was the simplest of the supply chain methodologies and the simplest of the demand patterns to program. A dummy vendor with infinite inventory was created to confirm that the retailer was updating its forecast, reordering when inventory levels fell below the reorder point, and achieving the target fill rates. Once the retailer was deemed to be operating as intended, the vendor in IS was programmed into the simulation model. The VMI and CPFR models were developed from this initial IS model with each feature of VMI and CPFR programmed one at a time, rather than all at once. With the IS, VMI, and CPFR models without promotions developed, the simulations with promotions could be created.

To aid in debugging the simulation model, a variety of input parameters were utilized to ensure that the model would behave as expected (Banks et al., 2005). For example, lower reorder points at the vendor would increase the number of lost sales at the retailer, increased POQ figures would increase the number of days between orders at the vendor, and an increased postponement amount would result in fewer orders being delayed at the vendor. As recommended by both Law and Banks et al., a trace was performed on each simulation at both the retailer level and the vendor level and compared to hand calculations to ensure that the simulation was operating as intended week after week. This made certain that inventories were being updated, replenishments were arriving on schedule, orders were placed when required, forecasts were being updated, and replenishment plans changed accordingly.

The simulation models were validated with the help of subject-matter experts, or SMEs (Law, 2007). SMEs in VMI and CPFR were consulted throughout the process of building the simulation model. Their insights into the behaviours of the retailer and vendor in VMI and CPFR were used to accurately model each respective supply chain system. An assumptions document was used to record the concepts, assumptions, and parameters used in the model and was reviewed with the SMEs. Case studies on VMI and CPFR were also consulted to determine how the supply chain methodologies operated in practice. Since output data from an existing system was not made available to us, we could not validate the results from the simulation model. As noted in the recommendations section of this thesis (Section 5.2.4), accomplishing this will be our next step.

3.4. Simulating with Empirical Data

The empirical data, provided by 3M, was weekly sales data for several consumer products from four different product categories. The categories were stationary, home care, first aid, and hardware products. The different products exhibited different sales patterns; some had steady demand throughout the year, while others had spikes in sales occurring one or more times during the year. Data exhibiting similar patterns to the demand in simulations using hypothetical data were selected to be simulated, data that showed no spikes in sales and data that showed several spikes throughout the year that could be a result of promotional activities by the retailer. Graphs of the sales data showing these patterns can be seen in Figure 3-5 and Figure 3-6.

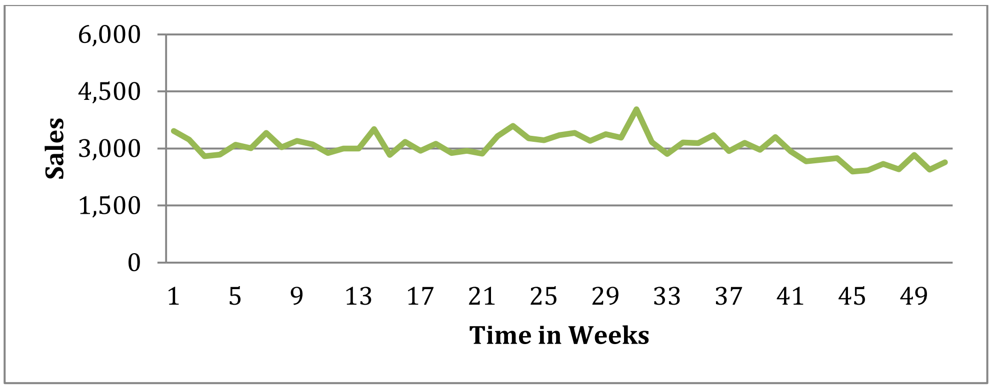

Figure 3-5: 3M Sales Data for a Product with Steady Demand

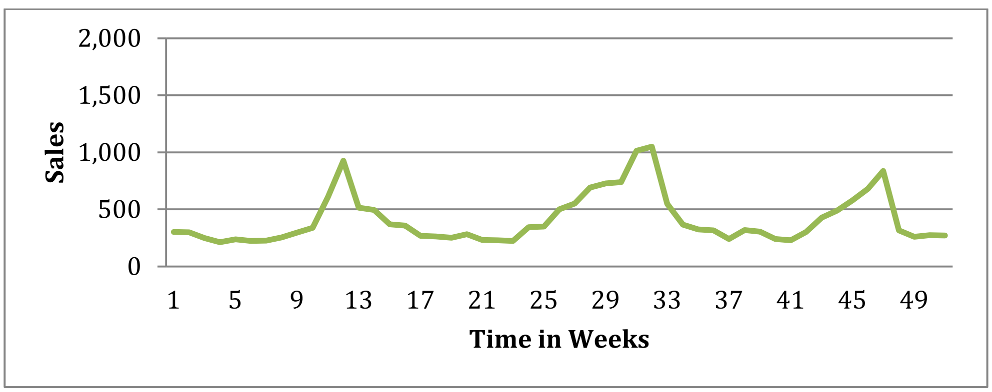

Figure 3-6: 3M Sales Data for a Product with Multiple Spikes in Demand

A third, interesting demand pattern was discovered among the sales data provided where a steep spike in demand occurred close to the end of the year. An example of this demand pattern can be seen in Figure 3-7. Products in this category experienced the majority of their demand during a few weeks of the year. Two products displaying this characteristic were also selected to be simulated.

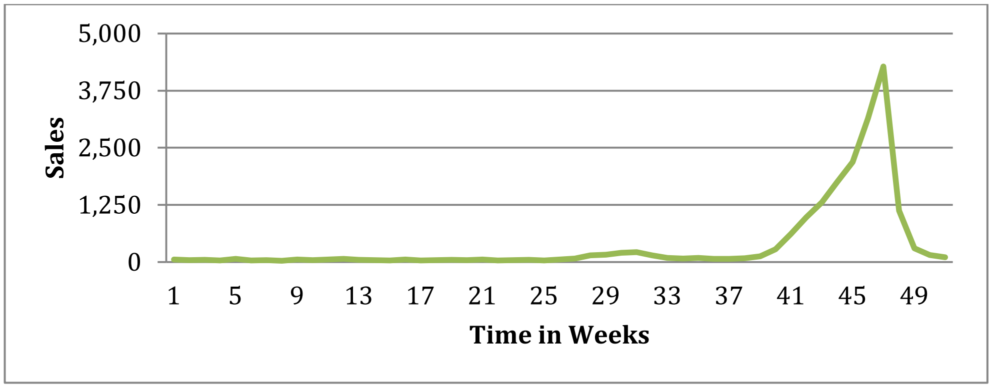

Figure 3-7: 3M Sales Data for a Product with One Large Spike in Demand

From the data provided, the mean and standard deviation of demand could be calculated. Demand was generated from normal distributions using these parameters since attempts to fit a distribution to the data were unsuccessful. Crystal Ball's distribution fitting feature was unable to find a distribution with a P-Value greater than 0.5 using the Kolmogorov-Smirnov test.

Table 3-4: Cost Parameters for Simulations using Empirical Data

Scenario	Level	A/r Ratio	r	Holding Cost ($/year)	Ordering Cost ($)
1	Retailer	100	0.25	1.25	25
	Vendor	200	0.1	0.3	20
2	Retailer	500	0.25	1.25	125
	Vendor	1000	0.1	0.3	100

Since cost parameters were not provided from 3M, exchange curves were used to find a ratio of the ordering cost and carrying cost (A/r) that would produce reasonable order times. Different A/r ratios and lead times were tested to examine how IS, VMI, and CPFR would perform under a variety of conditions. Two sets of A/r ratios were used in the simulations, A/r=100 for the retailer and A/r=200 for the vendor in Scenario 1 and A/r=500 for the retailer and A/r=1000 for the vendor in Scenario 2. The term 'Scenario' refers to a specific set of A/r ratios and lead times that were used in the simulations. The cost and lead time parameters used in Scenario 1 and Scenario 2 were consistent for each demand pattern. These cost parameters can be found in Table 3-4. A unit value of v=$5.00 for the retailer and v=$3.00 for the vendor was hypothesized based on the product categories. These parameters were different than those used in the simulations with hypothetical data.

Along with having a lower A/r ratio, products in Scenario 1 had shorter lead times, one week at the vendor and two weeks at the manufacturer. Products in Scenario 2 had a lead time of two weeks at the vendor and a lead time of four weeks at the manufacturer.

For each A/r scenario, a product with steady demand, a product with several spikes in its demand, and a product with one large spike in demand was selected. Products with steady demand were designated as products without promotions, while products with spikes in its demand were defined as having promotions.

Table 3-5: Demand Parameters for Simulations without Promotions - Empirical Data

Scenario	Product Category	Weekly Demand Mean	Weekly Demand Standard Deviation	Demand Coefficient of Variation
1	First Aid	72	24	0.333
2	Home Care	2518	278	0.110

In the case without promotions, a first aid product was chosen for Scenario 1 and a home care product for Scenario 2. The demand parameters determined from the 3M data can be found in Table 3-5. The weekly sales of the first aid product and the home care products can be found in Figure 3-8.

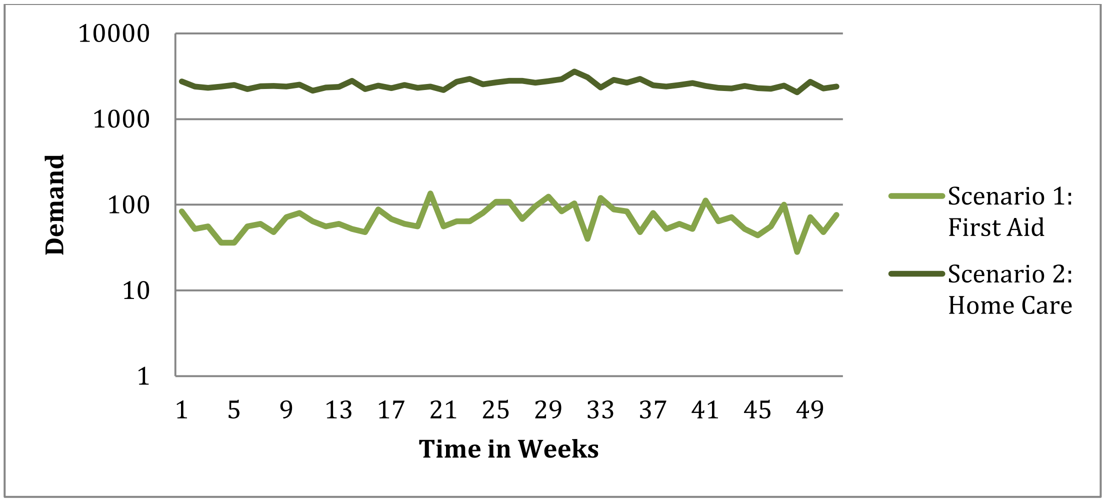

Figure 3-8: Weekly Demand for Products Selected for Simulations without Promotions

In simulations with multiple promotions, a stationary product was chosen for simulations using Scenario 1's cost and lead time parameters, while a hardware product was selected for simulations using Scenario 2's cost and lead time parameters. The demand parameters used for simulations with multiple promotions can be found in Table 3-6. Figure 3-9 shows the weekly sales data for the stationary and hardware product.

Table 3-6: Demand Parameters for Simulations with Multiple Promotions - Empirical Data

Scenario	Product Category	Demand Type	Weekly Demand Mean	Weekly Demand Standard Deviation	Demand CV
1	Stationary	Regular	213	54	0.254
		Promotional	590	213	0.255
2	Hardware	Regular	306	89	0.291
		Promotional	949	275	0.290

Similar to how the simulations using hypothetical data were setup, four promotional periods were planned throughout the year, one in each quarter. The timing of the promotion in each quarter was determined using a uniform distribution. Each promotional period would last for two weeks.

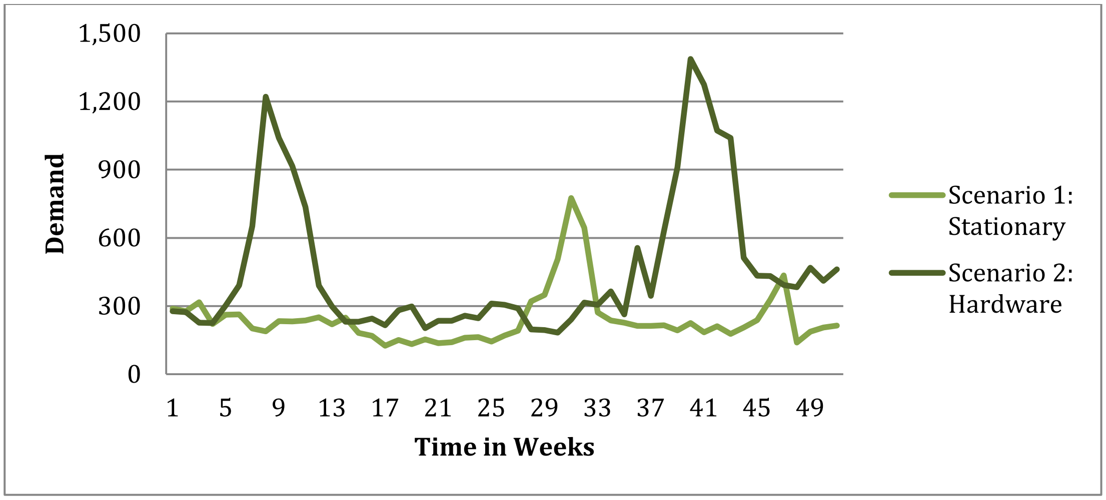

Figure 3-9: Weekly Demand for Products Selected for Simulations with Multiple Promotions

For simulations without promotions and simulations with multiple promotions, P_2 targets of 95% and 99% were modeled. As we did when using hypothetical data, the simulations were run for a total of 676 weeks or 13 years with the first 156 weeks or 3 years deleted to remove the effects of the transient stage. Both types of simulations were run for 2000 replications. Common random numbers were used for each supply chain strategy to reduce variance in the model.

A stationary product was selected to model Scenario 1 in simulations with one large, seasonal promotion. During the majority of the year, weekly demand had an mean of 88 and a standard deviation of 33, giving it a coefficient of variation of 0.398. In the sales data provided, demand peaked during the holiday season at 85,852 units/week. With a longer holiday season, and short lead times in Scenario 1, orders were placed weekly by the retailer during the holiday season. This allowed the retailer to update its forecast as the holidays progressed. It was assumed that as the holidays progressed, the retailer would gain more information on how the customers were reacting to the year's promotions and thus, better forecast demand.

For the second scenario, a hardware product was selected. Demand for most weeks of the year had a mean of 972, a standard deviation of 261, and a coefficient of variation of 0.269. In Scenario 2, lead times are longer and with a shorter holiday season, weekly replenishments would be fruitless. With longer lead times, it becomes more difficult to incorporate information learnt from recent demand. In the case of VMI and CPFR, the vendors would have to place their orders before the holiday season starts. Therefore, rather than weekly replenishments, the retailer in Scenario 2 places one order for the holiday season just prior to its start. A graph showing the weekly sales data for these two products can be found in Figure 3-10.

