Reducing the Demand Forecast Error Due to the Bullwhip Effect in the Computer Processor Industry

By
Emily Smith

Bachelor of Science in Mechanical Engineering,
Massachusetts Institute of Technology 2005

Submitted to the MIT Sloan School of Management and the Mechanical Engineering Department in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration
AND
Master of Science in Mechanical Engineering

In conjunction with the Leaders for Global Operations at the
Massachusetts Institute of Technology
June 2010

Signature of Author________________________
MIT Sloan School of Management
May 7, 2010

Certified by________________________
Charlie Fine, Thesis Supervisor
Chrysler Leaders for Manufacturing Professor of Management

Certified by________________________
Abbott Weiss, Thesis Supervisor
Senior Lecturer, Mechanical Engineering

Certified by________________________
Henry Marcus, Thesis Supervisor
Professor of Mechanical Engineering, Marine Systems

Accepted by________________________
David Hardt
Chairman, Committee on Graduate Students, Dept. of Mechanical Engineering

Accepted by________________________
Debbie Berechman
Executive Director of MBA Program, MIT Sloan School of Management

This page has been intentionally left blank

Reducing the Demand Forecast Error Due to the Bullwhip Effect in the Computer Processor Industry

By
Emily Smith

Submitted to the MIT Sloan School of Management and the
Mechanical Engineering Department on May 7, 2010 in Partial Fulfillment of the
Requirements for the Degrees of Master of Business Administration and
Master of Science in Mechanical Engineering

ABSTRACT

Intel's current demand-forecasting processes rely on customers' demand forecasts. Customers do not revise demand forecasts as demand decreases until the last minute. Intel's current demand models provide little guidance for judging customer orders when the market changes. As a result, during the economic downturn of Q3 and Q4 '08, Intel's model could not predict how much billings would decrease. The demand forecast had large amounts of error caused by the bullwhip effect (order amplification in a supply chain).

This project creates a new demand forecast model in two phases. The first phase investigated the supply chain of OEMs and Retailers. The second phase of the project used the supply chain information discovered in phase one to create a new demand forecast that reduces the error caused by the bullwhip effect.

The first phase determined that the average time it takes a CPU to go from Intel to end customer purchase is seventeen weeks. The first phase also indentified ownership of products throughout the supply chain and parties making purchase decisions. The supply chain information was then used in the second phase of the project to create a demand forecast model. The new model is a heuristic model that simulates quarterly purchase decisions of retailers and OEMs including lead times and inventory.

The resulting model allows Intel to monitor and react to consumption changes faster than waiting for customers to change their demand forecasts. The model also provides a better forecast during times of change. The model reduces the error due to the bullwhip effect and indentifies early when a downturn or upturn is going to happen in ordering behavior.

Thesis Supervisor: Charlie Fine
Title: Chrysler Leaders for Manufacturing Professor of Management

Thesis Supervisor: Abbott Weiss
Title: Senior Lecturer, Mechanical Engineering

Thesis Supervisor: Henry Marcus
Title: Professor of Mechanical Engineering, Marine Systems

This page has been intentionally left blank

Acknowledgments

I would like to thank the Leaders for Global Operations and Intel for this amazing experience. Thank you to Adam King for supporting an LGO thesis in Asia. My colleagues in Taiwan and Hong Kong were great support throughout this process. They helped me to find information about Intel and this project as well as to help me feel at home in Asia. My appreciation to the demand forecast team: Craig, Danny, Gidget, John, Pramod, Raphael, and Tom for helping me with the models and with gathering data. I also wanted to acknowledge H.C., Kitty, and Lancy for their instrumental help in collecting information about the computer supply chains. Special thanks to Joyce for showing me around Taiwan and treating me like family. My gratitude to Grace, Mavis, Pauline, Peggy, Sabrina, and Tony for helping make Taiwan feel like home. Thank you to Gidget for showing me where to shop in Hong Kong.

I want to thank my thesis advisors Charles Fine and Abbott Weiss for their mentorship during my internship. They helped me come up with creative solutions to mapping the supply chain.

Thank you to my parents Bev and Glenn Smith for the support. Not only have they been supportive during my time at school they also are always willing to help no matter how crazy the task. Thank you for all the pictures from retailers. Thank you to my LGO and Sloan classmates for helping me collect serial numbers to verify the supply chain length.

This page has been intentionally left blank

Table of Contents

This page has been intentionally left blank

Table of Figures

Table of Tables

Glossary

ADI: Avaiable Die Inventory. Raw die invenotry buffer before going into assembly............... 26

APAC: Asia-Pacific Region 61

ATM: Assemble-Test-Manufacturing Site. Facility where raw wafers are assembled into CPUs. 26

Bullwhip Effect: the fluctuation of ordering behavior in a supply chain due to the failure to account for the supply line length. As well as, the amplification of the flucations of orders as the change in ordering goes up the supply chain. 19

CCP: Customer commit process. Intel's allocation based sales process. 33

CW: component warehouse. Facility where finished CPUs and other products are kept. 26

Die: intermediate in the process of turning a silicion wafer into a CPU. 26

End Consumption: end consumer purchase of laptops. 37

EOH: End on Hand. Count of an item at the end of a quarter. 37

Ex-factory: the amount of product leaving the factory. 38

FSM: Fabricate, Sort and Manufacture. Intel facility where silicon wafers are turned into dies.26

GEO: Intels geographical division of sales areas. 34

Hub: company that warehouses and hold inventory. 29

Inventory Turn Time: time for an item of inventory to go from the bottom of the inventory to the top. 30

JD: Judged Demand. Finalized demand that is changed based on external knowledge before shipping to another part of the organization. 27

MMBP: Microprocessor Marketing and Business Planning. Intel Sales and Operations Group. 34

Throughput Time: time to go through a process from start to finish including the time to go through the inventory stage. 30

This page has been intentionally left blank

1. Introduction

1.1. Problem Introduction

In Q4 2008 Intel's sales were down 23.5% from Q4 2007 (Rogoway, 2009). Sales were down 20% compared to Q3, marking the second time in 20 years such a downturn had occurred (Deffree, 2009). Paul Otellini, Intel's president and CEO, said "The pace of the revenue decline in the quarter [Q4 2008] was dramatic and resulted from reduced demand and inventory contraction across the supply chain." (Rogoway, 2009) Intel was caught off guard by the sudden decline in orders and had no idea if and when orders were going to recover to their previous levels. Although Intel knew an economic recession had begun earlier that year, the drop in Intel's orders was greater than retail decline and occurred quarters later.

Two major lessons came out of the results of Q4 2008: Intel needed a forecast method that could better handle predicting demand during times of great changes, and Intel needed to know about its customers' supply chains. As noted by Paul Otellini, the supply chain had contracted, but Intel was not completely sure where it had contracted and by how much (Rogoway, 2009). This paper explores the customer supply chain and creates a model that uses the supply chain information to forecast demand better during dramatic changes.

1.2. Company Background

Intel® was founded in 1968 by Gordon E. Moore and Robert Noyce, both former employees of Fairchild Semiconductor. The company was initially named NM Electronics, after the founders, and was renamed Intel®, or "INTegrated ELectronics" later that year. Since its beginning, Intel has grown to nearly 84,000 employees and 300 facilities worldwide as of

December 2008. Intel is number 61 in the 2009 Fortune 500 rankings (Fortune 500, 2009). Its 2008 revenues were \$37.5 billion (Intel, 2008). Intel, the majority manufacturer of CPUs, receives over 80% of the total CPU market revenue and 87% of the mobile CPU market revenue (Shah, 2010). In 2008 Intel had \$11.4 billion in mobile CPU revenue (Intel, 2008).

1.3. Thesis Structure

Chapter 1 describes the problem Intel experienced that resulting in this study; this chapter also gives background on Intel.

Chapter 2 is a literature review that defines the bullwhip effect; provides evidence and consequences of the bullwhip effect in the personal computer industry; and describes methods for reducing the error in demand forecasting caused by the bullwhip effect.

Chapter 3 reviews the current supply chain from Intel's creation of a CPU to end consumer purchase. The order decisions points, inventory ownership, and time lags are also noted throughout the process.

Chapter 4 describes the current demand forecast process, as well as develops and analyzes two types of alternative demand forecast models. The implementation of a new demand forecast model is also described.

Chapter 5 reviews others actions that can reduce the error caused by the bullwhip effect in the demand forecast: vendor managed inventory, sales induced seasonality, and point of sale information sharing.

Chapter 6 summarizes the results of this project and next steps Intel can take to improve demand forecasting

2. Literature Review

2.1. Bullwhip Effect Definition

It has been shown that companies operating in serial supply chains can fail to account for feedback and actions that have not yet had a result (Sterman J. D., 1989). The bullwhip effect is the failure to account for inventory in the supply line and time it takes to receive an order. For example when there is an increase in demand, a party will order more and continue to order more until the demand is met. Once demand is met and the orders in the supply line keep coming in at the old level, the party cuts off or over decreases orders to compensate for the extra inventory arriving. This wave pattern of over and under correction of orders propagates up the supply chain. The further upstream the supply chain, the greater the amplification of the wave pattern becomes (Sterman J. D., 1989). Figure 1 shows an example of the bullwhip effect; the order variation amplifies as a change in demand by the consumer propagates through the supply chain.

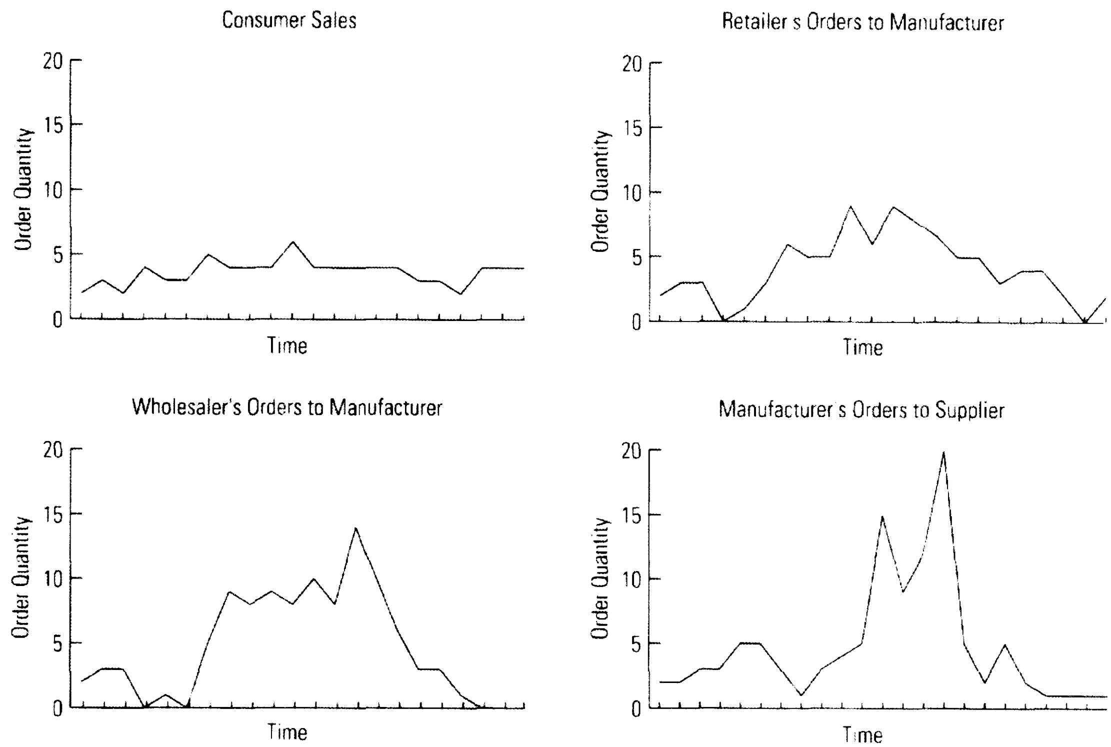

Figure 1: Example of the Increase in Order Variation Moving Up the Supply Chain (Lee, Padmanabhan, & Whang, 1997)

The amplification of order variance, the bullwhip effect, was first identified by Forester (1961). Sterman (1989) further characterized the effect as being caused by the difficulty of management to evaluate the complex feedback loops and time delays in supply chains. The bullwhip effect can result in costly over and under production. The automotive supply chain is a good example of the bullwhip effect; the GDP from 1961 to 1991 ranged from +2% to -3%, while automotive production swung from +/- 20% and machine tools fluctuated +/- 80% (Fine, 1998). The general magnitude of impact of the bullwhip effect is not known because it is highly dependent on the situation including: retailers' ordering patterns, channel structure, and inventory policies of supply chain members (Sahin & Robinson, 2002). The next section describes the channel structure and results of the bullwhip effect in the processor industry.