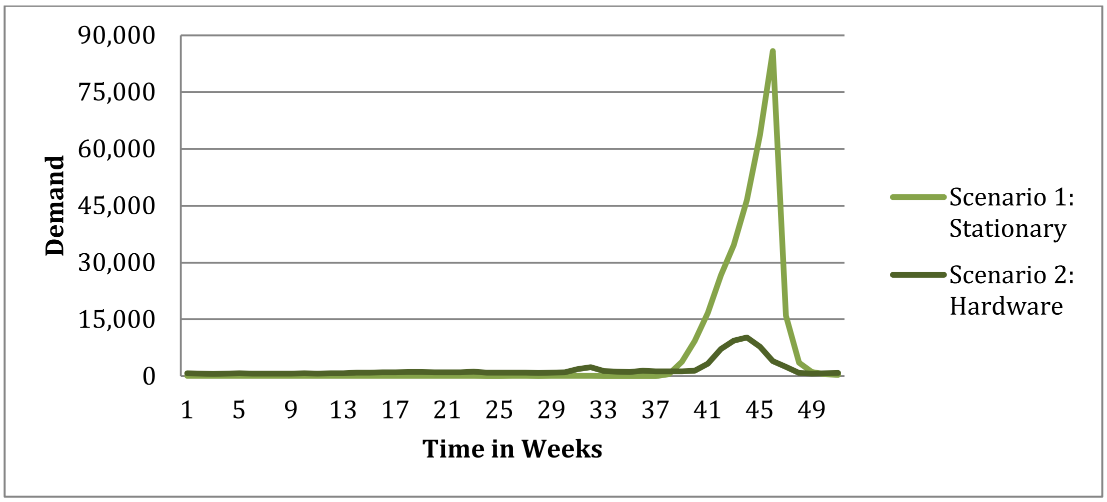

Figure 3-10: Weekly Demand for Products Selected for Simulations with One Large Seasonal Promotion

To model different intensities of retailer promotions from year-to-year, three demand peaks were created. If the retailer aggressively promoted the product during the holiday season, the demand peak would increase. If the retailer did not promote the product during the holiday season, the demand peak would decrease. The demand during the weeks leading up to and after the peak of demand were scaled accordingly. In Scenario 1, the demand peak increased or decreased by 50% while in Scenario 2, the demand peak increased or decreased by 20%. In each year modeled in the simulation, how the retailer promoted the product was determine by generating a number from the discrete uniform distribution, U(1,3). If a 1 was generated for a particular year, the retailer would not promote the product during the holiday season that year. If a 2 was generated, the retailer would promote the product and if a 3 was generated, the retailer would aggressively promote the product. The realized peak of demand was then generated from a normal distribution. The parameters used to generate the peak of demand for each level of promotion intensity are found in Table 3-7. The standard deviations were calculated using the regular demand's coefficients of variation; the coefficient of variation equaled 0.398 in Scenario 1 and 0.269 in Scenario 2.

Table 3-7: Demand Parameters used in Simulations with One Large, Seasonal Promotion

Scenario	Demand Peak Parameters	Promotional Plan		
		No Promotions	Regular Promotions	Aggressive Promotions
1: Stationary Product	Mean	42926	85852	128778
	Standard Deviation	17067	34135	51202
2: Hardware Product	Mean	10222	12778	15333
	Standard Deviation	2745	3431	4117

For simulations with one large, seasonal promotion, target fill rates of 95% and 99% were tested. These simulations were run longer than the previous simulations since the promotional season only occurred once a year. Scenario 1 was run for 884 weeks or 17 years, and Scenario 2 was run for 936 weeks or 18 years. In both cases, the first 156 weeks or 3 years were deleted to remove the effects of the transient stage. The lengths of each replication was limited by the computing resources available. 2000 replications were run for each scenario and each target fill rate. Once again, common random numbers were used for each supply chain methodology to reduce variance in the results.

The results from this and the other simulations detailed in this chapter can be found in the following chapter. The simulations using hypothetical data will be presented first in Section 4.1 and the results from the simulations using empirical data will be in Section 4.2.

Chapter 4: Results

4.1. Results for Simulations using Hypothetical Data

The subsequent types information was recorded following the sets of simulation runs:

- Averages
  - Costs at the retailer, vendor and system-wide (retailer + vendor)
  - Inventory levels at the retailer, vendor and system-wide
  - Stockout rates at the retailer and the vendor
- 90% Confidence Intervals
  - Cost differences between IS, VMI, and CPFR at the retailer, vendor and system-wide
  - Inventory differences between IS, VMI, and CPFR at the retailer, vendor and system-wide
  - Stockout rate differences between IS, VMI, and CPFR at the retailer and the vendor

4.1.1. Stockout Results for Simulations using Hypothetical Data

The achieved P_2 service levels, or fill rate, at the retailer were recorded to ensure that the targets of 95% and 99% were being achieved. This would allow for fair comparisons between IS, VMI, and CPFR to be made. 90% confidence intervals of the differences in stockout rates were also collected to make certain that there were no statistically significant differences between the stockout rates of the different supply chain strategies.

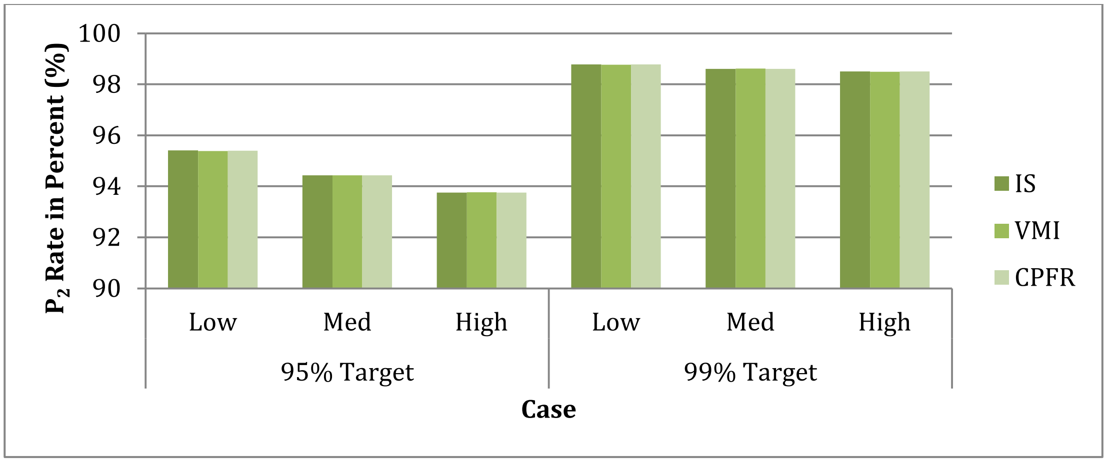

Figure 4-1: P_2 Service Rates Achieved by the Retailer in the Simulations without Promotions - Hypothetical Data

The stockout rates that were achieved in simulations without promotions can be found in Figure 4-1 and in Figure 4-2 for simulations with promotions. The targets of 95% and 99% were met and the confidence intervals showed that there was no statistically significant difference between the stockout rates of the retailers and between the stockout rates of the vendors in IS, VMI, and CPFR. Tables containing the confidence intervals can be found in Appendix A.

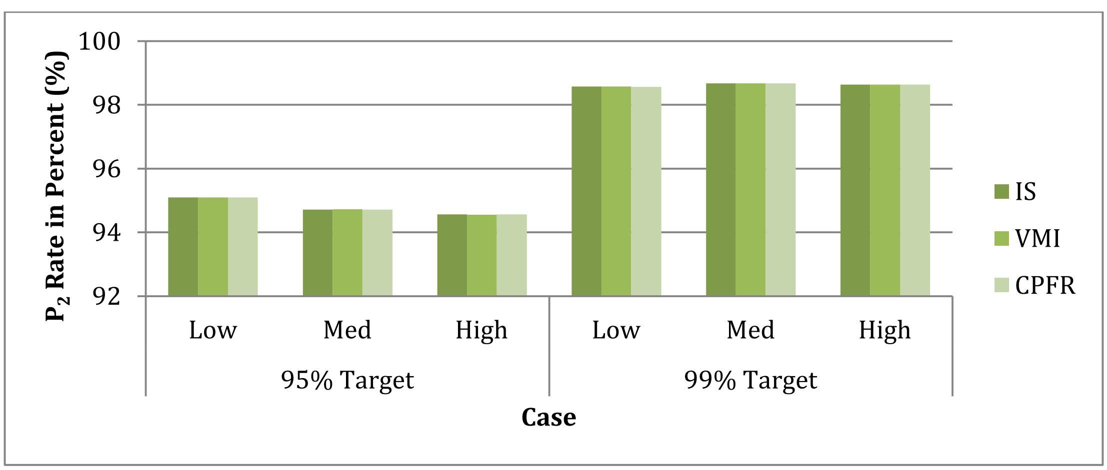

Figure 4-2: P_2 Service Rates Achieved by the Retailer in the Simulations with Promotions - Hypothetical Data

4.1.2. Cost Results for Simulations using Hypothetical Data

Costs were calculated by adding the holding costs and the ordering costs. Since stockouts at the vendor were eliminated, penalty costs for backorders were not required. There was no penalty cost for sales lost at the retailer. Costs were calculated over the length of the simulation which was 520 weeks. A graph with the average system-wide costs over the 2000 replications can be found in Figure 4-3 for simulations without promotions and in Figure 4-4 for simulations with promotions.

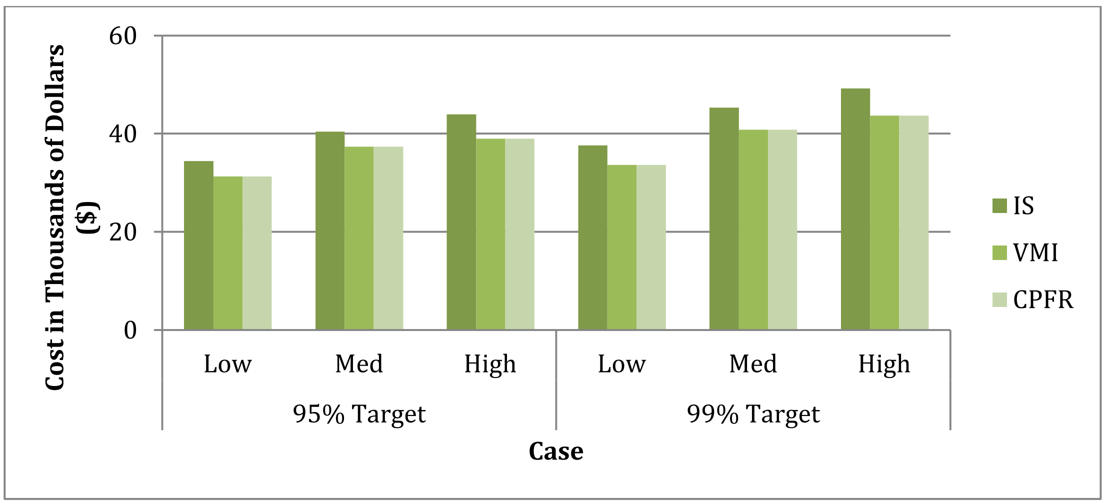

Figure 4-3: Average System-Wide Costs in Simulations without Promotions - Hypothetical Data

In simulations without promotions, both VMI and CPFR had lower system-wide average costs than IS. This was true for both P_2 targets and for all three levels of demand variability. The average system-wide costs were similar for VMI and CPFR. In the case with promotions, IS's system-wide average cost was the greatest in all cases. VMI had the second largest average cost and CPFR had the lowest in all cases.

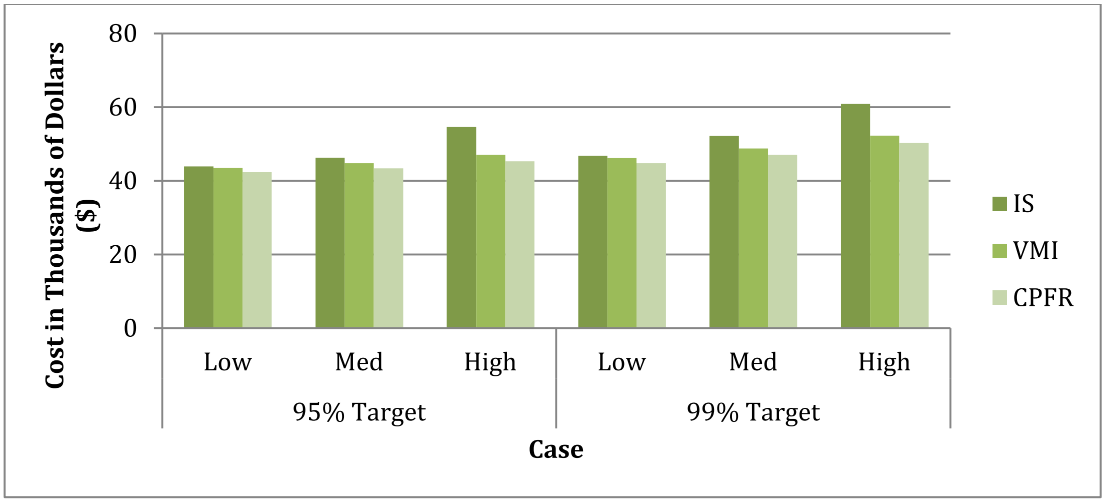

Figure 4-4: Average System-Wide Costs in Simulations with Promotions - Hypothetical Data

Table 4-1: Confidence Intervals for the Difference in System-Wide Cost ($ Thousands) for Simulations without Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Cost Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(0.79, 5.03)	(0.80, 4.96)	(-0.26, 0.25)
	99	(1.63, 5.39)	(1.64, 5.30)	(-0.40, 0.41)
Medium	95	(1.41, 5.12)	(1.48, 4.99)	(-0.41, 0.35)
	99	(2.86, 5.62)	(2.97, 5.53)	(-0.50, 0.48)
High	95	(3.68, 5.83)	(3.81, 5.84)	(-0.47, 0.46)
	99	(4.61, 6.65)	(4.63, 6.63)	(-0.65, 0.54)

From the confidence intervals in Table 4-1 for simulations without promotions, and in Table 4-2 for simulations containing promotions, we can see that these differences were often statistically significant. When promotions were not present, the confidence interval showed cost savings when utilizing VMI or CPFR over IS, since both the lower and upper limits of the confidence interval were positive. When this occurs, we will refer to the confidence interval as being "positive". When both the lower and upper limits of the confidence interval are negative, we will refer to the confidence interval as being "negative". As P_2 targets and demand variability increased, the savings in costs also increased. There was no statistically significant cost difference between VMI and CPFR. This observation held for the three different levels of demand variability and the two targeted fill rates.

Table 4-2: Confidence Intervals for the Difference in System-Wide Cost ($ Thousands) for Simulations with Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Cost Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(-0.22, 1.10)	(1.07, 2.21)	(0.67, 1.72)
	99	(-0.13, 1.34)	(1.38, 2.76)	(0.74, 2.16)
Medium	95	(0.77, 2.28)	(2.22, 3.60)	(0.76, 2.03)
	99	(2.51, 4.40)	(4.34, 6.03)	(0.90, 2.60)
High	95	(6.58, 8.49)	(8.50, 10.10)	(1.03, 2.47)
	99	(7.51, 9.71)	(9.74, 11.63)	(1.10, 2.97)

When promotions were present, (Table 4-2), there was no statistically significant difference between IS and VMI at the lowest level of demand variability. However, at higher levels of demand variability, there was a statistically significant cost improvement when VMI was implemented rather than IS. In all demand variability and target fill rate cases, CPFR was less costly than both IS and VMI. Again, the cost savings improved as the CV of demand increased and as the target fill rate increased.

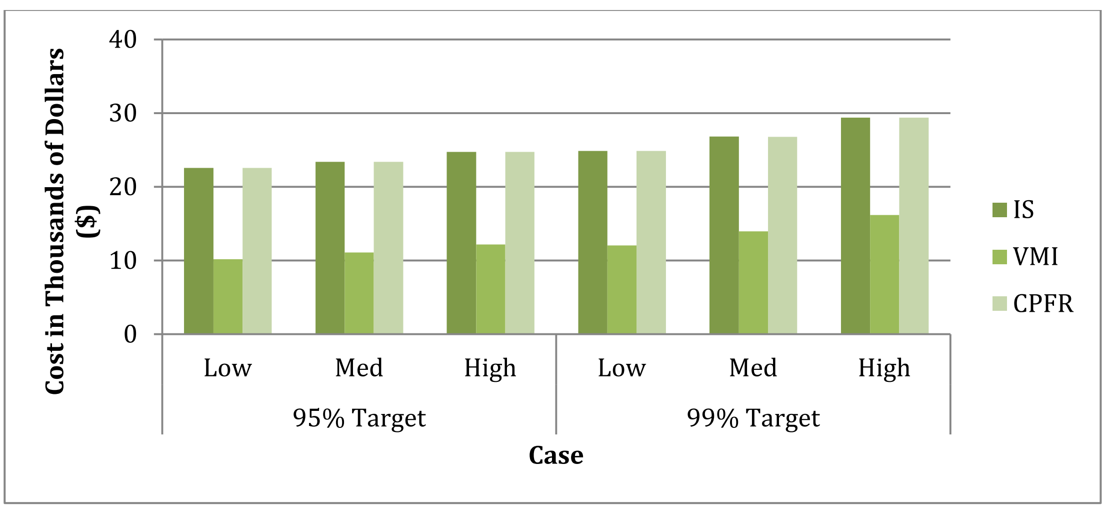

Figure 4-5: Average Retailer Costs in Simulations without Promotions - Hypothetical Data

Costs were recorded at the retailer and vendor levels, each comprised of the holding and ordering costs incurred by each party. As shown in Figure 4-5, very little costs are attributed to the retailer in VMI compared to the retailers in IS and CPFR since this only reflects the retailer's holding costs. The vendor in VMI seems to have much greater costs than the vendors of IS and CPFR (Figure 4-6) since this was comprised of the vendor's holding and ordering costs and the costs for the retailer's replenishments. The VMI retailer and vendor costs in simulations with promotions exhibited similar patterns. However, these costs did not take into account how the retailer and vendor shared any potential cost savings that may have resulted due to implementing VMI or CPFR. We assumed that the retailer and vendor had a good working relationship and that any cost savings realized through VMI or CPFR were shared to the satisfaction of both organizations. Rather than examining the retailer and vendor costs, we will focus our analysis on the system-wide costs and on inventory levels in the supply chain.

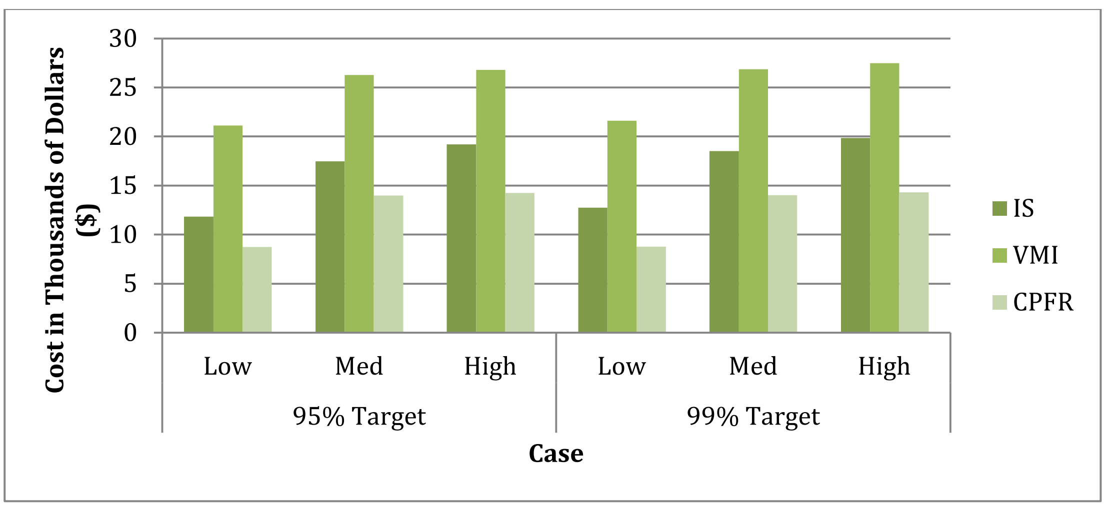

Figure 4-6: Average Vendor Costs in Simulations without Promotions - Hypothetical Data

4.1.3. Inventory Results for Simulations using Hypothetical Data

Inventory levels at the retailer and vendor level were also calculated over the 520 week period. The system-wide averages calculated over the 2000 replications are found in Figure 4-7 for simulations without promotional periods and in Figure 4-8 for simulations with promotional periods. The inventory results had patterns similar to the system-wide cost results.

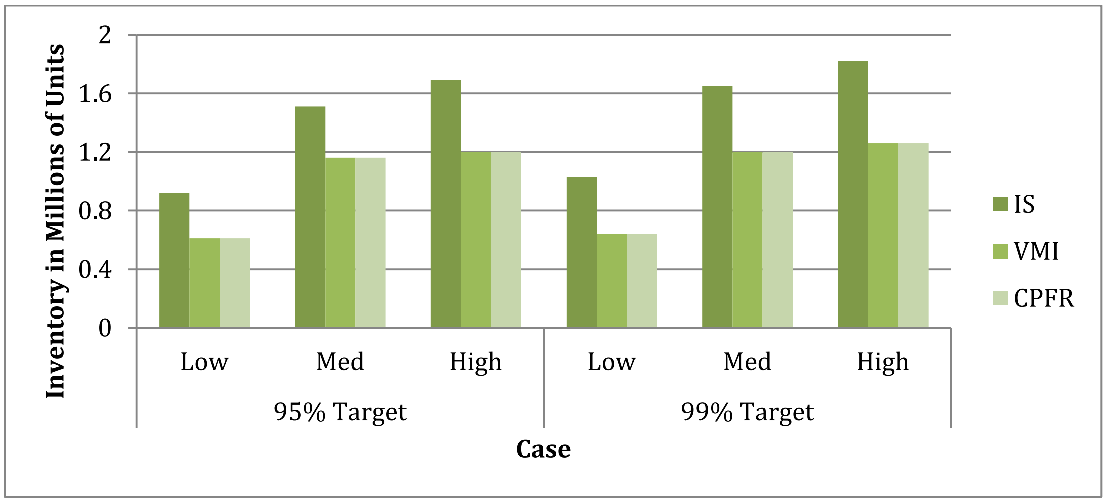

Figure 4-7: Average System-Wide Inventory in Simulations without Promotions - Hypothetical Data

On average, in simulations that did not include promotions, VMI and CPFR did not require as much inventory in the supply chain as IS did to achieve the same fill rate at the retailer. In simulations that included promotions, just as IS had the greatest system-wide costs, IS had the largest amount of inventory in the supply chain. The results for VMI and CPFR also mirrored the system-wide cost results; VMI held the second largest average amount of inventory and CPFR, the least amount of inventory.

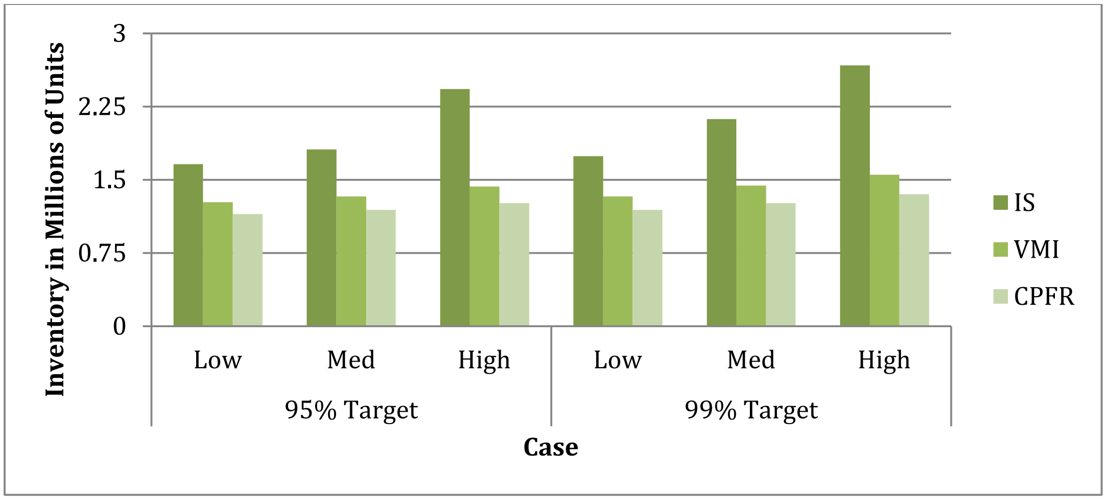

Figure 4-8: Average System-Wide Inventory in Simulations with Promotions - Hypothetical Data

These results were confirmed by the confidence intervals in Table 4-3 and Table 4-4. When promotions were not present, there was a statistically significant reduction in inventory levels when implementing IS rather than VMI or CPFR. These reductions improved as the demand CV increased and as the P_2 rate increased. There was no such difference however between VMI and CPFR. Their confidence intervals contained zero for all three levels of demand variability and both target fill rates.

Table 4-3: Confidence Intervals for the Difference in Inventory System-Wide (Millions) in Simulations without Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(0.08, 0.50)	(0.08, 0.50)	(-0.01, 0.01)
	99	(0.17, 0.53)	(0.17, 0.52)	(-0.01, 0.01)
Medium	95	(0.14, 0.51)	(0.15, 0.50)	(-0.02, 0.02)
	99	(0.28, 0.55)	(0.30, 0.54)	(-0.03, 0.02)
High	95	(0.37, 0.57)	(0.38, 0.58)	(-0.03, 0.03)
	99	(0.48, 0.65)	(0.47, 0.66)	(-0.03, 0.03)

When promotions were present in the simulation, both VMI and CPFR offered statistically significant inventory reductions over IS in all six cases. The reduction in inventory improved as demand variability and the target retailer fill rate increased. CPFR offered further inventory reductions over VMI and the drop in inventory grew as the demand variability and the target fill rate increased. These results were slightly different from the cost results; at the lowest level of demand variability, VMI offered statistically significant reductions in cost when chosen over IS. Whereas in the system-wide cost results, there was no statistically significant difference between the system-wide costs of IS and VMI.

Table 4-4: Confidence Intervals for the Difference in Inventory System-Wide (Millions) in Simulations with Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(0.33, 0.45)	(0.45, 0.57)	(0.11, 0.13)
	99	(0.34, 0.48)	(0.48, 0.62)	(0.13, 0.16)
Medium	95	(0.42, 0.56)	(0.56, 0.70)	(0.11, 0.16)
	99	(0.59, 0.76)	(0.77, 0.94)	(0.15, 0.20)
High	95	(0.91, 1.09)	(1.10, 1.26)	(0.14, 0.21)
	99	(1.02, 1.22)	(1.24, 1.42)	(0.17, 0.24)

Inventory levels were also recorded at the retailer and vendor level for IS, VMI, and CPFR. Figure 4-9 shows the levels of inventory held at the retailer for simulations without promotions. In all six test cases, the retailers of IS, VMI, and CPFR held similar amounts of inventory. This pattern was found in both simulations with and without promotions. Confidence intervals verified this observation; there was no statistically significant difference between the amounts of inventory held at the retailers of IS, VMI, and CPFR. Graphs and charts showing the average levels of inventory at the retailer in simulations with promotions and the differences in inventory held by the retailers in IS, VMI, and CPFR can be found in Appendix A.

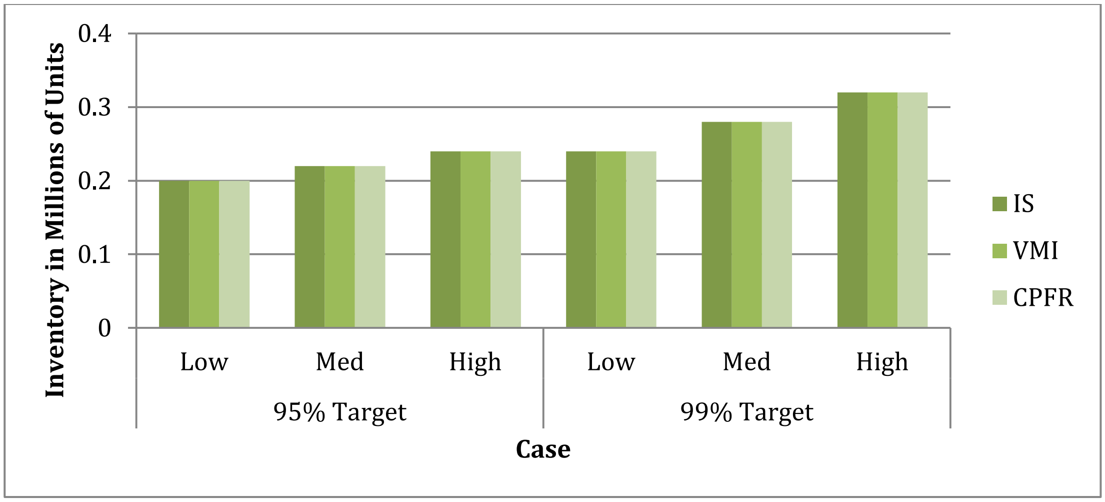

Figure 4-9: Average Inventory held by the Retailer in Simulations without Promotions - Hypothetical Data

Inventory at the vendor followed the same patterns found in the system-wide inventory results. Graphs of inventory at the vendor in both simulation with and without promotions had the same patterns as their respective system-wide inventory graphs (Figure 4-7 and Figure 4-8). When promotions were not present, the vendor in IS carried the most amount of inventory. The vendors in VMI and CPFR required far less inventory than in IS, and their averages were similar. When promotions were present, the vendor in IS held the most inventory and vendor in CPFR held the least amount of inventory. Graphs depicting the average amounts of inventory held by the vendors in IS, VMI, and CPFR for both simulations with and without promotions can be found in Appendix A.

Table 4-5: Confidence Intervals for the Difference in Inventory held at the Vendor (Millions) in Simulations without Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(0.08, 0.50)	(0.08, 0.50)	(-0.01, 0.01)
	99	(0.17, 0.53)	(0.17, 0.52)	(-0.01, 0.01)
Medium	95	(0.14, 0.51)	(0.15, 0.50)	(-0.02, 0.02)
	99	(0.29, 0.55)	(0.30, 0.54)	(-0.02, 0.02)
High	95	(0.38, 0.57)	(0.38, 0.58)	(-0.03, 0.02)
	99	(0.48, 0.65)	(0.47, 0.66)	(-0.03, 0.03)

The differences in inventory held by the vendors in the three supply chain strategies were, for the most part, statistically significant. When promotions were not present, the confidence intervals indicated that the vendor in IS carried more inventory than the vendors in VMI and CPFR. This difference in inventory grew as demand variability increased and the target P_2 rate rose. When promotions were present, VMI and CPFR both outperformed IS, and CPFR offered even greater reductions in inventory over VMI. The reduction in inventory grew as demand variability increased and as the target fill rate increased. The confidence intervals for simulations without promotions are in Table 4-5 and in Table 4-6 for simulations with promotions.

Table 4-6: Confidence Intervals for the Difference in Inventory held at the Vendor (Millions) in Simulations with Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(0.33, 0.45)	(0.45, 0.57)	(0.11, 0.13)
	99	(0.34, 0.48)	(0.48, 0.62)	(0.13, 0.16)
Medium	95	(0.42, 0.56)	(0.56, 0.70)	(0.12, 0.16)
	99	(0.59, 0.76)	(0.78, 0.93)	(0.15, 0.20)
High	95	(0.91, 1.09)	(1.10, 1.26)	(0.14, 0.20)
	99	(1.02, 1.22)	(1.24, 1.41)	(0.17, 0.23)

4.1.4. Summary of Results for Simulations using Hypothetical Data

Table 4-7 contains a summary of the system-wide cost results from the simulations with and without promotions. The table outlines whether the confidence intervals were positive or negative, whether zero was present in the confidence interval and if these observations were true for all cases. Table 4-8 offers a similar summary of the system-wide inventory results.