2.2. Bullwhip Effect in the Computer Industry

The channel structure of the computer industry has gone through many changes. When the personal computer industry began in 1985, computers were built by vertically integrated manufacturers. With IBM's standardization of protocols and connections, the supply chain shifted from being vertically integrated to having multiple suppliers. IBM and other manufacturers became assemblers (Dedrick & Kraemer, 2007). The change in the supply chain increased the number of parties and time involved in the manufacture of a computer. In 1995 Dell brought another shift to the market by changing the fulfillment supply chain from indirect through resellers to direct to customers (Fine, 1998). This shift reduced the amount of time it took to get a computer to a consumer. Around 2004 another shift occurred in the computer supply chain as computer manufacturers, OEMs, shifted the assembly process to contract manufacturers, ODMs (Gupte, 2009). The majority of computers are now manufactured in Asia. The move of manufacturing from the US to Asia increased time from final assembly to an end consumer's purchase. This change marked the return to outsourcing from vertical integration. The introduction of ODMs and assembly in Asia has decreased the manufacturing cost at the expense of increasing the amount of time it takes to complete a customer order. The new process has also added another party into the supply chain coordination. The full effects of adding another party to the supply chain coordination had not been evident until the recent economic downturn.

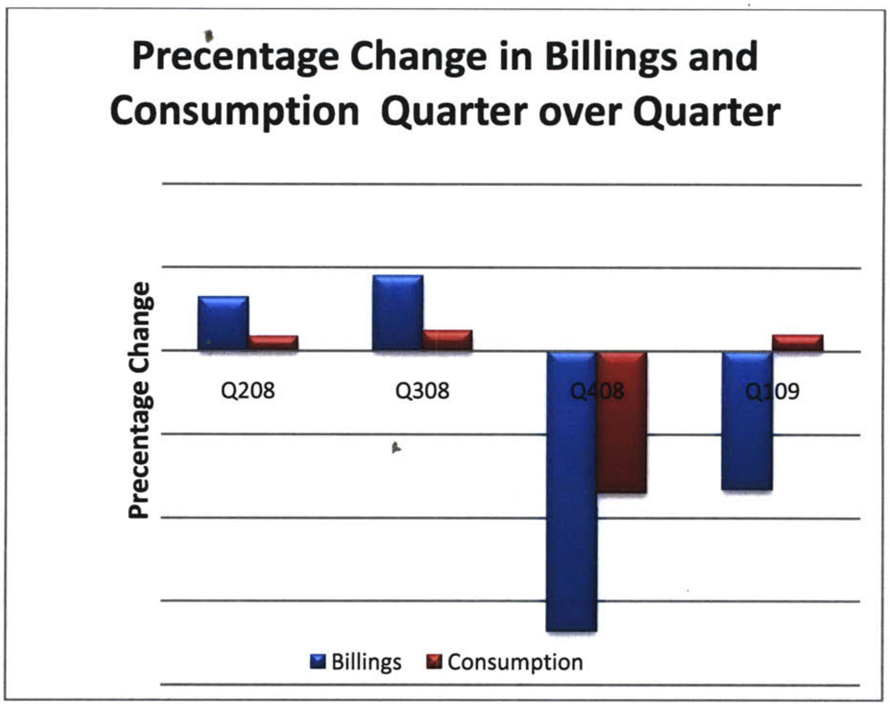

Figure 2: Bullwhip Effect in Intel's Orders during the Economic Downturn (International Data Corporation, 2009)

The recent economic downturn revealed the effect adding ODMs had on the bullwhip effect in the computer industry. The economic downturn is estimated to have begun around December of 2007 (National Bureau of Economic Reasearch, 2008). During the 2007-2009 economic downturn purchases of products containing chips fell 8%, product shipments fell 10%, and chip shipments fell 20% (Dvorak, 2010). Intel saw a much greater decrease in orders than the number of retail orders during Q3 and Q4 2008, Figure 2. The economic downturn changed the predictability of the market, and Intel needed to review its demand forecast techniques.

The next section reviews the causes and solutions of the bullwhip effect. These solutions will later be explored in the models in Chapter 4.

2.3. Bullwhip Effect Causes and Solutions

The bullwhip effect has five major causes: demand signal processing, lead times, order batching, price fluctuation, and shortage gaming (Lee, Padmanabhan, & Whang, 1997). The two main causes of the bullwhip effect have been shown to be lead times and demand signal processing (Disney & Towill, 2003). This project focuses on changing how Intel does its demand signal processing to reduce forecast error caused by bullwhip.

Many methods have been studied on how to reduce the bullwhip through demand signal processing. There are two main groups of study: one group explores the suppliers' demand forecast process, regardless of information sharing, and the other group explores modeling the decision processes within the supply chain.

The first area looks into supplier demand forecast techniques that can reduce demand signal fluctuations. Moving average and smoothing techniques of demand signals has been shown to reduce the bullwhip effect without needing information sharing (Graves, 1999). More sophisticated techniques of forecasting, including Holt's method and Brown's double-exponential smoothing, have also been shown to reduce the bullwhip effect even further (Wright & Xin, 2008). These methods require weekly or daily information for the entire supply chain to evaluate the results. Intel does not have this granularity of data to explore these techniques currently. A second method of signal processing was explored that could be done with the data Intel had available.

The second method of reducing the bullwhip effect is modeling the decision-making process throughout the supply chain. The majority of authors study an order-up-to policy. Chen (2000) defines an order-up-to policy as ordering demand plus a desired service level quantity without regard to lead times. Chen studied a simple scenario with two parties and found that this model can emulate the decision making of parties in the supply chain. Lee (2000) took Chen's model one step further by adding more parties to the length of the supply chain. Lee also compared the model results with and without information sharing. None of these models explored a case with multiple retailers and with Vendor Managed Inventory, VMI, like the model in this paper. The capability of these models in this scenario is explored in Chapter 4.

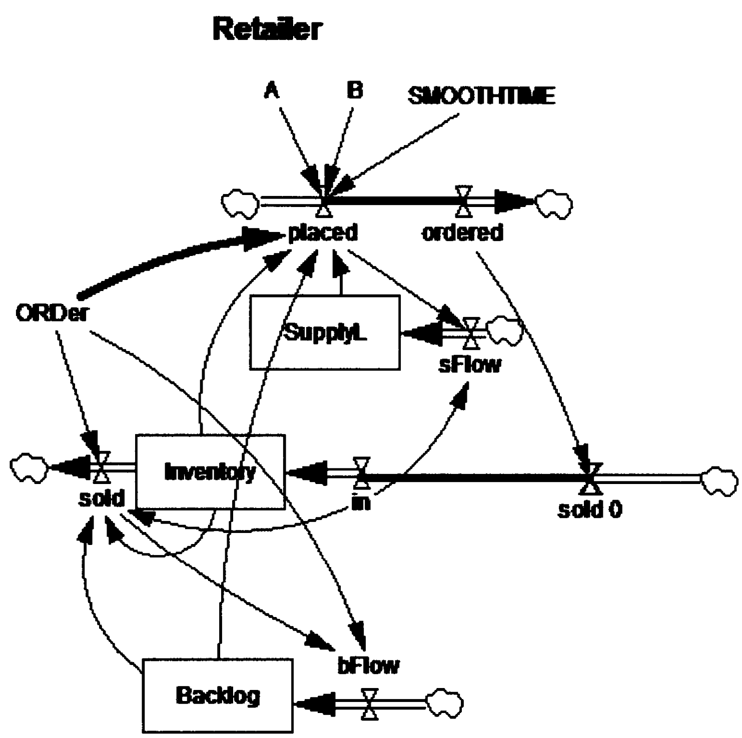

Figure 3: Heuristic Stock Flow Model

The next progression of decision making models includes lead times. Sterman (1989) creates a simple heuristic model. The model gives weights to the different information available to the supply chain party: time lag (supply line that has not arrived yet), optimal inventory level, and demand, Figure 3. In Sterman's model has variables A and B. The variable A represents the weight of inventory on hand in the new order process and B represents the weight of inventory in the pipeline in an order decision. For example, a supply chain party that fully recognizes the entire inventory in the supply line gives this B a weight of one, alternatively a party that does not take supply inventory into account gives this variable a weight of 0. This simple heuristic model was shown to be able to describe complex decision processes occurring within the supply chain.

2.4. Summary

This paper focuses on reducing forecast error caused by the bullwhip using supply chain decision making modeling. The first is an order-up-to method that focuses on maintaining the desired inventory level without time lags similar to Chen (2000). The second method studied focuses on customer ordering policies similar to Sterman (1989). The models in this paper add to the past research by modeling a situation with multiple retailers and a supply chain that has both VMI and non-VMI pathways. The model only accomplishes this for one industry and does so during an economic downturn. Further research in other industries over different time periods are needed to add to the results of this paper.

The next chapter begins the research into the time lags in the supply chain. These data are necessary to complete a model that analyzes the bullwhip effect in the Intel supply chain.

3. Computer Chip Supply Chain Background

3.1. Overview

The first step to improving the demand forecast process is to determine the supply chain from Intel to the end customer and the information flow within the supply chain. An accurate map of the supply chain is critical to determining how to model the bullwhip effect and what sort of demand forecast process to use. The following chapter describes the Intel production process, the customer supply chain, and each party's order process.

3.2. The Intel Supply Chain

The computer chip manufacturing process starts with a wafer. The first stop is a Fabrication facility or FSM. An FSM is where they Fabricate, Sort and Manufacture dies. A die is an intermediate step between a blank wafer and a computer chip. Dies have gone through lithography. The good dies are sorted and then shipped to an inventory location called available die inventory, ADI. Dies are then pulled out of the ADI to be turned into CPUs at the Assembly-Test-Manufacturing facility, ATM. Before a CPU is finally finished, there is another buffer called Semi Finished Goods Inventory, SFGI. Orders are finally pulled from SFGI and finished by burning in some of its features like speed and power consumption and shipped to a component warehouse, CW, for order fulfillment.

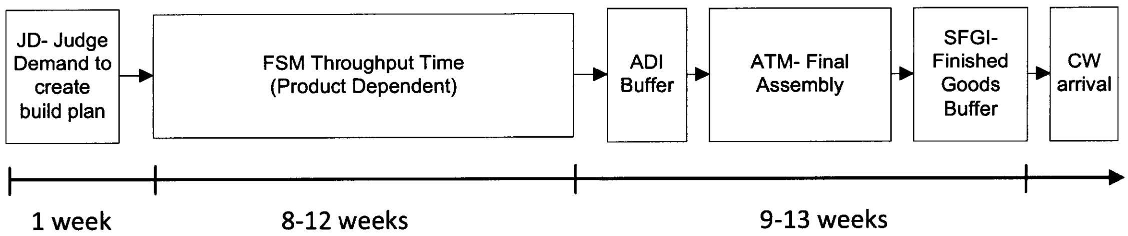

Figure 4: Timeline for the Production Process (Pai, 2009)

The entire manufacturing process from determining the build plan from judged demand, JD, to receiving finished CPUs at component warehouse takes 3-4 months, Figure 4. Expediting processes can reduce the amount of time to get CPUs manufactured and to a customer but usually production needs to start a 17-25 weeks ahead of order fulfillment.

The flow of information that kicks-off production processes is outlined in Pai (2009). The judged demand forecast is used to determine how many wafers to start at the beginning of the process. To ensure on time delivery of customer orders, production planning must work to a demand forecast that is at least one quarter ahead. The pull of dies and semi-finished goods is changed as the demand forecast is refined on a weekly basis.

A lot of work has already been done to improve and continue to improve the internal supply chain's responsiveness (Pai, 2009). To create an accurate supply chain model it is necessary to know when and how production uses the demand forecast and the time lags in production. The Intel manufacturing process information discussed in this section is used in Chapter 4. In the next section, we are going to continue following the flow of a CPU to an end customer.

3.3. The Customer Supply Chain

Once a CPU reaches the component warehouse, it is an average of another 14 weeks before a CPU reaches an end consumer as part of a mobile computer. A CPU goes through four major steps in the process of reaching an end consumer: component and manufacturing inventory, manufacture, shipping and retail inventory as shown in Figure 5.

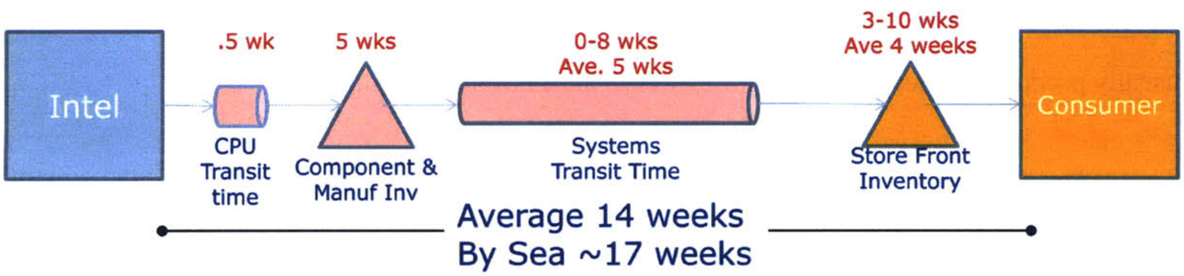

Figure 5: Overview of OEM Mobile Computer Supply Chain

The component and manufacturing inventory process includes either VMI or third party warehouse inventory and third party manufacturing. The third party manufacturing then ships the computer to a third party distributor in the receiving continent. The third party distributor takes care of the last mile distribution to retailers. Retailers then take care of moving inventory to the store front for purchase by consumers.