Table 4-7: Summary of System-Wide Cost Results in Simulations using Hypothetical Data

Case	Result	Confidence Interval: Cost Differences		
		IS-VMI	IS-CPFR	VMI-CPFR
Simulations without Promotions	Positive or Negative	Positive	Positive	Neither
	Contains 0	No	No	Yes
	For All Cases	Yes	Yes	Yes
Simulations with Promotions	Positive or Negative	Mostly Positive	Positive	Positive
	Contains 0	Mostly No	No	No
	For All Cases	No, when demand variability is low, the CI contains 0	Yes	Yes

From the table above, in simulations without promotions, the difference between system-wide costs for IS and VMI was positive, indicating that VMI was less costly than IS. The confidence interval did not contain zero and therefore, the result was statistically significant. This result held for all three levels of demand variability and both target P_2 fill rates. From the table below, again, in simulations without promotions, the difference between IS and VMI in system-wide inventory was positive. That is, VMI carried less inventory in the supply chain than IS. The confidence intervals did not contain zero indicating that this result was statistically significant. This result also held for all cases tested.

Table 4-8: Summary of System-Wide Inventory Results in Simulations using Hypothetical Data

Case	Result	Confidence Interval: Inventory Differences		
		IS-VMI	IS-CPFR	VMI-CPFR
Simulations without Promotions	Positive or Negative	Positive	Positive	Neither
	Contains 0	No	No	Yes
	For All Cases	Yes	Yes	Yes
Simulations with Promotions	Positive or Negative	Positive	Positive	Positive
	Contains 0	No	No	No
	For All Cases	Yes	Yes	Yes

Table 4-9 outlines the patterns displayed by the confidence intervals in the simulations run with the hypothetical data. As the coefficient of variation was increased or the target P_2 service rate was increased, the confidence interval could shift in a positive direction, a negative direction or not shift at all. The confidence intervals could also increase or decrease in width.

Table 4-9: Patterns in Confidence Intervals for Simulations using Hypothetical Data: Summary of System-Wide Costs and Inventory

Case	Action	Confidence Interval		
		IS-VMI	IS-CPFR	VMI-CPFR
Simulations without Promotions	CV increased	• interval narrowed • shifted positively	• interval narrowed • shifted positively	• interval widened • no shift
	P_2 target increased	• interval narrowed • shifted positively	• interval narrowed • shifted positively	• interval widened • no shift
Simulations with Promotions	CV increased	• interval widened • shifted positively	• interval widened • shifted positively	• interval widened • shifted positively
	P_2 target increased	• interval widened • shifted positively	• interval widened • shifted positively	• interval widened • shifted positively

From the table above, in simulations with promotions, when the demand coefficient of variation was increased, the IS-VMI confidence interval for both system-wide costs and inventory shifted in a positive direction. That is, the upper and lower limits of the confidence intervals increased. The width of the confidence interval also grew. The same patterns occurred when the target P_2 fill rate was increased; the confidence interval moved in a positive direction and the width of the confidence interval increased.

4.2. Results for Simulations using Empirical Data

Using the data provided by 3M, simulations were run and the same types of results were collected: costs, inventory levels and stockout rates. Both averages and 90% confidence intervals were documented from the 2000 replications that were run for each scenario. The various cases examined are summarized in Table 4-10. For each data trend, two products were selected, one for each A/r ratio and lead time scenario. In Scenario 1, A/r ratios were low and lead times short. In Scenario 2, A/r ratios were higher and lead times were longer than in Scenario 1.

Table 4-10: Summary of Cases Examined in Simulations Using Empirical Data

Data Trend	Scenario	Product Category	A/r ratios	Lead Times
No Promotions	1	First Aid	Low	Short
	2	Home Care	High	Long
Multiple Promotions	1	Stationary	Low	Short
	2	Hardware	High	Long
One Large Promotion	1	Stationary	Low	Short
	2	Hardware	High	Long

The achieved stockout rates for all three data trends and for both scenarios can be found in Appendix B.

4.2.1. Simulations without Promotions - Empirical Data Scenario 1: First Aid Product

The first aid product selected for Scenario 1 had a mean demand of 72 and a standard deviation of demand of 24, giving a CV of 0.333. The simulation results mirrored the results from the simulations using hypothetical data seen in Figure 4-3; IS had the greatest average system-wide costs for both target fill rates of 95% and 99%, while VMI and CPFR had similar system-wide costs that were lower than IS. A graph with the average costs can be found in Figure 4-10. As we did in the simulations using hypothetical data, we will focus our attention only to system-wide costs. It is again assumed that the retailer and vendor had an agreement to share any costs savings provided by VMI or CPFR that was to the satisfaction of both parties.

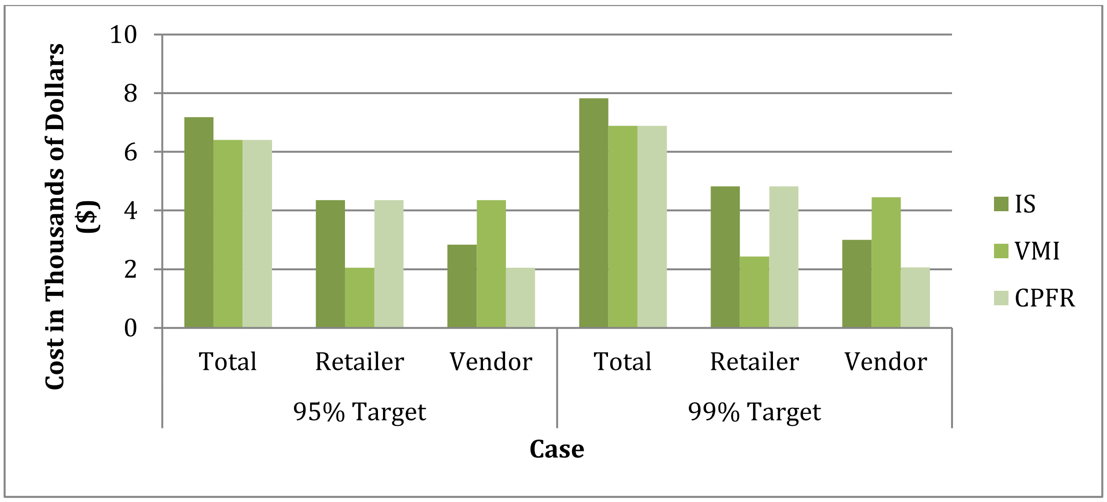

Figure 4-10: Average Costs for Simulations without Promotions - First Aid Product

Confidence intervals showed that the differences in system-wide costs between IS and VMI, and IS and CPFR were statistically significant. As expected, there was no statistically significant difference between the system-wide costs of VMI and CPFR. Table 4-11 contains the cost difference confidence intervals for the first aid product. As we saw in the confidence intervals from simulations using hypothetical data in Table 4-1, the differences in system-wide costs between IS and both VMI and CPFR increased as the target P_2 rate increased.

The levels of inventory held in the supply chain reflected the system-wide cost results. Graphing the levels of inventory resulted in a graph very similar to Figure 4-10. IS carried the greatest amount of inventory in the supply chain, and while both VMI and CPFR held less, there was no statistical significant difference in the amount of inventory held by either supply chain strategy. When the target fill rate was increased from 95% to 99%, the IS-VMI and IS-CPFR confidence intervals shifted in a positive direction and narrowed in width, just as it had in the simulations using hypothetical data (Table 4-9). Average inventory levels held by the retailer, vendor, and system-wide and confidence intervals for the inventory differences between IS, VMI, and CPFR can be found in Appendix B.

Table 4-11: Confidence Intervals for the Difference in System-Wide Costs ($ Thousands) in Simulations without Promotions - First Aid Product

Target Fill Rate (%)	Confidence Interval: Cost Differences		
	IS - VMI	IS - CPFR	VMI - CPFR
95	(0.33, 1.11)	(0.34, 1.10)	(-0.07, 0.06)
99	(0.57, 1.21)	(0.57, 1.21)	(-0.09, 0.07)

4.2.2. Simulations without Promotions - Empirical Data Scenario 2: Home Care Product

For the second scenario, a home care product with a mean demand of 2518 and a standard deviation of demand of 278 was selected. This gave the product a demand CV of 0.110. In the second scenario, lead times were longer and A/r ratios were larger. Graphs containing the average system-wide, retailer and vendor costs for IS, VMI, and CPFR in this scenario can be found in Figure 4-11.

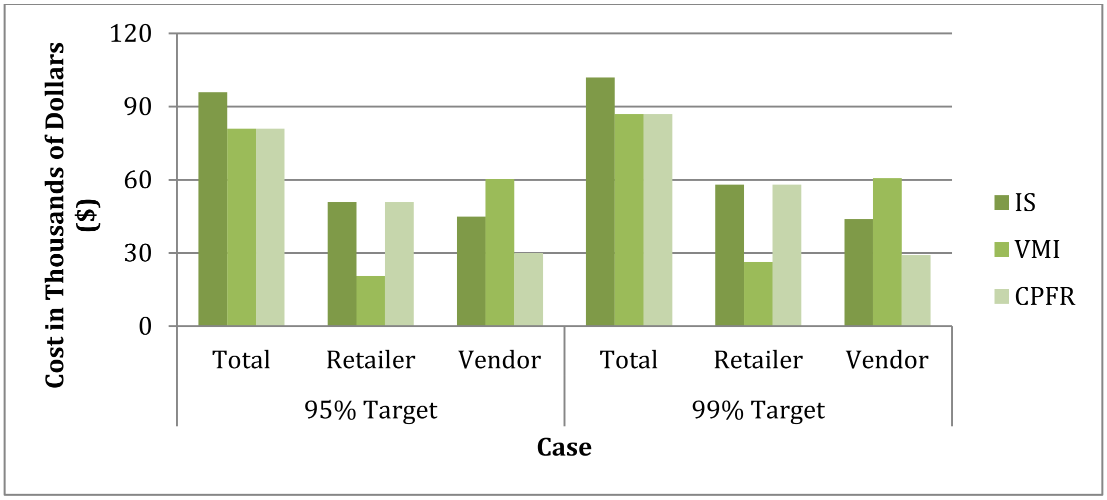

Figure 4-11: Average Costs for Simulations without Promotions - Home Care Product

As we saw in Scenario 1, despite having different lead times, A/r ratios and demand CV, IS had the greatest system-wide costs for both target P_2 fill rates. These differences in costs were statistically significant as seen by the confidence intervals in Table 4-12. As was the case in the other simulations without promotions discussed, there was no statistically significant cost difference between VMI and CPFR. However, the IS-VMI and IS-CPFR cost difference confidence intervals grew wider as the target fill rate increased which was unlike the patterns seen in the other simulations without promotions.

Table 4-12: Confidence Intervals for the Difference in System-Wide Costs ($ Thousands) in Simulations without Promotions - Home Care Product

Target Fill Rate (%)	Confidence Intervals: Cost Differences		
	IS - VMI	IS - CPFR	VMI - CPFR
95	(11.93, 16.13)	(11.97, 16.05)	(-0.60, 0.55)
99	(11.62, 11.53)	(11.90, 16.44)	(-1.15, -1.15)

The average quantity of inventory held by the retailer and vendor can be found in Figure 4-12. VMI and CPFR both held lower quantities of inventory than IS at the vendor and system-wide. The inventory carried by the retailer was similar in all three supply chain strategies.

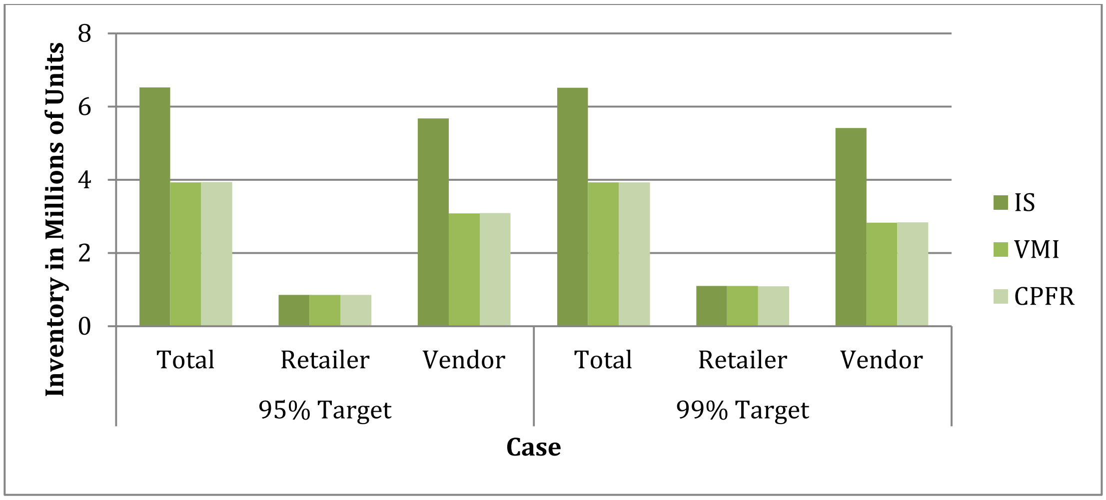

Figure 4-12: Average Inventory Levels for Simulations without Promotions - Home Care Product

Confidence intervals showed that in system-wide and vendor inventory, differences in inventory levels between IS and VMI, and IS and CPFR were statistically significant. Confidence intervals for Scenario 2 can be found in Table 4-13.

Table 4-13: Confidence Intervals for the Difference in Inventory (Millions) in Simulations without Promotions - Home Care Product

Target Fill Rate (%)	Level	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
95	System-Wide	(2.04, 2.78)	(2.09, 2.78)	(-0.09, 0.08)
	Retailer	(-0.02, 0.02)	(-0.01, 0.01)	(-0.02, 0.02)
	Vendor	(2.05, 2.78)	(2.08, 2.78)	(-0.09, 0.07)
99	System-Wide	(2.02, 2.77)	(2.05, 2.77)	(-0.08, 0.07)
	Retailer	(-0.05, 0.05)	(-0.03, 0.04)	(-0.05, 0.05)
	Vendor	(2.03, 2.76)	(2.06, 2.77)	(-0.06, 0.05)

4.2.3. Simulations with Multiple Promotions - Empirical Data Scenario 1: Stationary Product

For the first scenario, a stationary product with a regular demand mean of 213 and a regular demand standard deviation of 54 was selected. During promotions, the mean of weekly demand was 590 with a standard deviation of 213. These demand figures gave coefficient of variations of 0.254 and 0.255 respectively. Average system-wide, retailer, and vendor costs can be found in Figure 4-13.

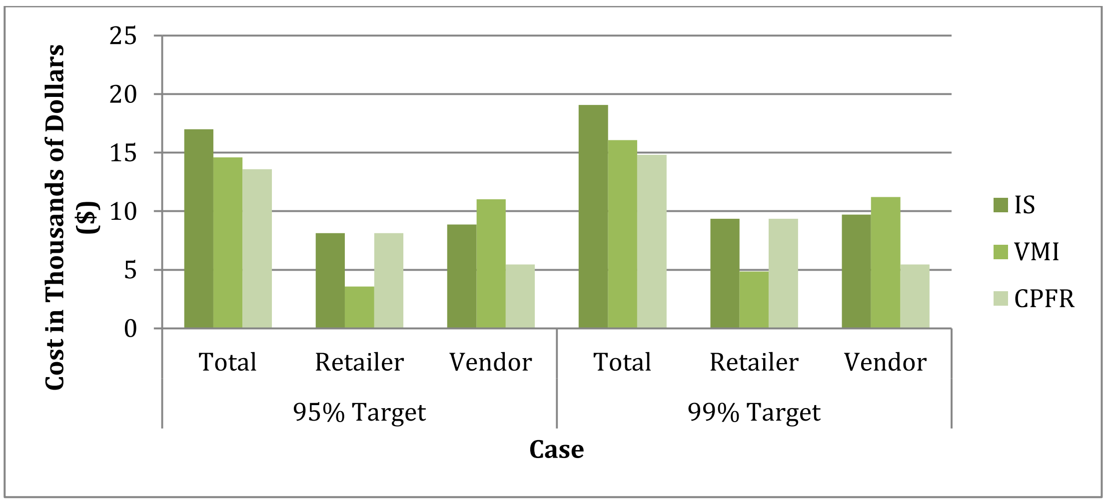

Figure 4-13: Average Costs for Simulations with Multiple Promotions - Stationary Product

Similar to the results of the simulations with promotions that used the hypothetical data in Figure 4-4, IS had the greatest system-wide cost for both target P_2 rates. VMI ranked second in system-wide costs and CPFR had the lowest. Confidence intervals confirmed that the difference in system-wide costs between IS, VMI, and CPFR were statistically significant. Table 4-14 contains the confidence intervals for the differences in system-wide costs for the stationary product. As the target fill rate increased, so too did the differences in costs between IS, VMI, and CPFR. This also occurred in the hypothetical data simulations with promotions (Table 4-9).