The majority of time is spent in shipping and in retail inventory. A CPU can be bought as part of an end consumer purchase in as little as eight weeks if retailers use air shipping and push items on the shelf as quickly as possible. On the opposite end of the spectrum, it can take 22 weeks before a CPU is purchased if a retailer chooses sea shipment and has a maximum of inventory of an item. The rest of this section details the pathways the CPUs take to reach a customer as part of a computer, the players involved, the ownership at different points, and decision making points.

3.3.1. Component Inventory

CPUs can take two major pathways from Intel to a manufacturing site. Some OEMs use Intel's VMI hub service and some utilize a third party hub company. In both situations it only takes a day to get CPUs to the distribution site and a day to get the CPUs to the manufacture site making the total shipping time half a week, Figure 5.

At Intel's VMI, customers submit weekly forecasts and Intel holds the weekly forecast of CPUs for the customer. At the end of the week Intel can release those CPUs to another customer. If the customer wants more CPUs there is no guarantee they can have them. In return for the VMI hub usage, customers are expected to share their total inventory information in their factory on a weekly basis. The result of this ordering policy is that OEMs using Intel's VMI services hold more inventory at the third party manufacturing sites, ODMs, versus those using a third party hub.

OEMs that do not utilize Intel's VMI services use a third party hub to hold their raw inventory. At a third party hub the OEM owns the inventory rather than Intel or the hub. The purpose of the hub is to hold the raw inventory in a central location before sending to the multiple locations that manufacture the mobile computers. A third party hub allows the OEM to postpone the manufacturing location decision and keep inventory low at ODMs.

The total amount of inventory at a distribution site and an ODM is equivalent regardless of the pathway chosen. To maintain enough buffer inventory for the shifts in demand, the total number of weeks of inventory, WOI, at the distributor and the ODM is the same. Therefore the choice to use a VMI hub or a third party hub does not change the supply chain model.

3.3.2. Manufacture Process and Inventory

The combined throughput time from Intel to shipment averages five weeks, Figure 5. The throughput time for a mobile computer is a week. A CPU goes through several buffers and assembly process to become a mobile computer, Figure 6.

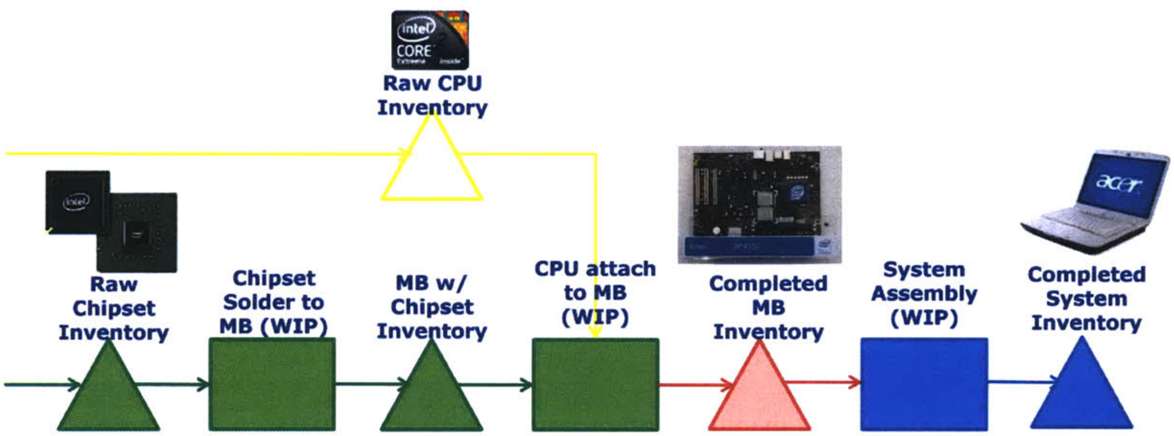

Figure 6: Manufacturer Process Flow

Most of the five weeks it takes to go through the component inventory and manufacture process mentioned in Section 3.2 is inventory turn time. CPUs are also kept in inventory attached to a motherboard. A completed motherboard has a chipset and CPU attached. The final configuration of the computer is not done until a completed motherboard and other components are assembled into a mobile computer. ODMs do not keep finished system inventory because of all the variations and the expense of completed systems versus components. Inventory is kept at semi-finished state until an order is placed. ODM's receive weekly demand forecasts several quarters out but actual orders are not finalized until 48-72 hours prior to assembly. To be able to build to order, BTO, but maintain a quick turnaround, ODMs keep completed motherboard inventory. Once the order is built, the final step for the ODM is to ship the order to the appropriate receiving country.

3.3.3. Shipping

Completed systems are shipped directly from ODMs to third party distributors in the receiving continent, Figure 7.

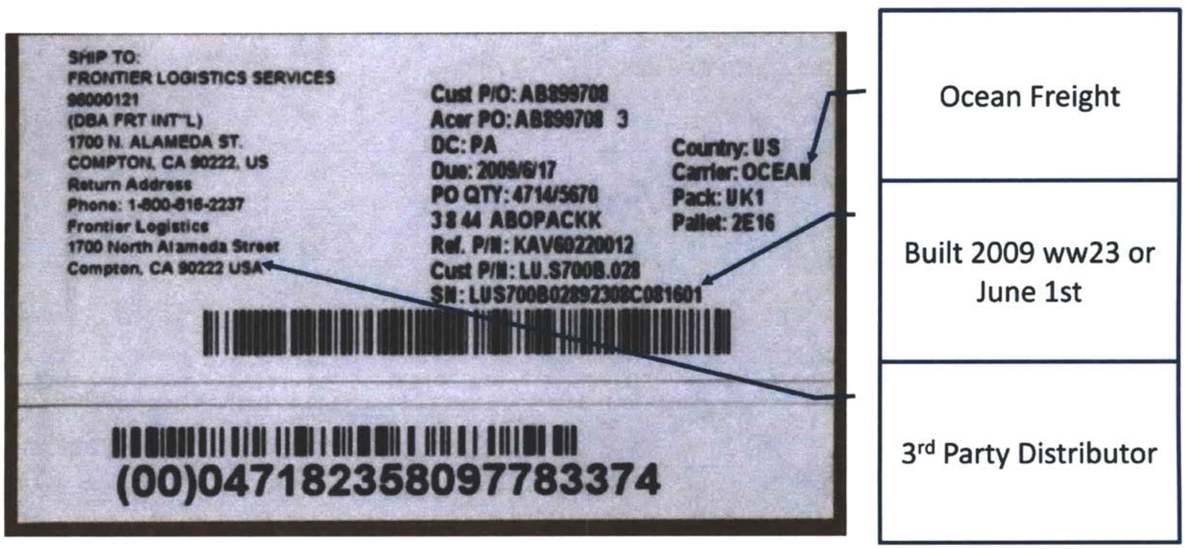

Figure 7: Shipping Label on Retail Mobile Computer Box at Wal-Mart

OEMs employ third parties to take care of the last mile of distribution from customs to the local retail hub. The distributor or logistics service breaks apart shipments into smaller batches for each local retail hub. Figure 7 shows the shipping label of an Acer mobile computer. The label shows that Acer employed Frontier Logistics to ship the computer from customs in California to a local Wal-Mart hub. The label also shows the shipment type and date of manufacture. These labels were used to verify the total throughput time found through many different interviews and research processes.

3.3.4. Retail Inventory

Upon arrival at the retailer's local hub, the retailers own the computers. The retailer then takes responsibility for sorting the computers again into even smaller batches that go to retail stores. At retail stores computers first go to the back room; then they can move up to the store front shelves. The amount of time spent at each of these areas depends on the amount of floor space a product is given and the speed of sales.

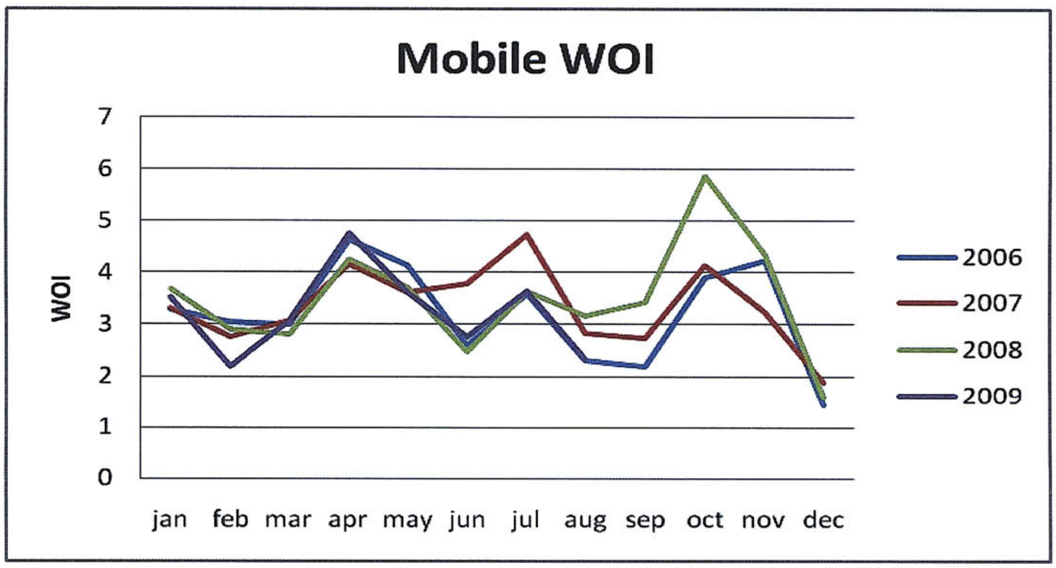

Figure 8: WOI in US Electronics Retailers by Month for the Past 4 Years (NPD Group, 2009)

The total time spent in inventory in all the retailers inventory locations can vary from 4-10 weeks. External market forces cause some variation in the WOI at a retailer but the majority of the variation in WOI during the year comes from the effects of seasonality, Figure 8. To prepare for seasonal demand spikes, retailers stockpile inventory.

The timing of this fluctuation and the amount varies for different regions spending habits. For example, in the US retailers start stockpiling inventory in September for increased sales in November and December. In Korea the big sales period is back to school in September.

3.3.5. Customer Supply Chain Summary

There are many parties involved in getting a CPU assembled into a computer and delivered to an end customer. Although five parties touch a computer in the supply chain, only two parties are making order decisions: OEMs and retailers. OEMs employee two third-party distributors: raw inventory distributors and final product distributors. The distributors are not making decisions on how much inventory to store; they are receiving materials based on the OEM orders and shipping materials based on OEM requests. Third parties employed by the OEM are not part of the decision making models.

3.4. The Sales Process

The demand for Intel CPU can often be greater than supply. For this reason Intel has created a sales process based on allocating the supply of chips. Intel's sales method is called the customer commit process, CCP. The CCP is used to ensure proper allocation of CPUs over different sales regions and companies. CCP also prevents Intel from selling more chips than available.

Figure 9: Intel Sales Process (Chow, 2004)

CCP starts with customers sending their forecasts for the upcoming quarter to the sales teams. The sales teams pass this data to demand forecast teams that aggregate each company's data and judge the demand to create a regional GEO forecast, Figure 9. The regional forecast is then sent to MMBP, Intel's global demand forecast team, and aggregated into a global forecast. Based on this forecast MMBP sends back to each GEO an allocation of CPU's. Customers can then book up to the allowed amount. Orders are cancelable up until 2 weeks before delivery. The order, change, and cancel process are further discussed in Chow (2004).

3.5. Summary

The total throughput time from a wafer to an end consumer is about 6 months; 3 months for CPU manufacturing and 3 months for mobile computer assembly, shipping and inventory. Anywhere from 5 to 6 parties are involved in the whole process, with 3 companies in charge of the order decisions.

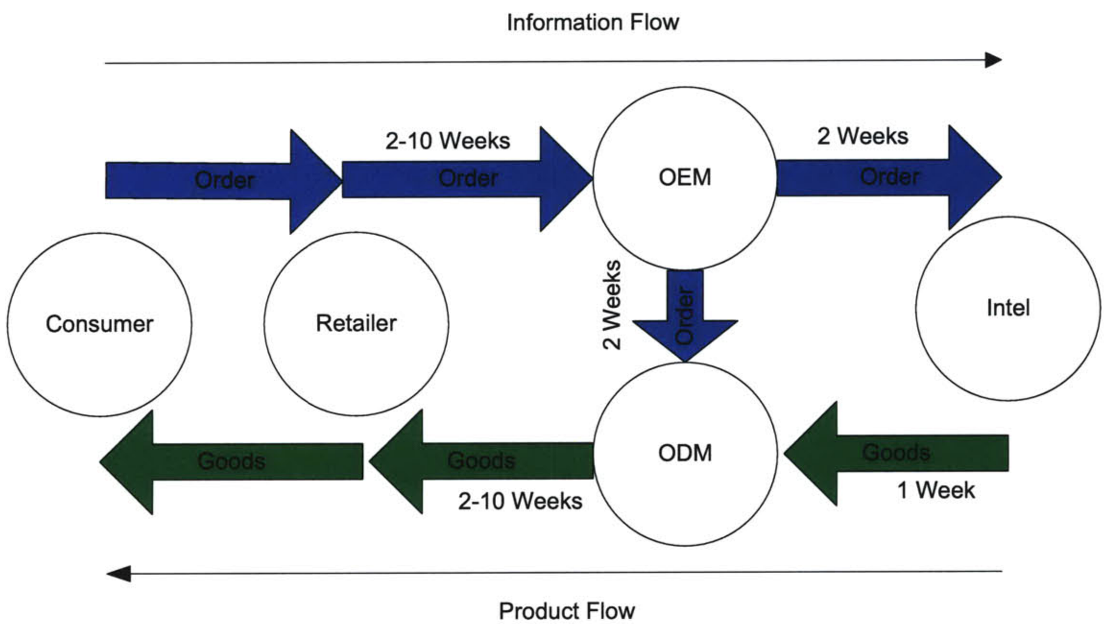

Figure 10: Flow of Goods and Orders in the Mobile Supply Chain

Working backward through the supply chain shows the information flows and how far ahead of time each party must place an order, Figure 10. It takes 2 to 10 weeks to get a mobile computer from an ODM to a retailer. This means retailers have to predict what demand will be in 10 weeks if they do not want to air ship computers. OEMs have to finalize CPU orders 2 weeks prior to delivery and computers take one week to manufacture. When OEMs order computers they are predicting retail demand 3 weeks ahead of orders and 13 weeks ahead of end customer demand. Because Intel needs 3 months to turn a wafer into a CPU; this means Intel is building wafers 3 months ahead of OEM demand and 6 months ahead of end customer demand.

The effect of the supply chain length and the multiple parties involved in the ordering decisions will be discussed in the next chapter on demand forecasting. Chapter 4 goes over how to use the information discussed in Chapter 3 to mitigate the bullwhip effect in the demand forecast error.

4. Demand Forecast Models

4.1. Overview

The current demand process relies upon OEMs demand forecasts. Chapter 4 looks at forecasting methods that reduce the reliance on OEMs forecasts and reduce the effects of order amplification.

The current demand forecast process is explored first. Then two types of models are created and evaluated on how they improve the current demand forecast process and reduce the error cause by the bullwhip effect. The first model is an order-up-to model like Chen (2000). An order-up-to model works with the current model in use by Intel. The second model is a heuristic model like Sterman (1989). Three versions of the heuristic model are explored as the variables used by the parties in the supply chain are revised to improve the similarity of the model to the actual decision making process.

4.2. Current Demand Forecast Process

The current demand forecast process combines a top-down, consumer demand to billings forecast, and bottoms-up, expected billings to consumer demand, approach to triangulate upon a billing's forecast. The process involves inputting the OEM billings forecast number and a consumption estimate, then reviewing the market assumptions.

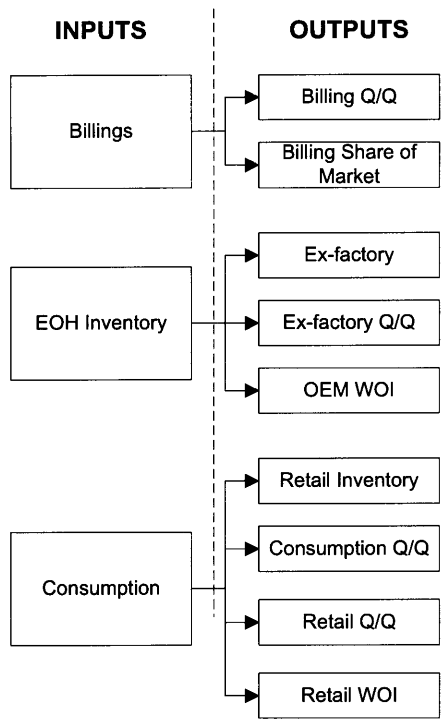

Figure 11: Inputs and Outputs of Current Demand Forecast Model

The current demand forecast model has three main inputs, Figure 11. Intel receives billings numbers from the customer, receives or calculates ending on hand, EOH, inventory from OEMs and takes an educated guess at the end consumption, consumer sales. From these numbers, Intel determines multiple outputs to decide if the OEM demand forecast is accurate.

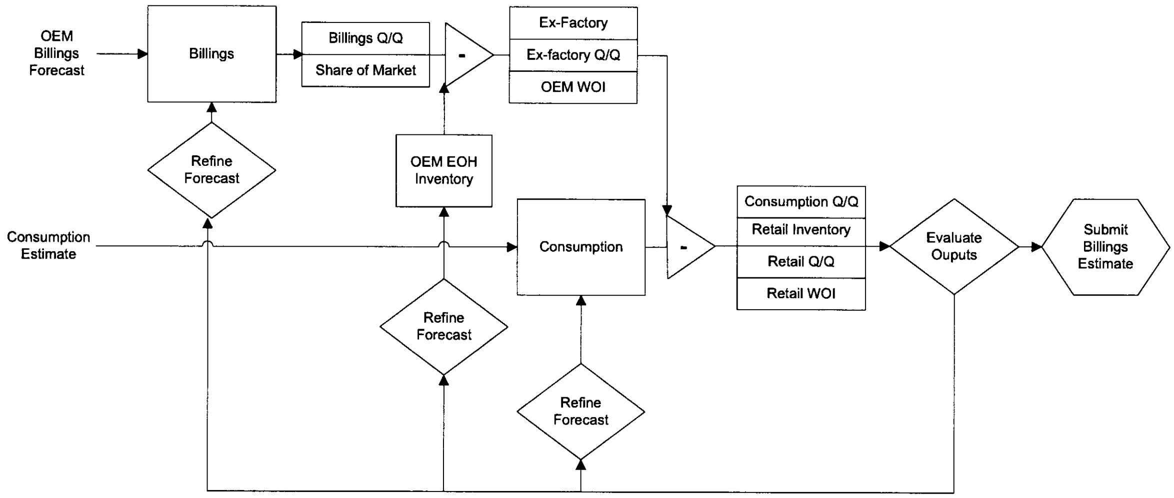

Figure 12: Demand Forecast Process Flow

The bottoms-up process begins by inputting the OEM's billing forecast into the market share model, Figure 12. The billings forecast is compared to prior quarter's billings numbers through reviewing the Q/Q growth and share of market. The billings number is then translated into an OEM factory output number (ex-factory) using the projected end on hand inventory. The factory output is the increase in inventory minus the billings. The ex-factory number is reviewed against prior quarters to ensure that billings and inventory numbers match to the OEM's inventory goals.

The top-down process begins with an end user consumption estimate. The estimate of end user consumption is combined with the ex-factory information to determine the retail inventory. Similar to how EOH inventory is determined, Intel calculates retail inventory as the difference between ex-factory and consumption; what was bought minus what was sold. The retail inventory number is checked for a building of retail inventory by comparing the WOI and Q/Q change in inventory. These numbers provide a check on the bottoms-up approach to determine if the factory numbers make sense for the retail market.

The output (retail consumption, factory shipments, and billings) numbers are reviewed for alignment between themselves and against what is going on within the market. Then the billings forecast number is submitted to MMBP. When the outputs do not align with the market, the billings, consumption and EOH inventory numbers are revised until the outputs align with market knowledge. This process is completed monthly for the current quarter and for one quarter ahead. During periods of dramatic changes the process is completed on a weekly basis.

4.3. Development of Methodology

The first model focuses on predicting inventory. The current process does a top-down and bottoms-up approach because there is no retail inventory information available, Section 4.2. The first model uses an order-up-to policy and focuses on predicting retail inventory levels to improve the demand forecast error. This model's process is outlined in Section 4.3.1.

The second version of the model simulates the decision-making process within the supply chain. This model is a top-down process; end user sales data is used to create billings forecasts. The evolution of this model is outlined in Sections 4.3.2-4.

4.3.1. Order-Up-To Model

The first model uses an order-up-to policy. The model assumes that retailers and OEMs make ordering decisions to keep a constant amount of inventory. Retailers and OEMs have set level of inventory that they work to maintain.

$$EOH_{inv} = bWOI \tag{4.1}$$

The goal for the retailer is to keep the inventory at the set level, Equation 4.1. Every quarter the goal is to have the ending on hand, EOH_{inv}, inventory be a specific amount, *b*, of WOI. In this case WOI is equal to average weekly end consumption. Average weekly end consumption is set equal to quarterly end consumption over the number of weeks in a quarter, Equation 4.2.

$$\text{WOI} = Cons_W = \frac{Cons_Q}{13} \tag{4.2}$$

Having a set amount of inventory at the end of the quarter is the goal but the actual ending on hand inventory, $EOH_{actualinv}$, is the beginning on hand inventory, BOH_{inv}, minus the excess in inventory, Equation 4.3. The excess inventory is the amount ordered during the quarter minus the amount sold.

$$EOH_{actualinv} = BOH_{inv} + Orders - Sales \tag{4.3}$$

An order will be the difference between the desired EOH inventory and the actual EOH inventory, Equation 4.4.

$$Order_{Q+1} = EOH_{inv} - EOH_{actualinv} \tag{4.4}$$

Substituting for the values of EOH_{inv} and $EOH_{actualinv}$ results in Equation 4.5.

$$Order_{Q+1} = b\frac{Cons_Q}{13} - (BOH_{inv} + Order_Q - Sales_Q) \tag{4.5}$$

Intel does not know the amount of retail sales and orders therefore, substitutions for $Order_Q$ and $Sales_Q$ need to be determined. In the case of $Sales_Q$ a market forecast will be used and in the case of $Order_Q$ a third party data source, IDC, will be used. The retail number is inputted back into the current Intel model to get a billings number.

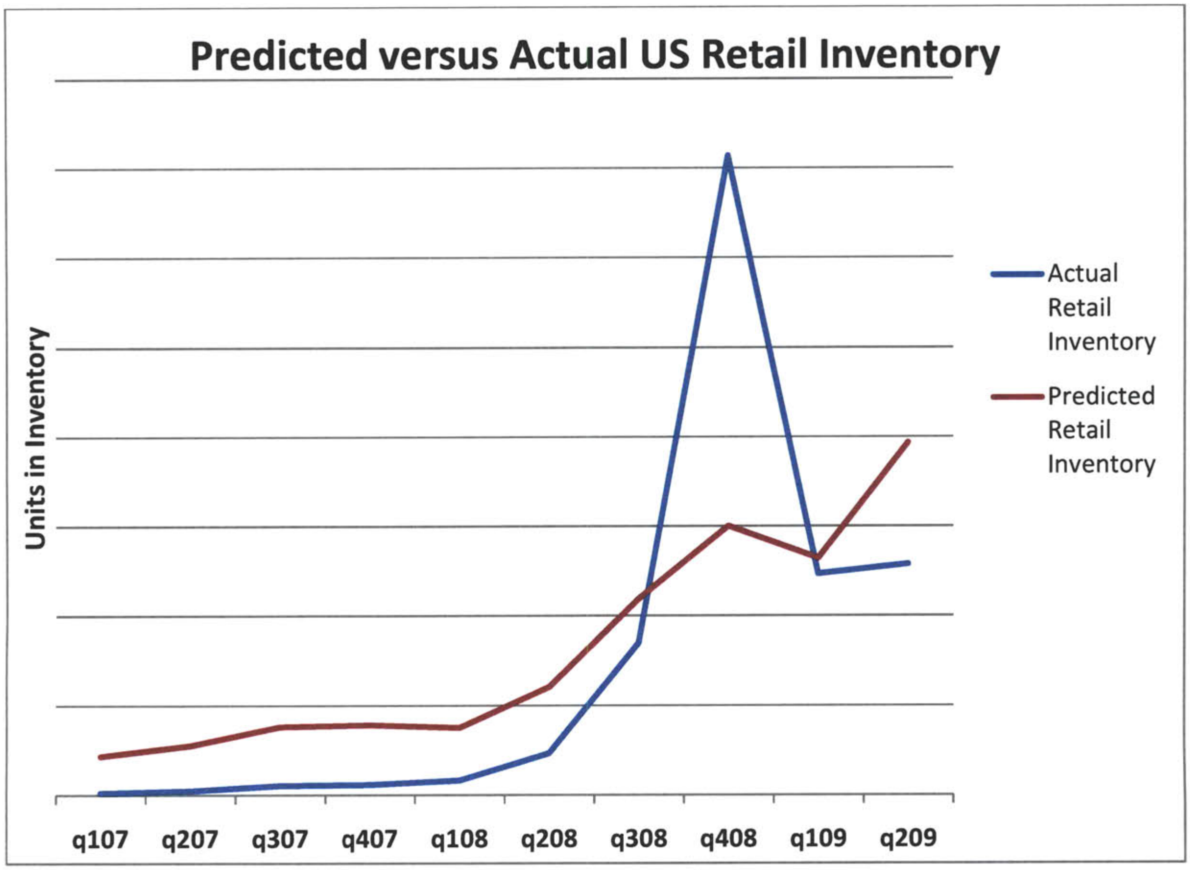

Figure 13: Predicted versus Actual Retail Inventory for Order-Up-To Model (NPD Group, 2009)

The models prediction of how much retail inventory is at the retail is approximately the same as the actual retail inventory until the economic downturn, Figure 13. During the economic downturn the amount of inventory is greatly under predicted by the model. The model predicts 20% of inventory reported by NPD (2009).