Table 4-14: Confidence Intervals for the Difference in System-Wide Costs ($ Thousands) in Simulations with Multiple Promotions - Stationary Product

Target Fill Rate (%)	Confidence Intervals: Cost Differences		
	IS - VMI	IS - CPFR	VMI - CPFR
95	(2.12, 2.65)	(3.15, 3.67)	(0.85, 1.19)
99	(2.68, 3.30)	(3.97, 4.55)	(1.04, 1.52)

The inventory results for the stationary product were also similar to results seen in the simulations using hypothetical data in Figure 4-8. IS carried the most amount of inventory in the supply chain, and was followed by VMI, who carried the second largest amount of inventory. CPFR carried the least amount of inventory at the system-wide and vendor level. Confidence intervals showed that the disparities in inventory levels were statistically significant. The system-wide cost and inventory confidence intervals behaved similarly to those in the simulations using hypothetical data. When the target fill rate was increased from 95% to 99%, the width of the confidence intervals widened and the intervals shifted in a positive direction. A graph showing the average amount of inventory held in IS, VMI, and CPFR, and a table with the confidence intervals for the differences in inventory levels can be found in Appendix B.

4.2.4. Simulations with Multiple Promotions - Empirical Data Scenario 2: Hardware Product

A hardware product was chosen for the second scenario of simulations with multiple promotions. The product had a regular demand mean of 306, a regular demand standard deviation of 89, and a regular demand CV of 0.291. During promotions, the weekly demand had a mean of 949, a standard deviation of 275 and a CV of 0.290. Figure 4-14 contains the average costs for IS, VMI, and CPFR. Similar to other simulations with promotions, IS had the greatest system-wide costs. CPFR provided the lowest system-wide costs while VMI fell in between IS and CPFR.

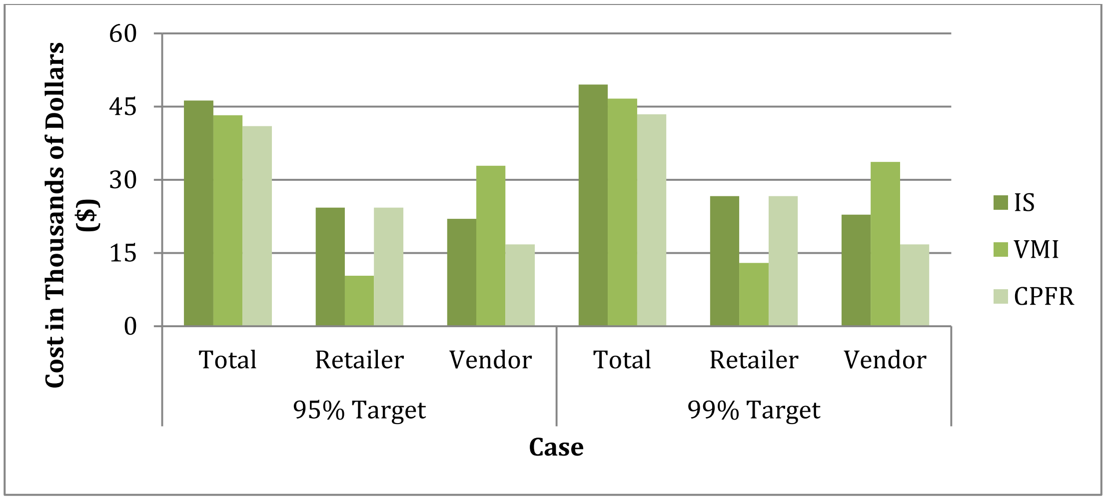

Figure 4-14: Average Costs for Simulations with Multiple Promotions - Hardware Product

The confidence intervals in Table 4-15 showed that the VMI offered a statistically significant cost reduction over IS and CPFR offered a statistically significant cost reduction over both IS and VMI in system-wide costs. When the target fill rate was increased, the width of all three confidence intervals increased. While IS-CPFR and VMI-CPFR shifted in a positive direction, IS-VMI shifted in a negative direction.

Table 4-15: Confidence Intervals for the Difference in System-Wide Costs ($ Thousands) in Simulations with Multiple Promotions - Hardware Product

Target Fill Rate (%)	Confidence Intervals: Cost Differences		
	IS - VMI	IS - CPFR	VMI - CPFR
95	(2.00, 4.04)	(4.27, 6.12)	(1.46, 2.83)
99	(1.82, 3.97)	(5.05, 7.18)	(2.34, 4.10)

Inventory results were comparable to the cost results; IS carried the largest amount of inventory system-wide, VMI held the second largest and CPFR, the lowest. The amount of inventory held by the retailers in IS, VMI ,and CPFR were alike, and the confidence intervals showed that there was no statistically significant difference between them. Just like in the simulations using hypothetical data, as seen in Table 4-4 and Table 4-6, the confidence intervals were positive, showing that there was a statistically significant difference between the amounts of inventory held system-wide and at the vendor in IS, VMI, and CPFR. Both VMI and CPFR held less inventory than IS, and CPFR held less inventory than VMI system-wide and at the vendor level. The inventory confidence intervals for system-wide and vendor inventory displayed similar patterns to the system-wide cost confidence intervals. As the target fill rate increased, the confidence intervals increased in width. Again, the IS-CPFR and VMI-CPFR confidence intervals shifted in a positive direction, while the IS-VMI confidence interval shifted in a negative direction. These confidence intervals can be found in Table 4-16 and averages inventory levels can be found in Appendix B.

Table 4-16: Confidence Intervals for the Difference in Inventory (Millions) in Simulations with Multiple Promotions - Hardware Product

Target Fill Rate (%)	Level	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
95	Total	(1.01, 1.32)	(1.39, 1.70)	(0.34, 0.42)
	Retailer	(-0.01, 0.01)	(-0.005, 0.005)	(-0.01, 0.01)
	Vendor	(1.01, 1.32)	(1.39, 1.70)	(0.34, 0.42)
99	Total	(0.93, 1.31)	(1.50, 1.84)	(0.47, 0.64)
	Retailer	(-0.02, 0.02)	(-0.01, 0.01)	(-0.02, 0.02)
	Vendor	(0.93, 1.30)	(1.50, 1.84)	(0.47, 0.64)

4.2.5. Simulations with One Large Promotion - Empirical Data Scenario 1: Stationary Product

Products in this category had the majority of their sales during a few weeks near the end of the year, i.e. the holiday season. A stationary product with this characteristic was selected for Scenario 1. During the last fourteen weeks of the year, 99.02% of its sales would occur. In the remaining weeks of the year, its weekly demand had a mean of 213 and a standard deviation of

54, giving a CV of 0.254. In Scenario 1, lead times were short and therefore, weekly replenishments would be feasible during the holiday season.

Average costs for IS, VMI, and CPFR can be found in Figure 4-15. At the lower P_2 target of 95%, IS had the greatest system-wide costs but at the higher P_2 target of 99%, VMI had the greatest system-wide costs. In both target fill rates, CPFR had the lowest system-wide costs.

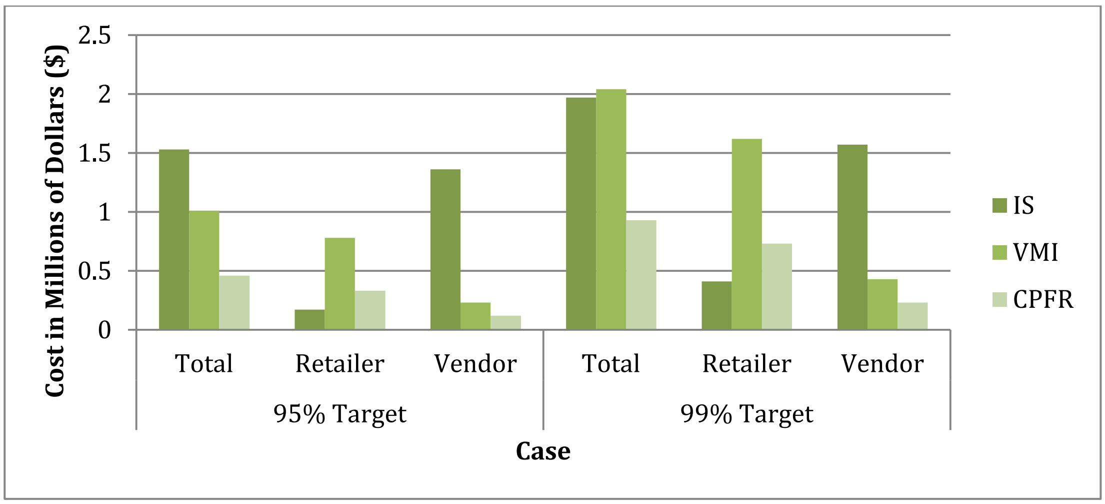

Figure 4-15: Average Costs for Simulations with One Large Promotion - Stationary Product

Confidence intervals for the differences in costs between IS, VMI, and CPFR can be found in Table 4-17. At the 95% target fill rate, IS had a higher system-wide cost than VMI and CPFR. When the target fill rate was increased to 99%, VMI no longer had a statistically significant, system-wide cost improvement over IS.

Table 4-17: Confidence Intervals for the Difference in System-Wide Costs ($ Millions) in Simulations with One Large Promotion - Stationary Product

Target Fill Rate (%)	Confidence Intervals: Cost Differences		
	IS - VMI	IS - CPFR	VMI - CPFR
95	(0.27, 0.76)	(0.90, 1.22)	(0.20, 0.90)
99	(-0.39, 0.24)	(0.79, 1.27)	(0.63, 1.58)

Average levels of inventory held at the retailer, vendor and system-wide can be found in Figure 4-16. For both target P_2 rates, IS held the largest amount of inventory in the supply chain. VMI carried the most inventory at the retailer, while IS carried the least. At the vendor level, IS carried the most inventory and CPFR, the least amount of inventory.

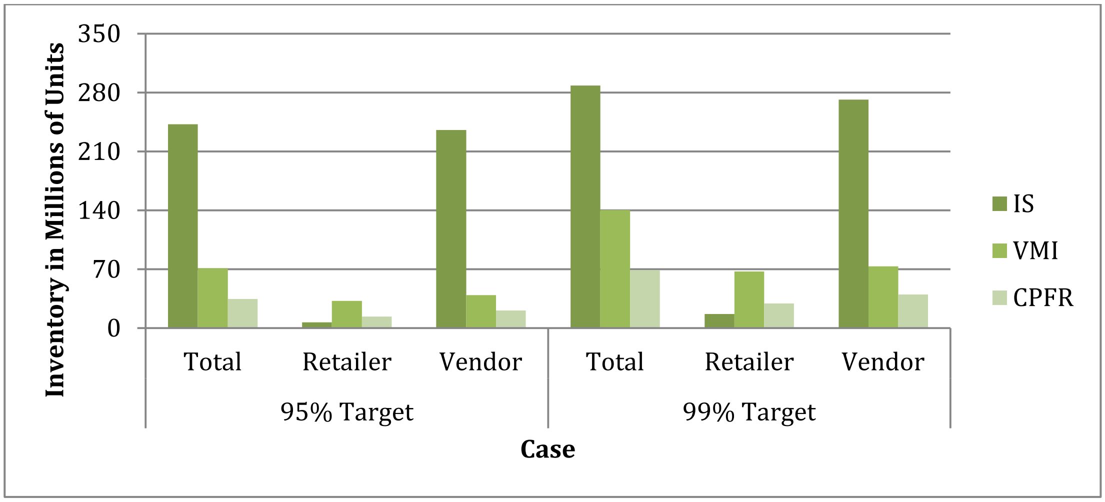

Figure 4-16: Average Inventory Levels for Simulations with One Large Promotion - Stationary Product

Confidence intervals, found in Table 4-18, confirmed the that the differences in inventory levels were statistically significant. CPFR offered statistically significant reductions to the amount of inventory in the supply chain when employed over IS. However, inventory at the retailer increased with CPFR. VMI also provided system-wide inventory reductions over IS but once again, inventory at the retailer increased. The VMI-CPFR confidence interval was positive at the retailer, vendor, and system-wide. CPFR offered inventory reduction over VMI for both target fill rates tested.

Table 4-18: Confidence Intervals for the Difference in Inventory (Millions) in Simulations with One Large Promotion - Stationary Product

Target Fill Rate (%)	Level	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
95	System-Wide	(158.07, 183.23)	(186.50, 226.68)	(13.16, 60.55)
	Retailer	(-35.96, -15.27)	(-11.00, -3.33)	(7.33, 30.03)
	Vendor	(185.15, 206.61)	(195.56, 232.21)	(5.84, 30.73)
99	System-Wide	(134.06, 162.06)	(189.10, 247.15)	(39.60, 104.47)
	Retailer	(-64.33, -36.53)	(-17.49, -7.86)	(22.36, 53.83)
	Vendor	(184.96, 211.87)	(204.50, 256.51)	(16.76, 50.84)

4.2.6. Simulations with One Large Promotion - Empirical Data Scenario 2: Hardware Product

For the Scenario 2, a hardware product that exhibited a large spike in sales was chosen. This product experienced 51.49% of its sales during seven weeks of the year. During the remainder of the year, weekly demand had a mean of 972 and a standard deviation of 261, which gave it a CV of 0.269. Longer lead times in Scenario 2 made weekly replenishments impractical, as any updates made to the demand forecast during the holiday season could not be utilized by the vendor. Instead, one order before the start of the holiday season was placed for the anticipated increase in demand.

Figure 4-17 contains the average costs for IS, VMI, and CPFR in Scenario 2. VMI had the greatest average system-wide costs for both target P_2 rates. VMI, was followed by IS and then CPFR. Table 4-19 contains the confidence intervals for the differences in system-wide costs. Though VMI's average system-wide cost was higher than that of IS, the difference in costs was not statistically significant as shown by the confidence interval. However, the differences in system-wide costs between CPFR and both IS and VMI were statistically significant.

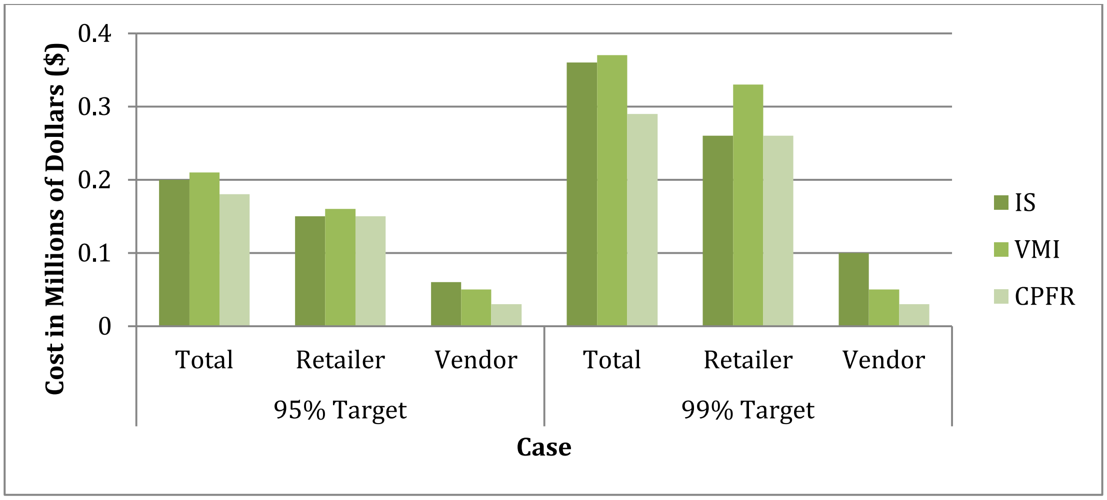

Figure 4-17: Average Costs for Simulations with One Large Promotion - Hardware Product

Table 4-19: Confidence Intervals for the Difference in System-Wide Costs ($ Thousands) in Simulations with One Large Promotion - Hardware Product

Target Fill Rate (%)	Confidence Intervals: Cost Differences		
	IS - VMI	IS - CPFR	VMI - CPFR
95	(-29.77, 14.41)	(16.45, 35.78)	(3.74, 63.75)
99	(-60.23, 39.45)	(31.94, 129.48)	(4.52, 162.06)

The inventory results were similar to the cost results. There were no statistically significant difference between the amount of inventory held in the supply chain by IS and the amount of inventory held by VMI. CPFR offered statistically significant reductions in supply chain inventory over IS and VMI. The inventory confidence intervals can be found in Table 4-20. VMI resulted in more inventory being held at the retailer than IS and CPFR for both target fill rates. There was no difference between the amount of inventory held by the retailers in IS and CPFR. At the vendor level, both VMI and CPFR carried less inventory than IS. There was also no difference between the inventory held by the vendors in VMI and CPFR.