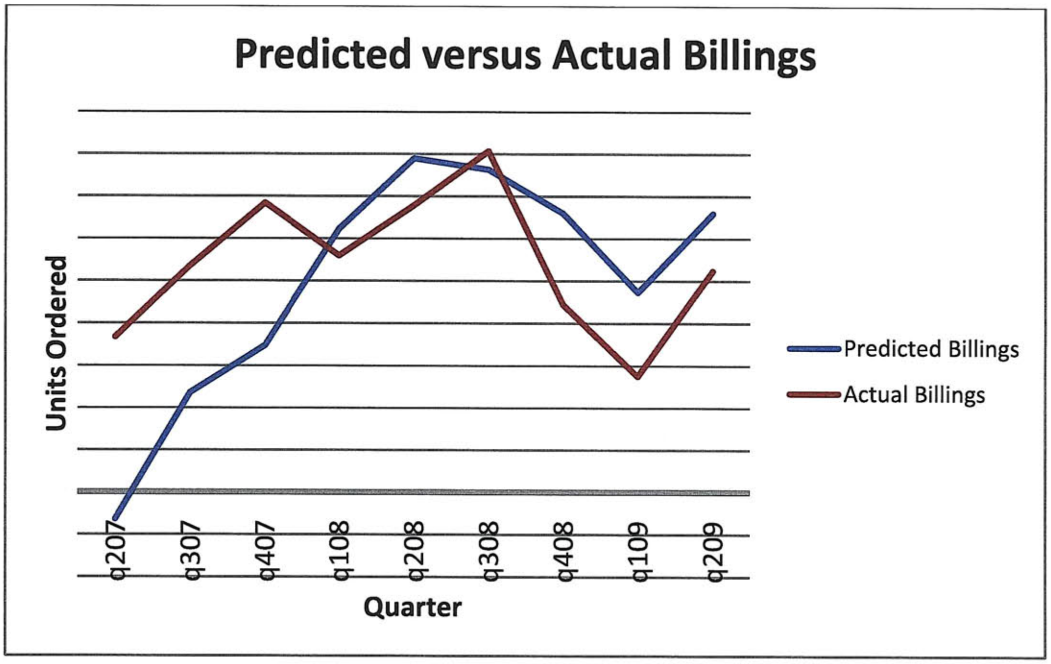

Figure 14: Predicted Billings from Order-Up-To Model versus Actual Billings

The order-up-to model does not provide an accurate billings forecast, Figure 14. The model under predicts the amount of billings prior to the economic downturn. The model then over-predicts during the economic downturn. The lack of trending shows that retailers and OEMs are not using this process to decide how much to order.

Table 1: Error Rate for the Order-Up-To Model

	Total Error	**Max**	**Min**
Billings Error	**45%**	**118%**	**-72%**

Table 1 confirms that the error rate for this model is too great for this to be how decisions are being in the supply chain.

4.3.2. Simple Heuristic Model with NPD Seasonality and Retail Inventory

The second model is similar to the heuristic model of Sterman (1989). This model adds in the time delays that the first model was missing.

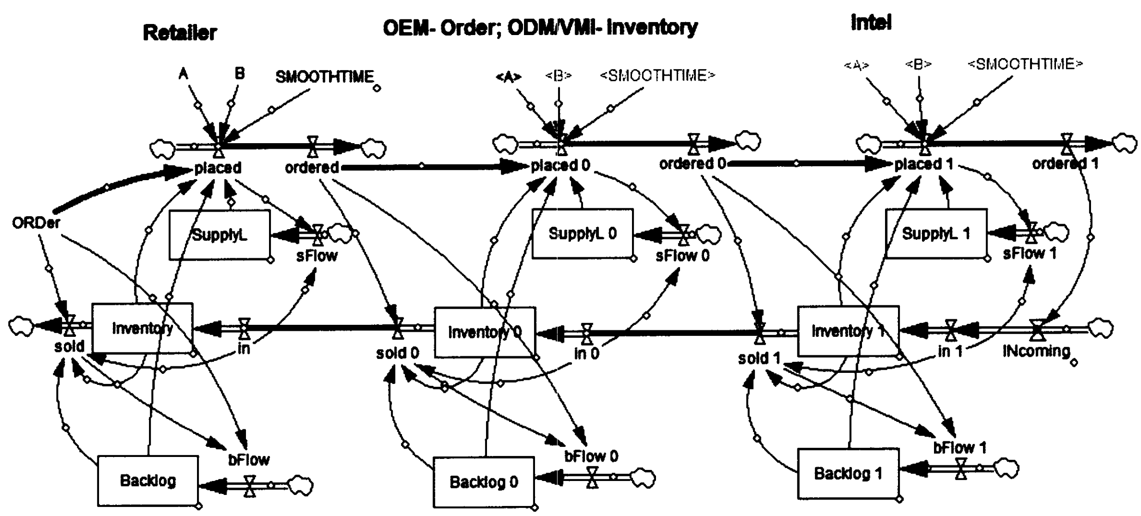

Figure 15: Heuristic Model of the Intel Supply Chain

The model works backwards through the retailer order process to the billings Intel receives. As shown in Figure 15, supply, demand, and backlog are considered in the order process. Supply is limited by the amount produced by the previous party. The amount of production or sales is limited by the amount of inventory. The arrival of orders into inventory is delayed by the amount of time it takes an order to arrive. Demand is smoothed by the function *smoothtime* and is weighted by variable A in the decision process of how much to order. As described in section 2.4, there are two order variables, A and B. Variable A is the weight of inventory on hand and B is the value of inventory in the supply chain. The value of B is how much an organization understands the time lags of the system.

The Intel supply chain model starts out using the information from Chapter 3 on the time delay retailers work to when ordering products. The first step is for retailers to predict demand one quarter ahead. They need to determine how much they think they are going to need when the product arrives, Equation 4.6.

$$DemandForecast_{retail} = \Delta Q\% * Sales_{Q} \tag{4.6}$$

This is the expected demand for retailers. Prior to placing an order the retailers adjust demand by the amount of excess inventory. This makes the value of variable A and B both one in Sterman's model. The final result is the demand forecast minus excess inventory, Equation 4.7.

$$Orders_{retailer} = \Delta Q\% * Sales_{Q} - \mathrm{ExcessInv}_{\mathrm{Q}} \tag{4.7}$$

The retailers' orders are then submitted to the OEMs. As described in Chapter 3 all computer are build to order. Since OEMs do not build computers until they are ordered by retailers there is no forecasting of what to manufacture for the retailers. The lag between order and delivery is included in the retailers' order to the OEMs. This means the retail order is directly the demand signal for the OEMs, Equation 4.8.

$$DemandForecast_{OEM} = Orders_{retailer} \tag{4.8}$$

Like retailers, OEMs also review their inventory levels prior to ordering. OEMs order their demand forecast minus any excess inventory, Equation 4.9.

$$Billings = Orders_{OEM} = Orders_{retailer} - \mathrm{ExcessInv}_{\mathrm{Q}} \tag{4.9}$$

Intel does not know the retailers' demand forecast process, so a substitute for a quarterly change,$\Delta Q\%$, must be determined to figure out how much retailers are ordering. In Figure 8 we saw that the WOI follows seasonal trends. Looking at sales on a quarterly basis, Figure 16, the sales pattern has a clear quarterly trend.

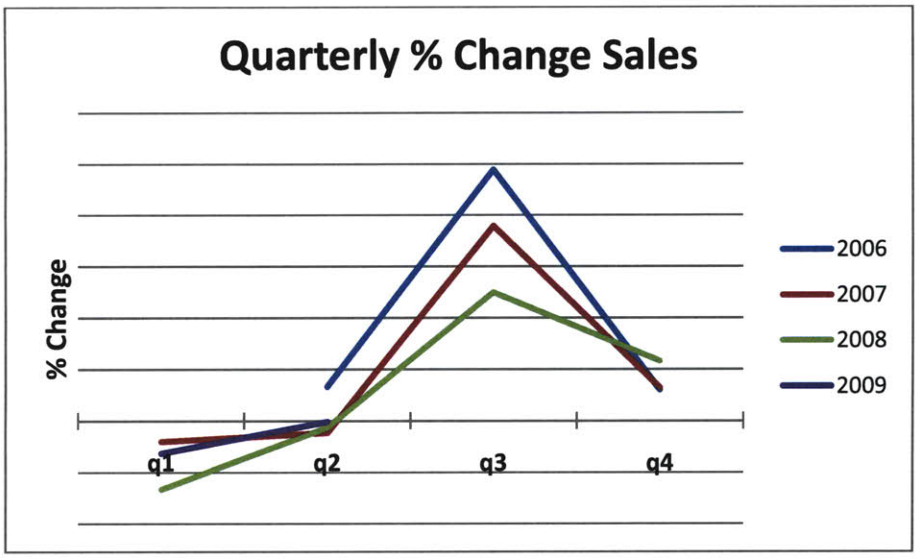

Figure 16: Quarterly Change in Sales Year over Year.

Substituting in the seasonality trend found above, results in Equation 4.10.

$$DemandForecast_{retail} = NPD\ Sesonality * Sales_Q \qquad (4.10)$$

Intel does not have access to retailers' sales and inventory data; therefore, the next step is to determine a substitute for sales and excess inventory that can be used in the model. Third party data and Intel's market share can be used as a quarterly sales number or market forecast. To determine the excess inventory we can use previous quarter's information. We can pick a quarter to be the baseline quarter and use this quarter as the level of inventory the retailer would like to keep. The excess inventory for that quarter would be what was ordered, $Orders_{retailer}$ minus what was sold, MSF_Q, Equation 4.11.

$$Orders_{retailerQ} =$$

$$NPD\ Sesonality * MSF_Q - (Orders_{retailer\ (Q-1)} - MSF_Q) \tag{4.11}$$

An alternative for OEM excess inventory also needs to be determined because OEMs do not give these numbers to Intel. Like in the retail case, a baseline quarter will be chosen and inventory changes after that quarter will be considered excess inventory. The same quarter that is used as the baseline for the retailers will be used as the baseline for the OEMs. Intel knows how many CPUs it sends to the OEMs, so the number of CPUs entering the OEMs inventory every quarter is known. Intel has agreements in place where OEMs send the number of CPUs leaving the factory every quarter, ex-factory. The excess inventory would be the difference in billings and ex-factory, Equation 4.12.

$$Billings = Orders_{retailerQ} - \left(Billings_{q-1} - Exfactory_{q-1}\right) \tag{4.12}$$

These equations make up the new heuristic forecast model. The timing of data used in the model and the results are analyzed next.

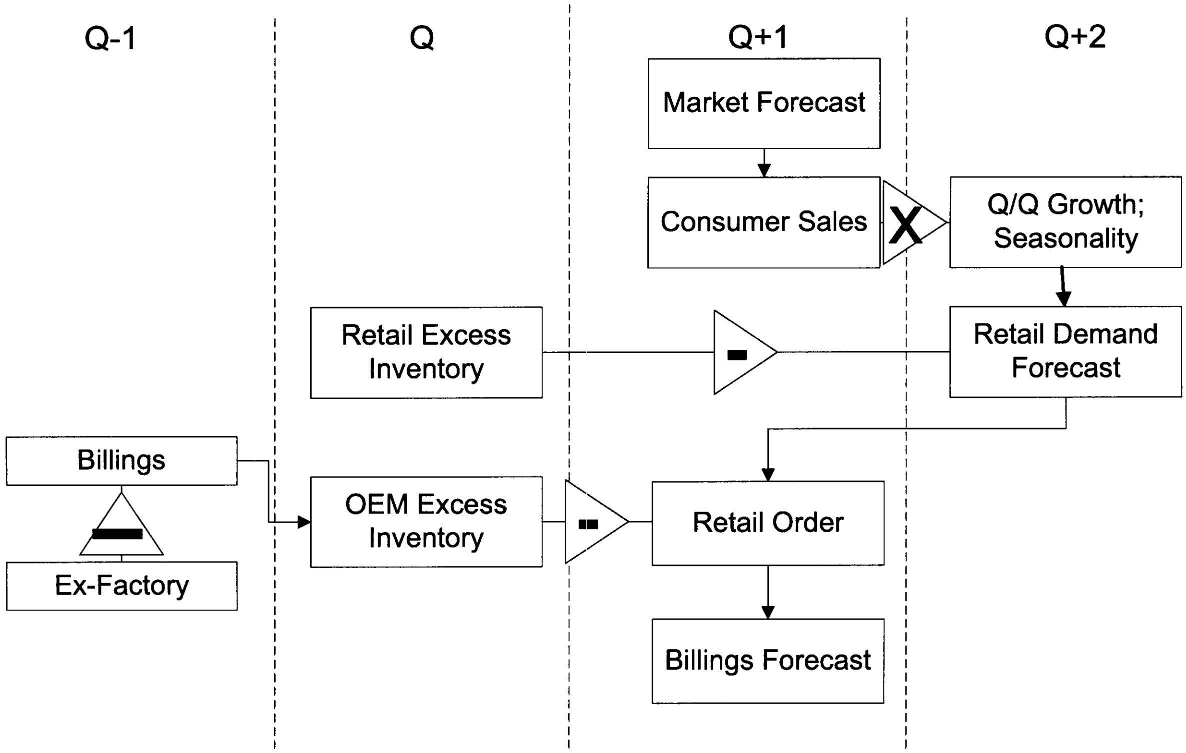

Figure 17: Control View of Heuristic Model

Figure 17 shows how the time lags are used in the model. The prior quarter's billings and ex-factory are used to calculate the current quarter's inventory, for OEMs. The retail excess inventory is calculated for the current quarter using the prior quarter's ex-factory and retail sales. The seasonality for the upcoming quarter is used to create a retail demand forecast.

To analyze the model a few third party sources are used. IDC is a third party database that tracks OEMs' reports of orders to retailers. The comparison of the model's retail orders prediction to IDCs record of retail orders results in Figure 18. Using a retail order comparison as well as a billings comparison allows for the pinpointing of error locations in the model's simulation of the supply chain's decision making process.

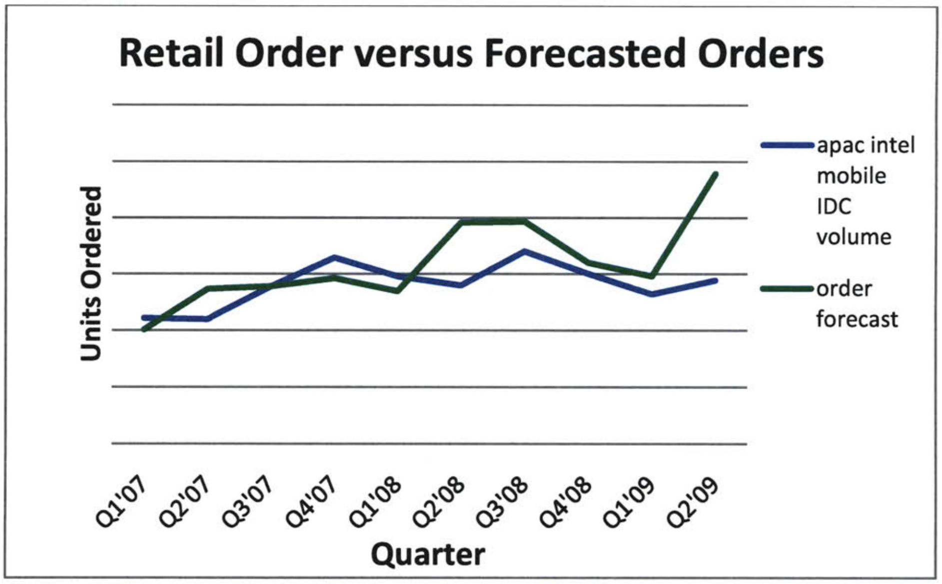

Figure 18: Retail Actual Orders versus Forecasted Orders

The comparison shows the model is capable of general trends but needs improved accuracy. The model order forecast is high and it peaks a quarter before the actual ordering peak. The ordering trends are also flip-flopped during the more stable sales period prior to 2008.

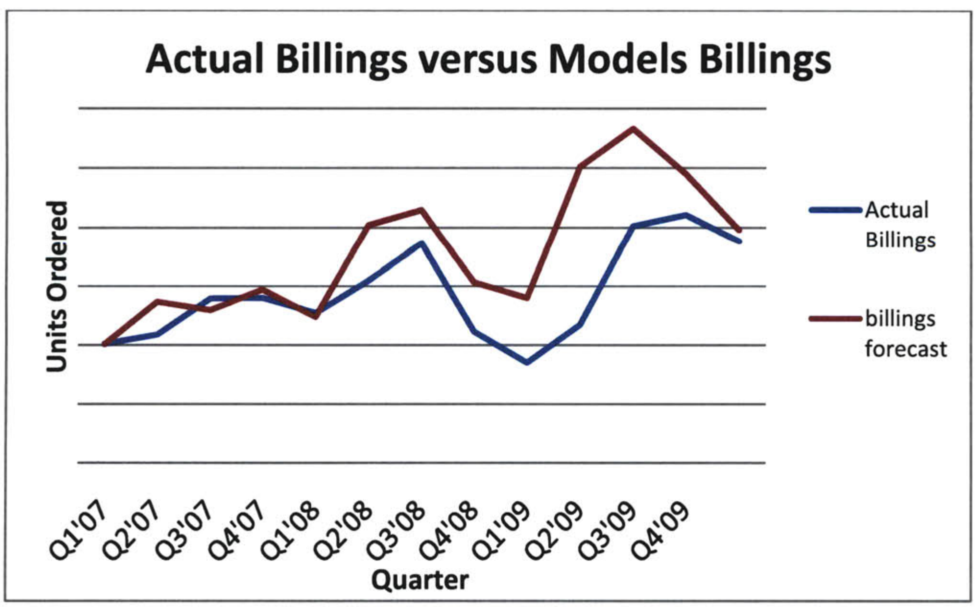

Figure 19: Actual Billings versus Model Billings for Heuristic Model with Retail Inventory

Figure 19 shows that the model over forecasts the amount of billings throughout the economic downturn. Table 2 shows the error rates for the NPD seasonality model with both OEM and retail inventory.

Table 2: Error Rates for NPD Seasonality with Inventory Model

	average	max	min
Retail Error	18.8%	4%	-23%
Billings Error	24.1%	55%	-20%

While the error percentage of the new model is reasonable, it is not sufficient. The next step is to adjust the variables used to model the decision-making process of the retailers and OEMs.

4.3.3. Simple Heuristic Model without Retail Inventory

Reviewing the first heuristic model, reveals that the number being used as retail inventory may be increasing the overall model error. The next step is to explore why this may be causing an increase in the error and what the error in the model would be without the retail inventory.

Comparing the OEM decision process to the retail order process, it becomes clear that there are far more retailers than OEMs. The order predicted is the sum of all the retailers' orders, Equation 4.13.

$$Orders_{predicted} = \sum_{All\ Retailers} \text{Orders}$$

$$= \sum_{All\ Retailers} [Predicted\ Sales - Excess\ Inventory] \tag{4.13}$$

The predicted retail order is an aggregation of all the individual retailers' decisions. Sterman's model assumes one retailer judging the weight of future demand, inventory and time lags. In this model, the value for predicted retail order is a sum of this decision of each individual retailer. No one retailer owns the entire excess inventory. Each retailer owns a portion of the excess inventory and must make decisions on how to value this inventory in its next order process. While one retailer may place heavy weight on the excess inventory in its possession, another may not concern itself with its excess inventory. The two retailers' valuations of their inventory negate each other. The summation and average of thousands of retailers' decisions makes variable *B*, the weight of inventory at retailers in the decision-making process, zero.

Another problem with the way excess retail inventory is calculated is that the numbers for retailer inventory are highly estimated. The number being used for sales data is an estimation made by Intel using multiple third party sources, economic indicators, and other market knowledge. While this estimation is better and more complete than any third party database, it still a significant portion of the error. Using this number to calculate inventory double counts the error in the sales estimation. Not only does the sales number include the error but the inventory information includes the error.

For this second model weight of the retail excess inventory, variable A, was changed to zero and the model was analyzed again. The resulting model flow is shown in Figure 20.

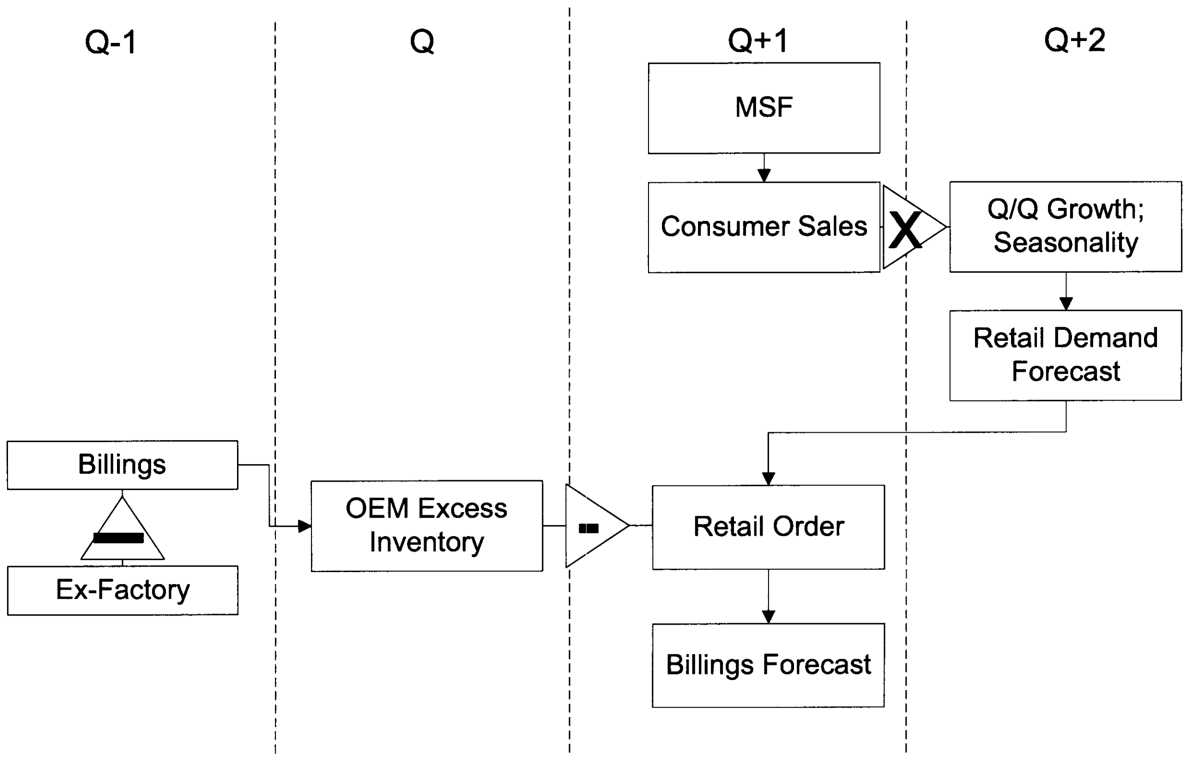

Figure 20: Control View of Model without Retailer Excess Inventory

The model still contains the excess inventory for OEMs for two reasons. There is data on the excess inventory of OEMs from the OEMs. The OEMs report their ex-factory information to Intel, and Intel knows how many chips were sold. Also, the amount of arbitrage of CPUs is very small for Intel due to their contracts and other governance. Therefore, arbitrage is not a factor in calculating the excess inventory of the OEMs, and the calculated number can be considered accurate.

The second reason why the weight of the OEM excess inventory remains unchanged is that there are far fewer OEMs. One OEM can have up 60% of the excess inventory. With such a significant portion of inventory owned by each OEM, the inventory has a considerable reduction on the OEM's order to Intel. The OEMs fully value the inventory on hand; the value of variable A will remain equal to 1 for the OEMs.

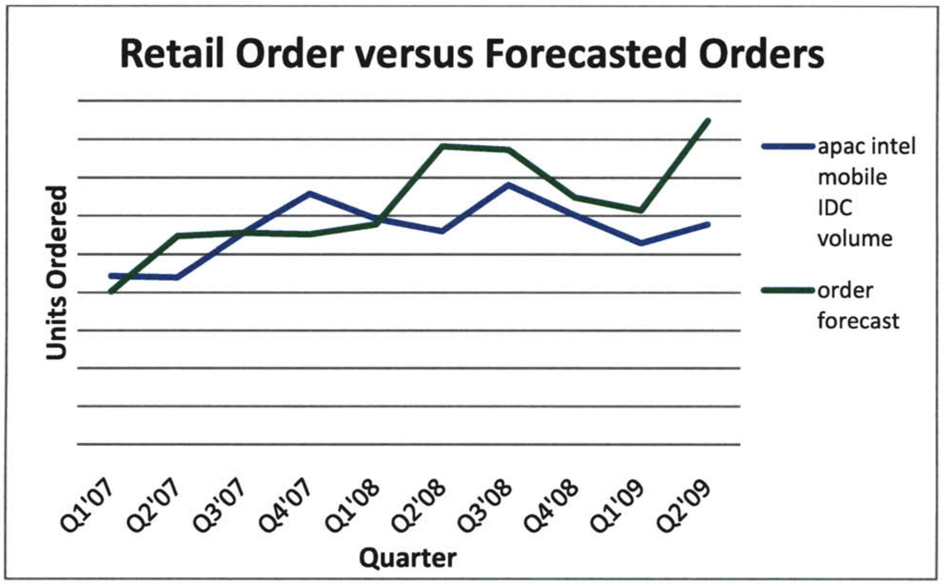

Figure 21: Retail Orders Predicted by the Model without Inventory versus Actual Orders

The new model has very similar trending of retailer orders versus the previous model, Figure 21. The main difference is that the new model better predicts the actual orders of retailers during the recovery quarters. The similarity shows that the retailer excess inventory does not have a significant effect on the retailers' orders. It also shows that putting too much weight on the value of retail excess inventory can have a negative effect on accuracy during the economic recovery.