Table 4-20: Confidence Intervals for the Difference in Inventory (Thousands) in Simulations with One Large Promotion - Hardware Product

Target Fill Rate (%)	Level	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
95	System-Wide	(2.75, 4.42)	(3.46, 7.03)	(0.09, 3.16)
	Retailer	(-2.78, -0.23)	(-0.07, 0.06)	(0.23, 2.79)
	Vendor	(3.56, 6.74)	(3.47, 7.02)	(-0.25, 0.51)
99	System-Wide	(5.13, 16.89)	(6.25, 23.33)	(0.05, 7.32)
	Retailer	(-6.99, -0.32)	(-0.13, 0.12)	(0.29, 6.96)
	Vendor	(6.40, 23.16)	(6.21, 23.32)	(-0.57, 0.83)

4.2.7. Summary of Results for Simulations using Empirical Data

Table 4-21 and Table 4-22 contain summaries of the system-wide cost and inventory results from the simulations using empirical data. The table outlines whether the confidence interval was positive or negative, if the confidence interval contained zero, and if the results were consistent for both scenarios and both target fill rates tested.

Table 4-21: Summary of System-Wide Cost Results in Simulations using Empirical Data

Case	Result	Confidence Interval: Cost Differences		
		IS-VMI	IS-CPFR	VMI-CPFR
No Promotions	Positive or Negative	Positive	Positive	Neither
	Contains 0	No	No	Yes
	For All Cases	Yes	Yes	Yes
Multiple Promotions	Positive or Negative	Positive	Positive	Positive
	Contains 0	No	No	No
	For All Cases	Yes	Yes	Yes
One Large Promotion	Positive or Negative	Mostly Neither	Positive	Positive
	Contains 0	Mostly Yes	No	No
	For All Cases	No, in Scenario 1, when target P_2 is 0.95, the CI is positive	Yes	Yes

From Table 4-21, in simulations with one large promotion, the IS-VMI confidence interval was for the most part, neither positive or negative. In most cases, the confidence interval contained zero, indicating that the difference between IS and VMI's system wide costs were not statistically significant. This result held for all but one case; in Scenario 1, when the target fill rate was 95%, the confidence interval was positive, that is, both the upper and lower bounds of the confidence interval were positive numbers. Therefore, in that particular case, there was a statistically significant reduction in system-wide costs when VMI was chosen over IS.

Table 4-22: Summary of System-Wide Inventory Results in Simulations using Empirical Data

Case	Result	Confidence Interval: Inventory Differences		
		IS-VMI	IS-CPFR	VMI-CPFR
No Promotions	Positive or Negative	Positive	Positive	Neither
	Contains 0	No	No	Yes
	For All Cases	Yes	Yes	Yes
Multiple Promotions	Positive or Negative	Positive	Positive	Positive
	Contains 0	No	No	No
	For All Cases	Yes	Yes	Yes
One Large Promotion	Positive or Negative	Positive	Positive	Positive
	Contains 0	No	No	No
	For All Cases	Yes	Yes	Yes

Summaries of the patterns found in the confidence intervals for system wide cost and inventory levels can be found in Table 4-23 and Table 4-24. The tables detail what happens to the confidence intervals as the target fill rate is increased from 95% to 99%. A confidence interval can either increase or decrease in width and shift in a positive or negative direction, or not shift at all.

As shown in Table 4-23, in the case with multiple promotions, in Scenario 1, the IS-VMI confidence interval widens when a higher target fill rate is specified. The confidence interval shifts in a positive direction as well. In Scenario 2, the IS-VMI confidence interval widens as it did in Scenario 1. However, instead of shifting in a positive direction, the confidence interval shifts in a negative direction when the target fill rate is increased.

Table 4-23: Patterns in Confidence Intervals for System Wide Costs - Simulations using Empirical Data

Case	Scenario	Confidence Interval: Cost Differences		
		IS-VMI	IS-CPFR	VMI-CPFR
No Promotions	1	• interval narrows • shifts positively	• interval narrows • shifts positively	• interval widens • no shift
	2	• interval widens • no shift	• interval widens • no shift	• interval widens • no shift
Multiple Promotions	1	• interval widens • shifts positively	• interval widens • shifts positively	• interval widens • shifts positively
	2	• interval widens • shifts negatively	• interval widens • shifts positively	• interval widens • shifts positively
One Large Promotion	1	• interval widens • shifts negatively	• interval widens • no shift	• interval widens • shifts positively
	2	• interval widens • shifts negatively	• interval widens • shifts positively	• interval widens • shifts positively

Table 4-24: Patterns in Confidence Intervals for System Wide Inventory Levels - Simulations using Empirical Data

Case	Scenario	Confidence Interval: Inventory Differences		
		IS-VMI	IS-CPFR	VMI-CPFR
No Promotions	1	• interval narrows • shifts positively	• interval narrows • shifts positively	• interval widens • no shift
	2	• interval widens • shifts negatively	• interval widens • shifts negatively	• interval narrows • no shift
Multiple Promotions	1	• interval widens • shifts positively	• interval widens • shifts positively	• interval widens • shifts positively
	2	• interval widens • shifts negatively	• interval widens • shifts positively	• interval widens • shifts positively
One Large Promotion	1	• interval widens • shifts negatively	• interval widens • shifts positively	• interval widens • shifts positively
	2	• interval widens • shifts positively	• interval widens • shifts positively	• interval widens • shifts positively

In the next chapter, we will analyse the results presented in this chapter, discuss their implications for practitioners of CPFR and make recommendations.

Chapter 5: Analysis and Discussion

5.1. Simulations without Promotions

In simulations without promotions, both VMI and CPFR reduced inventory and lowered costs in the supply chain. In simulations using hypothetical data, this result held for the two target fill rates and for the three levels of demand variability that were tested (Figure 4-3, Figure 4-7). In simulations using the empirical data, both scenarios gave this result despite the different demand parameters, lead times, and A/r ratios being used. This result of reduced inventory and lowered costs was expected, as it mirrors case studies found in literature. When given access to the retailer's POS and inventory data, the vendors in VMI and CPFR were able to better anticipate when the retailer would order. This allowed those vendors to reduce the amount of inventory they carried to meet the retailers' needs.

As demand variability increased in the simulations using hypothetical data, the IS-VMI and IS-CPFR confidence intervals for system-wide costs and inventory shifted in a positive direction (Table 4-9). This indicated that VMI and CPFR's ability to reduce costs and inventory improved as demand variability increased. In most cases, a positive shift in the IS-VMI and IS-CPFR confidence intervals also occurred when the target fill rate was increased from 95% to 99%. However, this did not occur in Scenario 2 of the simulations run with empirical data due to this case's low demand CV. From this result, we can conclude that the effectiveness of VMI and CPFR at reducing costs and inventory diminishes when demand variability is very low. At low levels of demand variability, the vendor in IS does not struggle as much to forecast demand and the benefits of the information shared in VMI and CPFR decreases.

There was no reduction in the retailer's inventory when implementing VMI or CPFR (Figure 4-9), but this was expected as information was only shared in one direction, from the retailer to the vendor. In our simulation model, the vendor does not have insights into end-customer demand that could be shared with the retailer. It is assumed that the savings experienced by the vendors in VMI and CPFR will be shared with the retailer, but how those savings are shared is outside of the scope of this research.

While VMI and CPFR performed better than IS, there was no distinction between the savings offered by VMI and the savings offered by CPFR (Table 4-1, Table 4-3). There was no statistically significant difference between the system-wide costs of VMI and CPFR, nor was there a statistically significant difference between the amounts of inventory held by each supply chain. The vendors in VMI and CPFR were privy to the same information. Neither had access to data that the other did not. In the case without promotions, CPFR thus, does not offer an advantage over VMI.

Other than its abilities to reduce costs in a supply chain, when promotions are not present, either VMI or CPFR would appeal to vendors for their ability to reduce inventory levels and to retailers for their ability to reduce the occurrence of backorders from the vendor. However, there are aspects of CPFR and VMI not accounted for in this research that may lead a practitioner to favour one over the other. For example, VMI could in practice, offer an cost advantage over CPFR; when the vendor orders on behalf of the retailer, it is assumed that the vendor can place orders with itself at a lower cost than the retailer can with the vendor. Since the exact reduction in ordering cost will depend on the particular situation, we did not lower the ordering cost, and instead used the retailer's ordering cost as an upper limit of what the ordering cost would be for the vendor. While VMI may offer lower ordering costs, CPFR could, in practice, offer lower transportation costs. Several CPFR case studies noted lowered transportation costs as a result of advance notice of replenishments. By knowing beforehand, the size and timing of each replenishment, the vendors were able to more accurately forecast their transportation needs and reduce the amount of rushed shipments (Seifert, 2003; Cederlund et al. 2007). Depending on the situation, it is conceivable that either VMI or CPFR could be less costly than the other.

5.2. Simulations with Promotions

5.2.1. Simulations with Multiple Promotions

In simulations with multiple promotions, VMI was able to provide reductions to cost and inventory over IS in most cases. In simulations using hypothetical data, there was only two cases where this was not so. When demand variability was low and for both target fill rates, there was no statistically significant difference between the system-wide costs of IS and VMI. However, in this case, VMI provided a statistically significant reduction in system-wide inventory over IS. The vendor in VMI placed more orders with the manufacturer than the vendor in IS. In VMI, the vendor knew that promotions would eventually occur, and therefore, placed extra orders in preparation. When demand variability was low, the increase in ordering costs was not always offset by the decrease in the vendor's inventory costs, which produced confidence intervals that contained zero. When demand variability was higher, the vendor in VMI was more effective at reducing inventory costs and this more than compensated for the increase in ordering costs. In simulations using the empirical data, VMI was able to reduce system-wide inventory levels and costs in all scenarios.

When simulating with the hypothetical data, VMI's performance over IS improved when the target fill rate was increased from 95% to 99%. The IS-VMI confidence intervals for the differences in costs and in inventory levels shifted in a positive direction as the target fill rate was increased. This result was also found in Scenario 1 when simulating with empirical data (where A/r ratios were lower and lead times shorter). However, in the second scenario simulated with empirical data, where A/r ratios were higher and lead times longer than in Scenario 1, VMI's cost and inventory performance over IS fell as the target fill rate was increased. In Scenario 2, demand during promotions was approximately three times larger than regular demand. In the first scenario simulated with empirical data and in the simulations using hypothetical data, demand during promotions was approximately double the regular demand. As promotions grew more extreme and higher fill rates were desired, VMI's ability to reduce costs diminished.

While VMI was able to provide cost and inventory savings, it could not match the performance of CPFR when multiple promotions were present in the simulations. CPFR outperformed VMI and IS in all the cases tested, reducing both costs and inventory levels in the supply chain. In simulations using the hypothetical data, CPFR offered lower costs and decreased inventory for the low, medium, and high levels of demand variability and for the 95% and 99% P_2 target fill rates that were examined. In simulations that used the empirical data, despite different lead times and demand and cost parameters being employed, the results were similar to those of the simulations using hypothetical data. CPFR lowered the costs and levels of inventory in the supply chain for both target P_2 rates simulated. While both the retailer and vendor would be attracted to CPFR's ability to reduce costs in the supply chain, the vendor would also prefer it over IS and VMI for its ability to reduce inventory levels when promotions are present. The retailer would also prefer CPFR as it greatly improves the vendor's abilities to meet the retailer's orders, ensuring that end customer demand can be met.

In all cases tested, in simulations using hypothetical data as well as simulations using empirical data, CPFR's ability to reduce costs and inventory over IS and VMI improved as the target P_2 fill rate was increased. As demand variability was increased in the simulations using hypothetical data, the IS-CPFR and VMI-CPFR confidence intervals shifted in a positive direction. That is, the upper and lower bounds of the confidence intervals increased in value, indicating that CPFR can be especially valuable to companies facing highly variable demand.

In all simulations, whether using hypothetical data or empirical data, the vendor in VMI was not given access to the same information that the vendor in CPFR was given. While the vendor in CPFR was aware of the retailer's promotions plans far in advance, the vendor in VMI was not. The vendor in CPFR knew exactly when promotional periods were approaching and could order extra inventory accordingly. This reduced the amount of inventory the vendor in CPFR had to carry. The vendor in VMI knew that the retailer had a promotion planned for each quarter of the year, but exactly when the promotion would occur in the quarter was not shared. The vendor in VMI had to carry more inventory to ensure it would have sufficient stock when a promotion occurred. While able to lower inventory levels and costs with the aid of the retailer's point-of-sale and inventory data, VMI could not operate as efficiently as CPFR since the vendor in VMI was unaware of the timing of promotions. This outcome is in agreement with observations found in literature, that some retailers ended their VMI implementation due to its inability to effectively handle promotions (Sari, 2008a).

5.2.2. Simulations with One Large, Seasonal Promotion

In the second promotional demand pattern simulated using the empirical data, there was one large, seasonal promotion wherein the majority of the year's demand occurred. In these simulations, the system-wide operating costs of CPFR were lower than that of VMI and IS. Utilizing CPFR also resulted in less inventory being carried in the supply chain. The vendor in CPFR was able to use the information provided by the retailer to order adequately for the anticipated holiday seasons. The vendor in CPFR was forewarned that the retailer expected higher than or lower than average sales, and how much the retailer planned to order throughout the holiday season.

In our simulations, sharing the demand forecast with the vendor had a small negative effect on the retailer. It was assumed that more recent demand information created more accurate demand forecasts than older information, and in IS, the retailer could alter their demand forecast right up until the moment an order was placed. However, in CPFR the retailer froze its forecast earlier, and was unable to incorporate the most recent demand information. Since the vendor needed to order weeks ahead of the retailer, the demand forecast was frozen to ensure that both parties were working from the same information. The retailer in CPFR carried more inventory than the retailer in IS, but the negative effect was small, and savings throughout the supply chain more than compensated for this increase. Whether or not this negative effect occurs in practice will depend on the particular situation. Despite this small increase in inventory, the retailer would prefer CPFR over IS or VMI as it would result in more reliable replenishments from the vendor which is especially important when the majority of demand occurs over a few weeks of the year. As was the case when multiple promotions were present, the vendor would also prefer CPFR for its abilities to reduce inventory levels.

Forecasts were frozen earlier in time during simulations with multiple promotions, but the retailer's forecast was not as negatively impacted. Unlike simulations with one large, seasonal promotion, demand was not rapidly changing. Instead, demand was relatively steady, and having the most up-to-date demand information was not necessary to forecast demand.