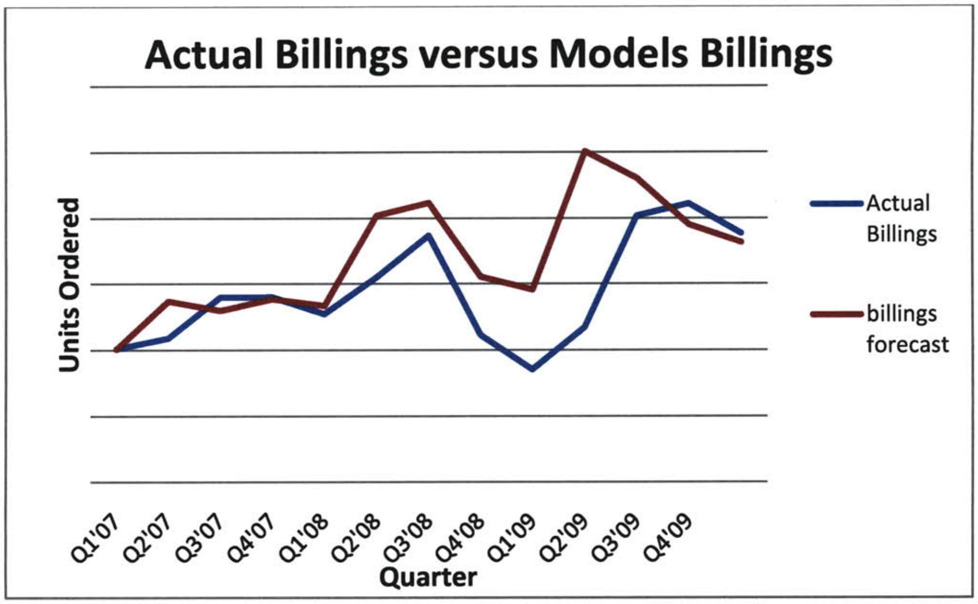

Figure 22: Predicted Billings of the Model without Inventory versus the Actual Billings

The billings prediction of the new model follows the same trend as the initial heuristic model prior to the economic downturn, Figure 22. The new model works better at predicting the billings after the economic downturn. Removing the excess inventory calculation for retailers has removed the over prediction of the economic recovery.

Table 3: Error Rates for the Model without Inventory

	average	max	min
Retail Error	**15.0%**	**32%**	**-19%**
Billings Error	**19.9%**	**48%**	**-8%**

Table 3 shows the error rates for the new model. Both the retail and billings error has been improved. The billings error has gone from 24% to 19%. The error range has also been decreased for both billings and retail error.

4.3.4. Simple Heuristic Model with Global Weighted Seasonality

The next revision of the model adjusts the seasonality. The previous models used seasonality based on US sales data. US seasonality provides a good representative number because the US is such a large amount of sales, but the US is not the only country consuming products from Asia Pacific-OEMs. A global seasonality provides better representation of the quarterly changes in sales.

$$Weighted\ Seasonality = \sum \frac{IDC\ Intel\ GEO_{Qx}}{IDC\ OEM\ Intel_{Q}} * GEO_x Seasonality_{Q+1} \tag{4.14}$$

A new weighted seasonality was developed that creates seasonality each quarter based on the percentage of sales to each GEO. Each GEO has a predictable seasonal demand. The quarterly changes in seasonal demand for each GEO are multiplied times the amount of sales in that GEO that quarter to determine the global retail seasonality, Equation 4.14.

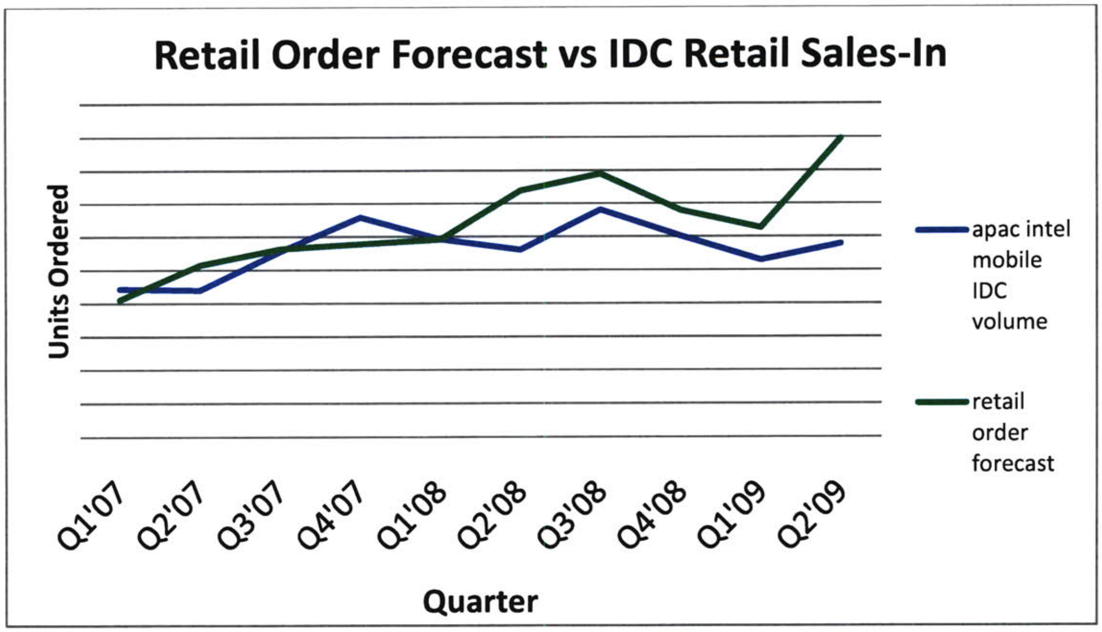

Figure 23: Predicted Retail Orders of the Weighted Seasonality Model versus Actual Orders

The predicted orders of the new model still follow the same trend and are comparable to actual orders, Figure 23. The predicted retail orders are improved by the weighted seasonality model. Although it is only slight improvement, the retail order trend is better than the previous model.

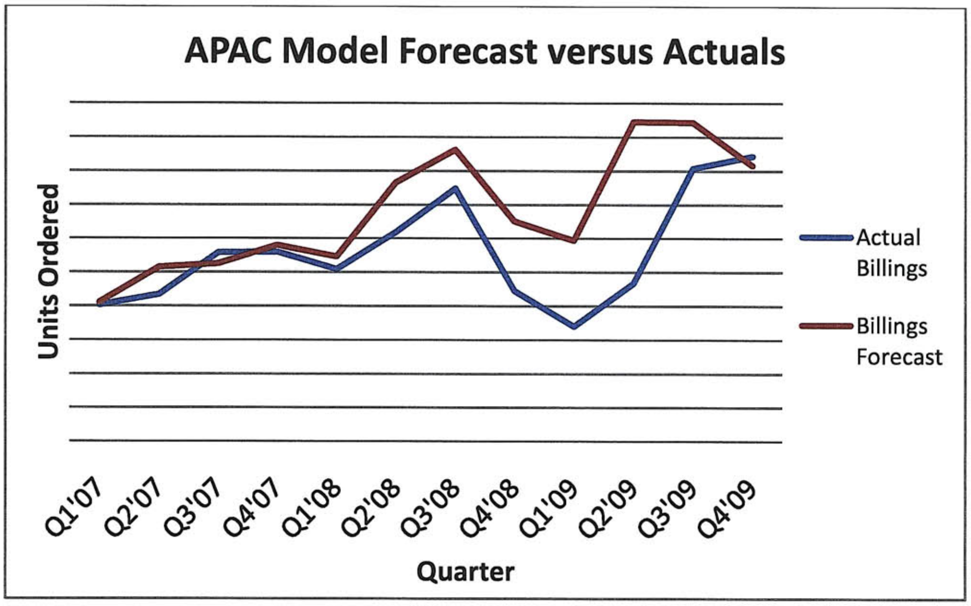

Figure 24: Predicted Billings of the Model with Weighted Seasonality versus Actual Orders

Figure 24 shows the new billings trend. The trends are similar to the previous models and follow the actual trends. The model is still over forecasting during the economic downturn but the final number is more accurate than the previous model. There is not a large under forecast for Q4'09.

Table 4: Error rates for the Model with a Weighted Seasonality

	average	max	min
Retail Error	16.8%	35%	0%
Billings Error	17.4%	50%	-6%

The error rates have been improved for the new model, Table 4. The billings error rate decreased 2% from 19.9% to 17.4%. The range of error has also been improved.

4.4. Measurement of Results

The four different models each successfully made improvements in the error in forecasting, Table 4.

Table 5: Error rate Comparison

	Billings Error	Retail Error
Order-Up-To Model	45%	n/a
NPD Seasonality + Retail Inventory	24.1%	18.8%
NPD Seasonality	19.9%	15.0%
Weighted Seasonality	16.8%	17.4%

There is a significant improvement in the model error when changing from an order-up-to model to a simple heuristic model. The removal of the retail excess inventory also made a large improvement in the model error.

Sections 4.1-4.3 outline how the model was optimized to match the actual decision-making process used by parties within the supply chain as best as possible with the data available. This optimization was done using historical estimates of the demand created by Intel once they received third party data. Historical estimate data worked best to ensure that the model reflected the actual decision making process by reducing the error of the consumer demand estimation.

Unfortunately, Intel does not have third party data for the current quarter available when making demand forecasts. Intel receives data thirteen weeks or more after a quarter closes. Intel uses prior quarter data and other external factors to predict consumer demand for the past quarter up to four quarters ahead. The predicted data is used by the demand forecast team when determining their billings estimate. To accurately compare the capability of the new model to the current process the same input data must be used. Changing the input data to predicted data will increase the error but will allow for an accurate comparison of the two billings prediction processes.

Changing to the sales forecast number to one that would have been available at the time the model would be run, will allow the model to be compared to the current forecast process. This means a sales forecast number created one quarter prior to the forecasted quarter should be used.

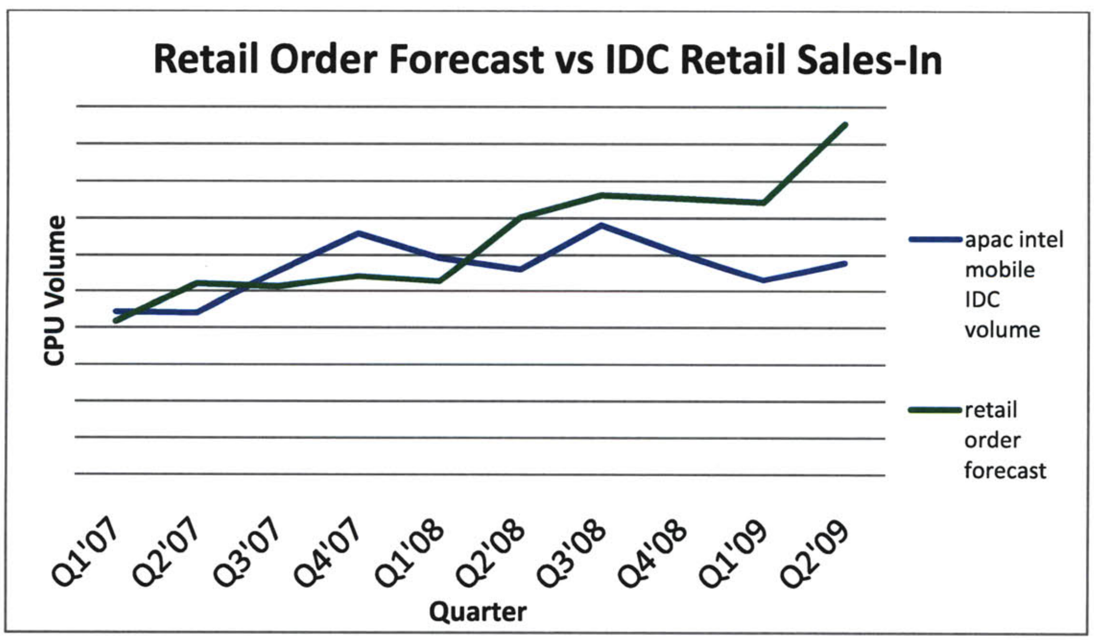

Figure 25: Retail Orders versus Models Orders using One Quarter Ahead Sales Forecast

The resulting prediction of retail orders is shown in Figure 25. The error in retail orders has increased versus the version of the model using the latest sales forecast numbers. The sales forecast has not only increased orders during the economic downturn but during the recovery as well.