In simulations with one large, seasonal promotion, VMI was rarely able to reduce costs. Only in Scenario 1, when the target fill rate was 95%, was there a statistically significant reduction in costs when VMI was chosen over IS. In all other cases, there was no statistically significant difference between the system-wide costs of VMI and IS. Yet, VMI was able to reduce the amount of inventory carried in the supply chain. The higher costs were a result of larger amounts of stock being held at the retailer rather than at the vendor, who had lower holding costs. Being unaware of the retailer's promotional plans, the vendor sent large amounts of inventory to the retailer to ensure that the target fill rates were attained. It is unlikely that in practice, a retailer would be pleased with this result, but to achieve the desired target rates, the vendor had little choice but to send large amounts of inventory to the retailer in case there were higher-than-average sales.

The disadvantage that can arise from withholding information can be clearly seen in the simulations with one large seasonal promotion. Despite knowing that sales would dramatically increase around the holidays each year, larger amounts of stock were necessary in VMI to reach the target fill rates. It may seem innocuous to a retailer to hold back information regarding how aggressive their marketing campaigns or sale prices will be, but the results of our simulations show that this can hamper the success of a supply chain strategy. If a vendor is unsure of how large each upcoming promotion will be or when the next promotion will occur, its only recourse is to hold extra inventory to safeguard against these possibilities.

5.3. Conclusions

In our research, we found that VMI and CPFR would perform similarly in cases without promotions. That is, when the information shared with the vendors in CPFR was the same as the information shared with the vendors of VMI, there was no statistically significant difference between the performance of the two methodologies. When promotions were present, we found that CPFR outperformed VMI, reducing inventory levels and costs in the supply chain. In the case with promotions, the retailer shared more information with the vendor in CPFR than with the vendor in VMI. In the case with multiple promotions, the vendor in VMI was unaware of the timing of promotions and in the case with one large, seasonal promotion, the vendor was unaware of the intensity of the promotions. In both of these cases, because information was withheld from the vendor, VMI could not perform as efficiently as CPFR.

In the research literature, the distinction between VMI and CPFR is unambiguous; they are defined as two unique information sharing, supply chain methodologies. VMI does not involve sharing forecasts and promotional plans, and the vendor in CPFR does not place orders on behalf of the retailer. However, in practice, the lines between VMI and CPFR can become blurred. Speaking with a CPFR practitioner, our client in this research, it was discovered that their CPFR partnership had elements of a VMI implementation. In practice, where a VMI strategy ends and where CPFR begins can be difficult to determine. As trust is an important factor in the success of CPFR, it may be difficult to properly implement CPFR without having a prior partnership. As trust between a retailer and vendor grows, a retailer in VMI may share more information, such as demand forecasts and promotional plans, with the vendor without fear that it will be leaked or shared with competitors.

Regardless of what methodology a supply chain utilizes, whether it be VMI, CPFR, or a hybridization of the two, complete information sharing is crucial. CPFR does not automatically provide better results than VMI. As was shown in simulations without promotions, when VMI and CPFR were given access to the same information, they achieved similar reductions in costs and inventory. Rather, the results of the simulations with promotions have shown that withholding information about upcoming demand will hamper the ability of a supply chain to operate efficiently. CPFR cannot live up to its promises if supply chain partners withhold valuable information from one another. Companies wishing to start CPFR partnerships must be prepared to share information freely. If the trust required to do so does not yet exist between the supply chain partners, supply chain strategies that do not require as much information sharing should be considered to help build that necessary trust.

5.4. Recommendations

The simulation models can be improved through additional validation. Currently, the simulation has been tested for face validity, but comparing the model's output with an existing system and using quantitative techniques would greatly enhance its credibility (Law, 2007). Obtaining output data from an existing system to compare with the simulation model's results will be our next task in this research.

The product that was simulated in our model had an infinite shelf life and therefore, best-before dates were not taken into consideration. There are many products, especially in the grocery industry, that have a short shelf life, which could greatly affect the simulation's results. The penalties for holding large amounts of inventory would be much higher if the product could not be sold after a period of time. Since supply chain methodologies like VMI and CPFR are of great interest to the highly competitive grocery industry, accounting for easily perishable products would be worth undertaking.

There are a variety of demand patterns that have not been tested in this simulation model, including demand with increasing or decreasing trends. The demand tested in the model for this research was normally distributed, steady, and without any trends. It would be beneficial to run the simulation model with different demand patterns, from different distributions to determine if the same conclusions can reached in these difference situations. There are also many activities that could impact future demand other than promotions. The introduction of a new product, the discontinuation of a product, and the promotion of a complementary product or a substitute product can impact demand. The impact may be negative, rather than positive, as was the case in our research.

Our model only explored VMI and CPFR in a supply chain with a single retailer. Examining supply chains with more retailers would be beneficial to vendors serving multiple customers. The vendor may not find CPFR to be as beneficial as we found it to be if the timing of the retailers' promotions are not correlated; that is, promotions at different retailers are not likely to occur at the same time. If this is the case, aggregating the demand from multiple retailers may result in the vendor experiencing a less variable demand. However, if the retailers' promotions are likely to occur at the same time, as may be the case in demand that is seasonal, the swings in demand would likely become more intense. Situations such as this may greatly benefit from implementing CPFR.

References

Aichlmayr, M., 2000. Who manages inventory in a Value Chain? Transportation & Distribution, October (2000), 60-66.

Aviv, Y., 2001. The Effect if Collaborative Forecasting on Supply Chain Performance. Management Science, 47 (10), 1326-1343.

Aviv, Y., 2002. Gaining benefits from joint forecasting and replenishment processes: The Case of Auto-Correlated Demand. Manufacturing & Service Operations Management, 4 (1), 55-74.

Aviv, Y., 2007. On the Benefits of Collaborative Forecasting Partnerships Between Retailers and Manufacturers. Management Science, 53 (5), 777-794.

Banks, J., Carson, J.S., Nelson, B.L., and Nicol, D.M., 2005. Discrete-Event System Simulation (4th Ed). Pearson Prentice Hall.

Barratt, M., and Oliveira, A., 2001. Exploring the Experiences of Collaborative Planning Initiatives. International Journal of Physical Distribution & Logistics Management, 31 (4), 266-289.

Boone, Y., Ganeshan, R., and Stenger, A., 2002. The Benefits of Information Sharing in a Supply Chain: An Exploratory Simulation Study. In: Geunes, J., Pardalos, P., Romeijn, E.(Eds.), Supply Chain Management Models, Applications and Research Directions. Kluwer Publishers, Dordrecht.

Cederlund, J.P., Kohli, R., Sherer, S.A., and Yao, Y., 2007. How Motorola put CPFR into Action. Supply Chain Management Review, October (2007), 28-35.

Chung, W.W.C., and Leung, S.W.F., 2005. Collaborative planning, forecasting and replenishment: a case study in copper clad laminate industry. Production Planning & Control, 16 (6), 563-574.

Danese, P., 2010. Towards a contingency theory of collaborative planning initiatives in supply networks. International Journal of Production Research, 49 (4), 1081-1103.

Denend, L., and Lee H., 2005. West Marine: Driving Growth through Shipshape Supply Chain Management. Case: GS-34, Stanford Graduate School of Business: 35.

Disney, S.M., Naim, M.M., and Potter, A., 2004. Assessing the Impact of e-Business on Supply Chain Dynamics. International Journal of Production Economics, 89, 109-118.

Fu, H.P., Lin, S.W., and Chen, C.R., 2009. A Study on Impact Factors of Retailing Implementing CPFR – A Fuzzy AHP Analysis. Global Perspective for Competitive Enterprise, Economy and Ecology, 2009, (11), 609-619.

Harrington, L., 2003. 9 Steps to Success with CPFR. Transportation & Distribution, April (2003), 50-52.

Hill, S., 1999. CPFR Builds the United Partnerships of Apparel. Apparel Industry Magazine, 1999, 60, 54-58.

Hill, S., Jr., and Mathur, S., 1999. A Net Gain. Manufacturing Systems, 17 (11), 34-43.

Ireland, R., and Bruce, R. 2000., CPFR: Only the Beginning of Collaboration. Supply Chain Management Review, September/October (2000), 80-88.

Law, A. M., 2007. Simulation modeling and analysis (4th Ed). New York: McGraw-Hill.

Lee, H.L., Padmanabhan, V., and Whang, S., 1997. Information Distortion in a Supply Chain: The Bullwhip Effect. Management Science, 43 (4), 546-558.

Min, H., & Yu, W.B.V., 2008. Collaborative planning, forecasting and replenishment: demand planning in supply chain management. International Journal of Information Technology and Management, 7(1), 4-20.

Raghunathan, S., 1999. Interorganizational Collaborative Forecasting and Replenishment Systems and Supply Chain Implications. Decision Sciences, 30 (4), 1053-1071.

Sari, K., 2008a. On the benefits of CPFR and VMI: A comparative simulation study. International Journal of Production Economics, 113, 575-586.

Sari, K., 2008b. Inventory inaccuracy and performance of collaborative supply chain practices. Industrial Management & Data Systems, 108, (4), 495-509.

Seifert, D., 2003. Collaborative Planning, Forecasting and Replenishment: How to Create a Supply Chain Advantage. New York: AMACOM.

Silver, E.A., Pyke, D.F., and Peterson, R., 1998. Inventory management and production planning and scheduling (3rd Ed). New York: Wiley.

Silver, E. A., Naseraldin, H., and Bischak, D.P., 2009. Determining the reorder point and order-up-to level in a periodic review system so as to achieve a desired fill rate and a desired average time between replenishments. Journal of the Operational Research Society, 60 (9), 1244-1253.

Smaros, J., 2007. Forecasting Collaboration in the European Grocery Sector: Observations from a Case Study. Journal of Operations Management, 25, 702-716.

Steermann, H., 2003. A Practical Look at CPFR: The Sears-Michelin Experience. Supply Chain Management Review July/August (2003), 46-53.

Waller, M.A., Johnson, M.E., and Davis, T., 1999. Vendor Managed Inventory in the Retail Supply Chain. Journal of Business Logistics 20 (1), 183-203.

Appendix A: Additional Results for Simulations using Hypothetical Data

Additional Results for Section 4.1.1 Stockout Results for Simulations using Hypothetical Data

Table A-1: Confidence Intervals for the Difference in Retailer Stockout Rates (%) for Simulations without Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Stockout Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(-0.75, 0.69)	(-0.50, 0.52)	(-0.74, 0.73)
	99	(-0.52, 0.50)	(-0.32, 0.41)	(-0.48, 0.52)
Medium	95	(-0.98, 1.00)	(-0.75, 0.75)	(-0.98, 0.98)
	99	(-0.67, 0.66)	(-0.56, 0.52)	(-0.67, 0.69)
High	95	(-1.06, 1.17)	(-0.75, 0.77)	(-1.11, 1.11)
	99	(-0.80, 0.76)	(-0.64, 0.59)	(-0.77, 0.79)

Table A-2: Confidence Intervals for the Difference in Retailer Stockout Rates (%) for Simulations with Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Stockout Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(-0.68, 0.69)	(-0.55, 0.62)	(-0.70, 0.68)
	99	(-0.56, 0.57)	(-0.46, 0.39)	(-0.60, 0.55)
Medium	95	(-0.91, 0.91)	(-0.66, 0.66)	(-0.91, 0.90)
	99	(-0.67, 0.68)	(-0.51, 0.55)	(-0.68, 0.67)
High	95	(-0.99, 0.95)	(-0.66, 0.65)	(-1.01, 1.03)
	99	(-0.68, 0.71)	(-0.54, 0.58)	(-0.72, 0.71)

Additional Results for Section 4.1.3 Inventory Results for Simulations using Hypothetical Data

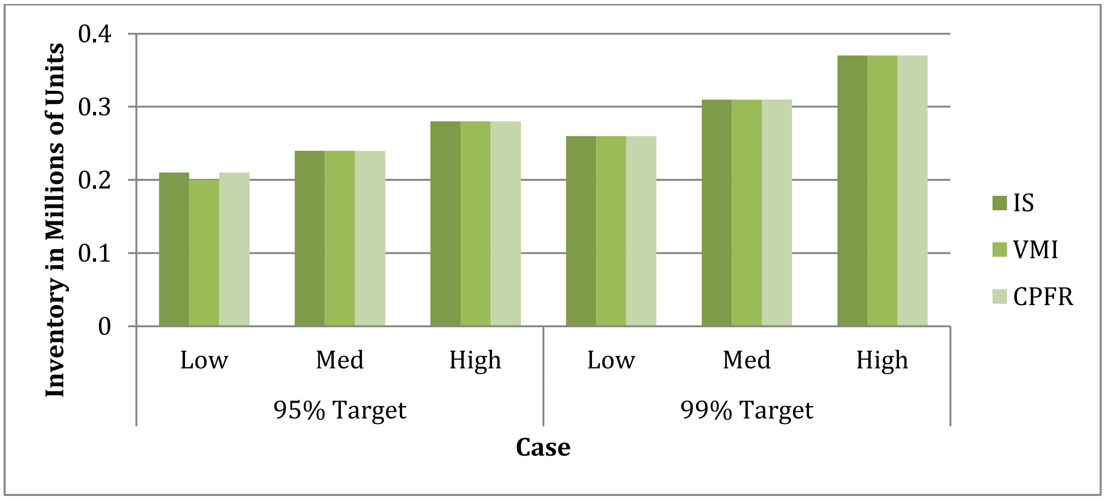

Figure A-1: Average Inventory held by the Retailer in Simulations with Promotions - Hypothetical Data

Table A-3: Confidence Intervals for the Difference in Inventory held at the Retailer (Thousands) in Simulations without Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(-3.84, 4.04)	(-2.20, 2.05)	(-3.83, 3.76)
	99	(-7.18, 7.15)	(-4.82, 4.47)	(-7.31, 7.16)
Medium	95	(-5.98, 6.57)	(-3.06, 3.20)	(-6.15, 5.99)
	99	(-8.88, 9.15)	(-5.37, 6.04)	(-8.76, 8.83)
High	95	(-7.55, 7.58)	(-4.00, 4.08)	(-7.81, 7.78)
	99	(-10.90, 10.66)	(-6.69, 6.75)	(-11.10, 10.73)

Table A-4: Confidence Intervals for the Difference in Inventory held at the Retailer (Thousands) in Simulations with Promotions - Hypothetical Data

Demand Variability	Target Fill Rate (%)	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
Low	95	(-3.86, 4.00)	(-1.92, 1.88)	(-4.07, 3.77)
	99	(-7.27, 7.57)	(-4.25, 4.61)	(-7.55, 7.64)
Medium	95	(-6.82, 6.25)	(-3.79, 3.55)	(-6.42, 6.31)
	99	(-10.85, 10.98)	(-6.38, 6.06)	(-10.96, 10.35)
High	95	(-8.83, 8.62)	(-4.60, 4.47)	(-8.79, 8.60)
	99	(-12.87, 13.95)	(-8.37, 8.45)	(-13.40, 13.35)

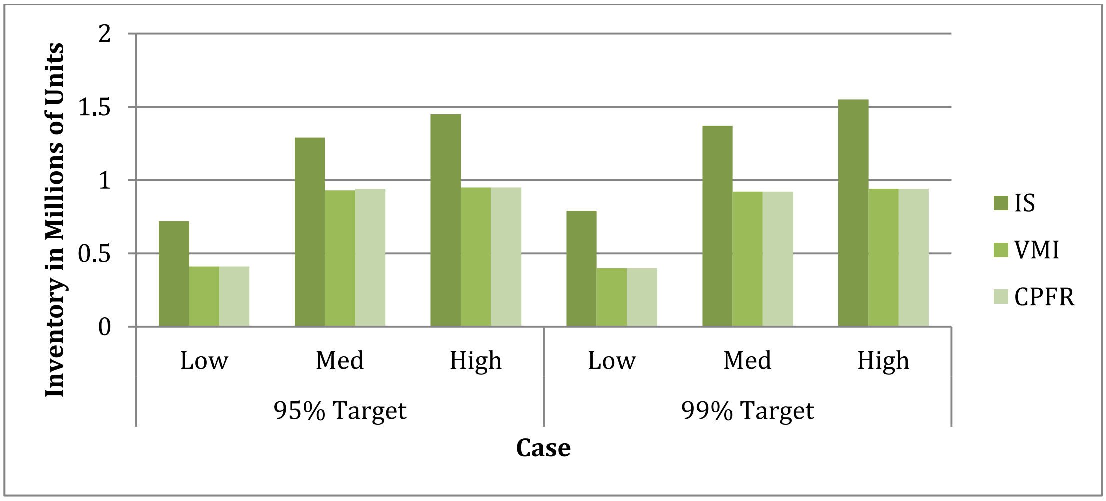

Figure A-2: Average Inventory held by the Vendor in Simulations without Promotions - Hypothetical Data

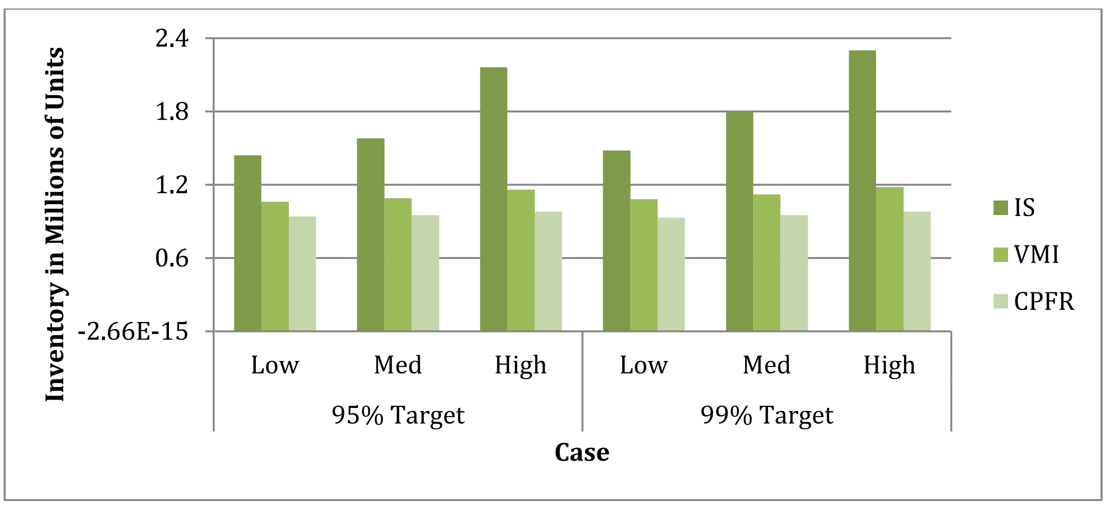

Figure A-3: Average Inventory held by the Vendor in Simulations with Promotions - Hypothetical Data

Appendix B: Additional Results for Simulations using Empirical Data

Additional Results for Section 4.2 Results for Simulations using Empirical Data

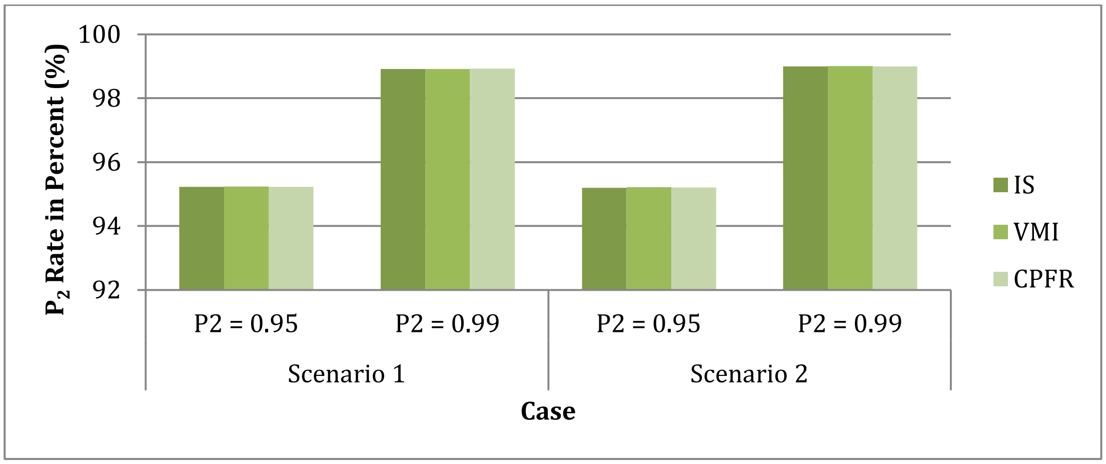

Figure B-1: P_2 Service Rates Achieved by the Retailer in the Simulations without Promotions - Empirical Data

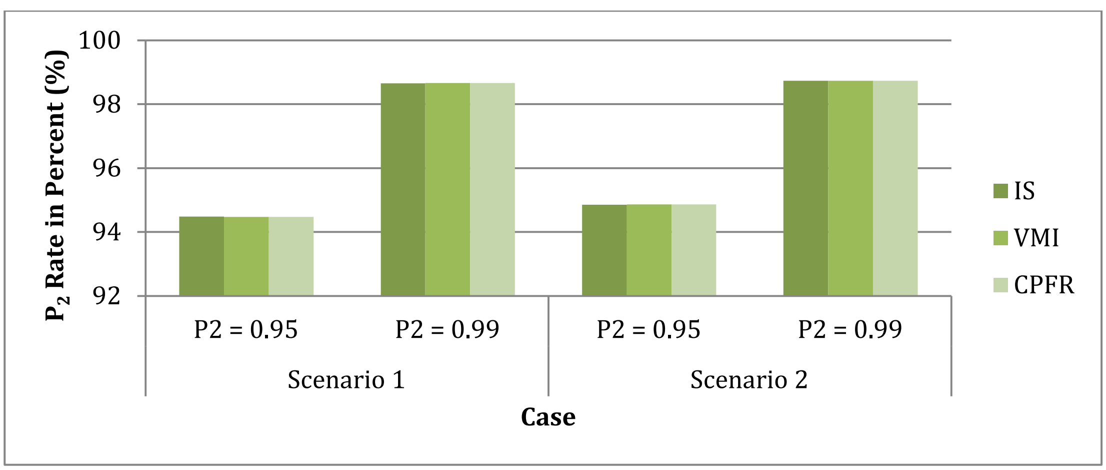

Figure B-2: P_2 Service Rates Achieved by the Retailer in the Simulations with Multiple Promotions - Empirical Data

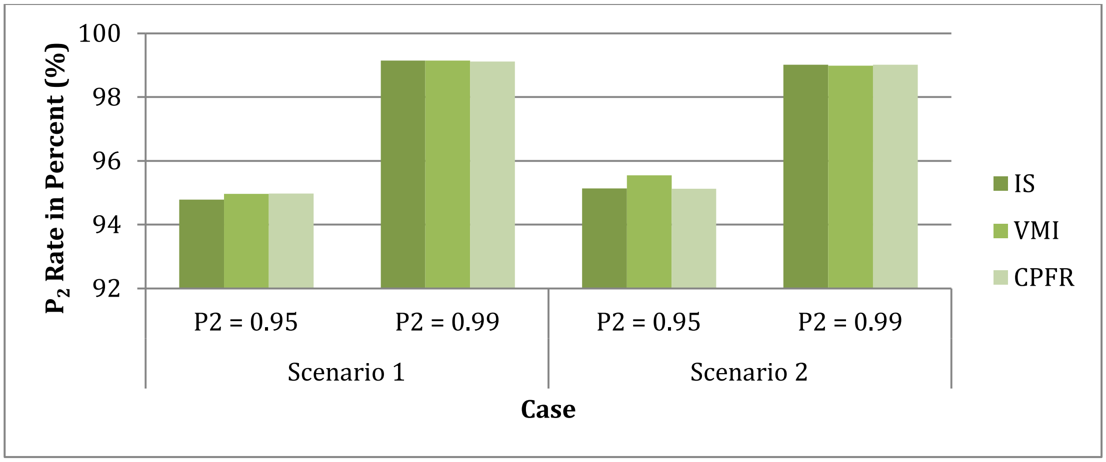

Figure B-3: P_2 Service Rates Achieved by the Retailer in the Simulations with One Large, Seasonal Promotion - Empirical Data

Additional Results for Section 4.2.1 Simulations without Promotions - Empirical Data Scenario 1: First Aid Product

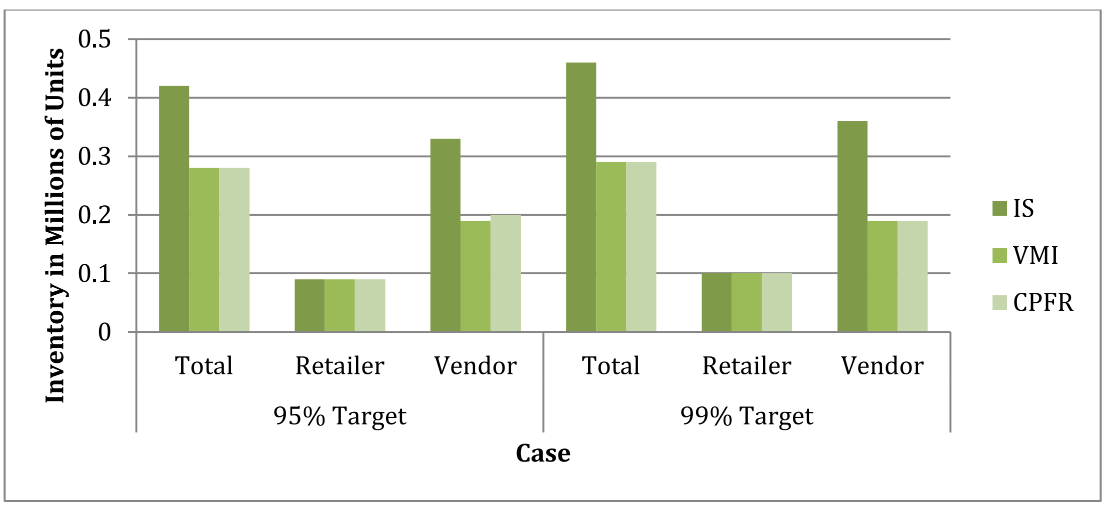

Figure B-4: Average Levels of Inventory for Simulations without Promotions - Scenario 1: First Aid Product

Table B-1: Confidence Intervals for the Difference in Inventory (Thousands) in Simulations without Promotions - Scenario 1: First Aid Product

Target Fill Rate (%)	Level	Confidence Intervals: Inventory Differences		
		IS - VMI	IS - CPFR	VMI - CPFR
95	Total	(57.21, 192.72)	(58.55, 190.82)	(-6.49, 4.55)
	Retailer	(-1.87, 2.25)	(-1.21, 1.17)	(-2.19, 1.93)
	Vendor	(57.95, 192.39)	(58.76, 190.64)	(-5.41, 3.72)
99	Total	(96.93, 209.14)	(100.25, 207.85)	(-7.00, 5.15)
	Retailer	(-2.68, 3.12)	(-1.98, 1.88)	(-3.13, 2.74)
	Vendor	(98.51, 208.81)	(99.03, 207.48)	(-5.49, 4.32)

Additional Results for Section 4.2.3 Simulations with Multiple Promotions - Empirical Data

Scenario 1: Stationary Product

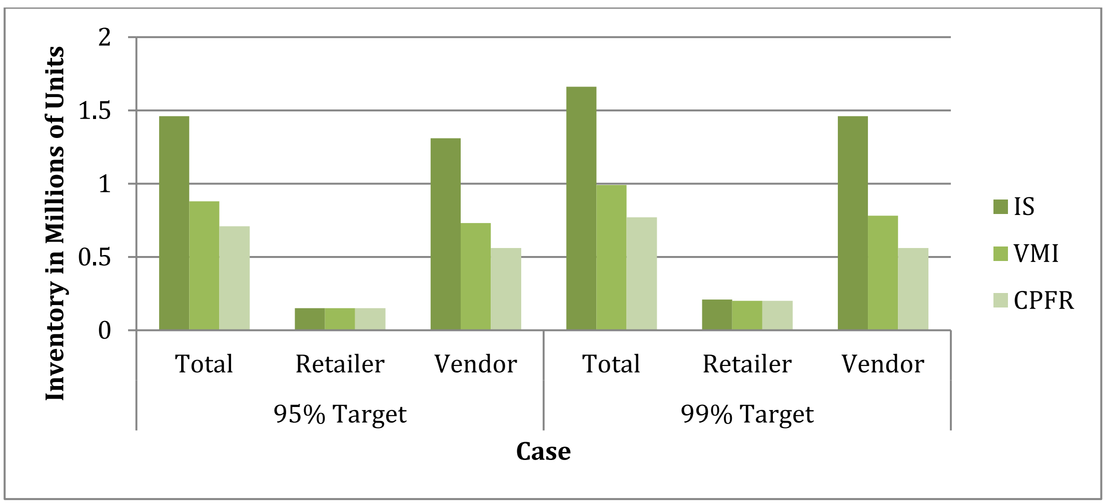

Figure B-5: Average Levels of Inventory for Simulations with Multiple Promotions - Scenario 1: Stationary Product

Table B-2: Confidence Intervals for the Difference in Inventory ($ Thousands) in Simulations with Multiple Promotions - Scenario 1: Stationary Product

Target Fill Rate (%)	Level	Supply Chain Strategy		
		IS - VMI	IS - CPFR	VMI - CPFR
95	Total	(0.53, 0.62)	(0.71, 0.79)	(0.16, 0.19)
	Retailer	(-0.003, 0.003)	(-0.002, 0.002)	(-0.003, 0.003)
	Vendor	(0.53, 0.62)	(0.71, 0.79)	(0.16, 0.19)
99	Total	(0.62, 0.72)	(0.84, 0.94)	(0.20, 0.24)
	Retailer	(-0.006, 0.006)	(-0.004, 0.004)	(-0.006, 0.006)
	Vendor	(0.62, 0.72)	(0.85, 0.94)	(0.21, 0.24)

Additional Results for Section 4.2.4 Simulations with Multiple Promotions - Empirical Data

Scenario 2: Hardware Product

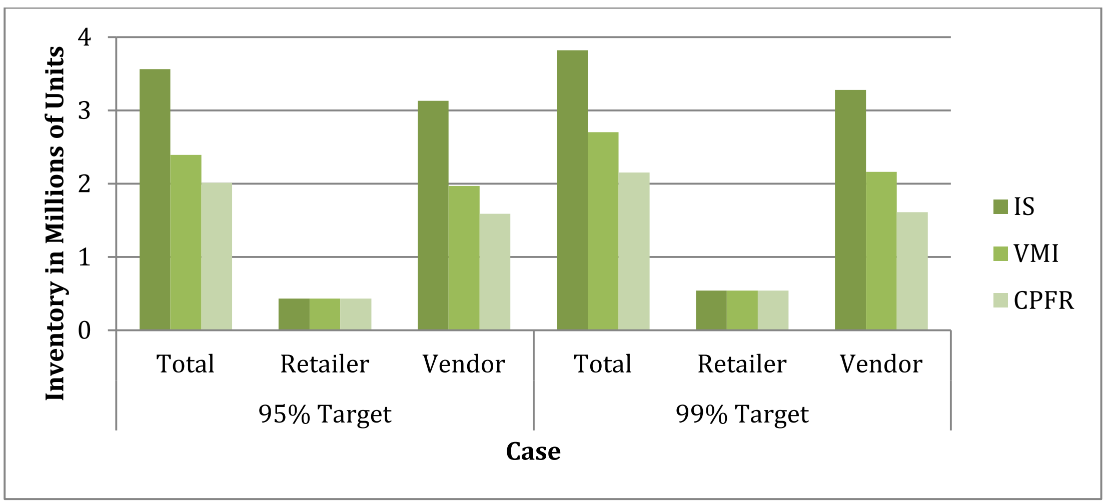

Figure B-6: Average Levels of Inventory for Simulations with Multiple Promotions - Scenario 2: Hardware Product