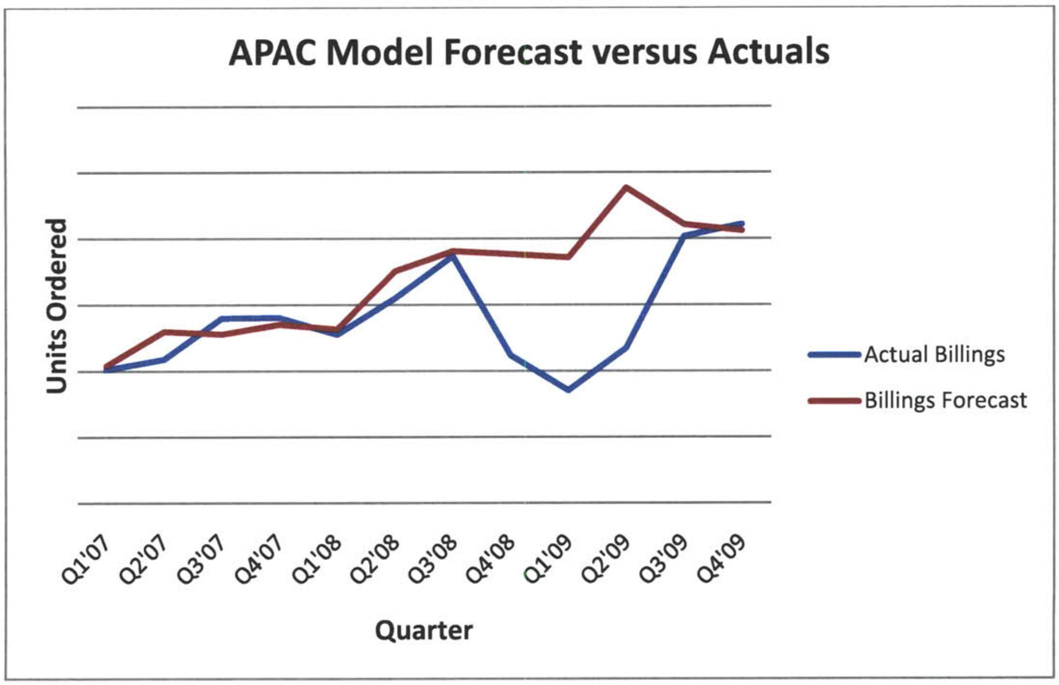

Figure 26: Billings Predicted by the Model versus Actual Billings using One Quarter Ahead Sales Forecast

Figure 26 shows the new billings prediction. Similar to the new prediction of retail orders, the billings prediction is much greater than actual billings during the economic downturn and recovery. The next step is to compare this result with the current demand forecast process error and determine the difference between the two methods.

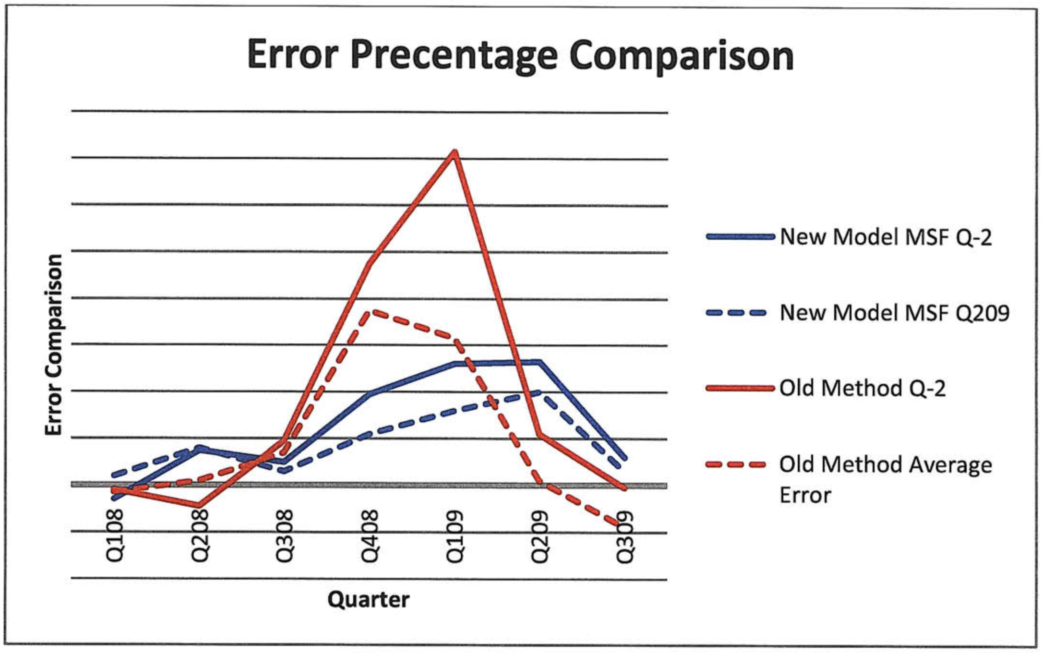

Figure 27: Comparison of Error of New Model versus Past Methods

Figure 27 shows the error rates of the two models. The blue lines are the error rates for the newly proposed model and the red lines are the error rates for the current forecast method. The solid lines are for the first submitted forecast one quarter prior to order submission. The dotted line is the average error for all forecasts submitted. The results show that both models improve their error from first to last submission. The new model greatly improves the forecast error during the economic downturn but it does not do as well during the recovery and the period prior to the economic downturn.

The error during the economic recovery can be explained by the changes in the model caused by using a new earlier sales forecast number. The sales forecast number over predicted the amount of orders during the economic downturn and recovery. The Intel sales forecast during the economic recovery had significant impact on the model's accuracy. The model is heavily dependent on the accuracy of Intel's sales forecast.

Prior to the economic downturn Intel's model worked better than the new model. Intel is very good judging demand during stable periods. The new model does not use judgment processes. Judging the model's demand based on external market knowledge should improve the new model during stable economic periods. The normal demand has a very stable and predictable pattern. The economic downturn had the same pattern, but when comparing quarterly sales, there was bit of a slow-down. The increase in inventory created a ripple through the system as retailers' decreased orders more than the change in demand and OEMs decreased orders more than retailers.

The slow-down showed up in the form of a pile up of inventory in the OEMs that then decreased their orders when they realized they failed to account for the incoming inventory in their previous order. This is why the inventory in the current quarter is not used in the model; the inventory of the previous quarter is used instead. OEMs do not account for CPUs ordered until they arrive. Then when the OEMs realize they have too many, they cut orders in the quarter they received this inventory. The model accounts for this change while Intel's does not. This is why the model was better during the downturn.

4.5. Implementation

The model was supplied to Intel as a multi-sheet excel workbook,

Figure 28. There is one sheet to input common APAC (Asia-Pacific GEO) demand information, another for common variables, and the last APAC model sheet has the APAC GEO demand forecast results. Then there are tabs for each regional account that are used to input the account demand data and show each accounts demand forecast.

Inputs

	Q1'07	Q2'07	Q3'07	Q4'07	Q1'08	Q2'08	Q3'08	Q4'08	Q1'09
MSF Q+2 Sales	3,000	2,500	3,500	2,800	3,000	2,500	3,500	2,800	3,000
MSF Q-1 to Q-4 Sale	2,200	2,200	2,900	3,500	2,200	2,200	2,900	3,500	2,200
IDC total Mobile	3,000	3,000	3,200	3,300	3,000	3,000	3,200	3,300	3,000
apac mobile total	600	500	500	800	600	500	500	800	600
apac intel mobile	300	400	300	200	300	400	300	200	300
ASMO	75%	60%	60%	60%	75%	60%	60%	60%	75%
EMMA	10%	10%	10%	10%	10%	10%	10%	10%	10%
APAC	15%	30%	30%	30%	15%	30%	30%	30%	15%
ex-factory	5,000	5,000	4,000	2,000	5,000	5,000	4,000	2,000	5,000
Billings	3,000	4,000	3,000	2,000	3,000	4,000	3,000	2,000	3,000

APAC Inputs / Results- Billing Forecast / variables / apac prediction model / account inputs / account results / Regional Account Billings Model / A

Outputs

	Q1'07	Q2'07	Q3'07	Q4'07	Q1'08	Q2'08	Q3'08	Q4'08
billings forecast	511	420	2,453	1,500	1,309	431	2,380	1,110
End Consumption fore	-	511	420	453	500	309	431	380
retail error %		5%	34%	60%	3%	7%	21%	-82%
billings % error		-853%	-22%	-33%	-129%	-827%	-26%	-80%

billings % error ave	258%
Billlings % error max	100%
billings % error min	-853%

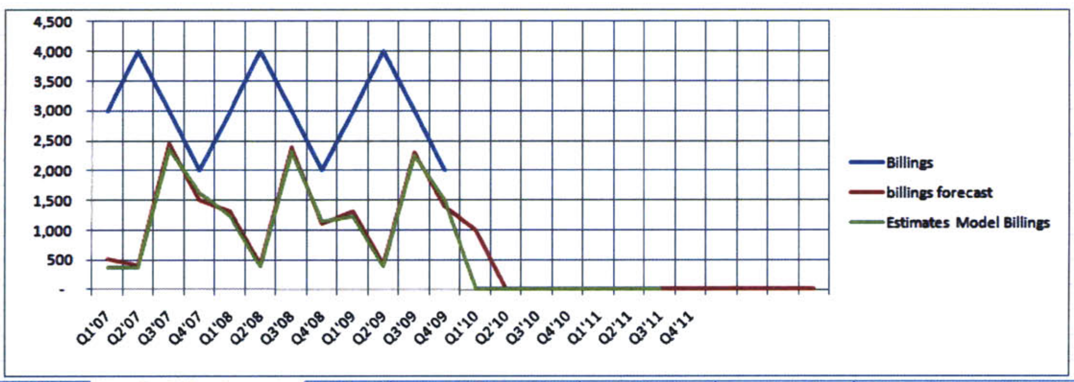

APAC Inputs / Results- Billing Forecast / variables / apac prediction model / account inputs / account results / Regional Account Billings Model

Figure 28: View of Model Excel Sheet APAC Input Tab and Results Tab

The task of inputting the data into the new forecast model has been divided amongst the OEM sales teams and demand forecast teams. The demand forecast team is responsible for the total Asia-Pacific GEO, APAC, model and inputting the third party sales data. The OEM sales teams add in the ex-factory and OEM inventory information for their OEM.

The individual OEM tabs give the OEM sales teams a reference to compare the demand forecasts from the OEMs. OEM sales teams can determine quickly if the numbers they are getting form the OEMs are accurate. The benefit of improved ability to judge OEM forecasts encourages the sales teams to fill out the model.

The OEM tabs are then sent back to the demand forecast team. The completed tabs give the demand team consolidated OEM inventory information. The demand forecast team then analyzes the forecasts and creates an APAC demand forecast. The APAC demand forecast is then compared against the total OEM submitted demand forecasts, and the APAC demand forecast is revised once more. Once the demand forecast judgment is complete, it is sent to MMBP for the creation of global demand forecast as described in Chapter 3.

The research from this project and the model has resulted in improved communication between the regional sales teams and the APAC demand forecast team. The model has shown the importance in understanding customer demand to better understand the order Intel receives. The project has also brought to attention the need for centralized information sharing between APAC customer sales teams, regional sales teams, and the APAC demand forecast team. The model improves the information sharing between customer sales teams and the demand forecast team, but it does not improve the information sharing with the regional sales teams. The future work needed and outcomes of that work are discussed further in Chapter 5.

5. Observations and Recommendations

5.1. Overview

The current model and proposed model are dependent upon the accuracy of the information Intel receives from OEMs and retailers on inventory and sales. The reliability of the external inventory and sales data can be improved in several ways: improving the VMI (vendor managed inventory) relationships and improving retail sales information sharing. The following section reviews the opportunities to improve the model input data reliability and the expected results of improving the data.

5.2. Vendor Managed Inventory

Lead time is a major cause of the bullwhip effect (Disney & Towill, 2003). VMI removes another party from the supply chain, thus decreasing supply chain length after a CPU leaves Intel. Decreasing the supply chain after Intel's component warehouse gives Intel greater visibility of the customers supply chain. VMI increases coordination between customers and suppliers and has been shown to reduce bullwhip costs by 4.7% compared to 2% for traditional information sharing (Aviv & Federgruen, 1998). VMI should be reducing costs and improving forecast error for Intel but these results have not been seen.

Customers currently using Intel's VMI are keeping more inventory at their manufacturing sites as a result of Intel's VMI policies described in Section 3.3.2. Keeping more inventory at ODMs is counter to expected results of using a VMI hub. The greater the inventory at the ODM, the lower the coordination between Intel and its customers and the less insight Intel has into the customer supply chain. If Intel provides the buffer inventory, they know what is going on with demand signals; if an OEM or ODM holds the inventory than Intel does not know about changes in inventory levels until the next order.

One of the keys of a VMI is the increased communication and trust needed between customers and suppliers and the expectation that a pre-determined quantity be in stock when requested (Curtis, 2002). Another problem with VMI is that it is less useful when product demand is stable or has long planning periods (Smaros, Lehtonen, Appelqvist, & Holmstrom, 2003).

To improve the current VMI usage and effectiveness in providing coordination with the customer, Intel needs to re-evaluate its VMI policies. The biggest change that needs to take place is an increase in the expected availability of materials. Intel will have to trade holding more buffer inventory for having reduced swings in orders. The relationship and interaction with customers also needs to improve beyond that of their current sales process. Intel needs to develop services in logistics coordination and order fulfillment similar to those distribution hubs currently provide. Without these services the value of Intel's VMI will not be worth losing the logistics skills of a distributor.

5.3. Approaches to Improve Demand Information Flow

The lack of information sharing is having a lot of different effects on Intel's order fulfillment process. For example, the current demand forecast model and the one developed in this paper are limited by the accuracy of the sales information. For the demand forecast process there are many levels of information sharing that can be done.

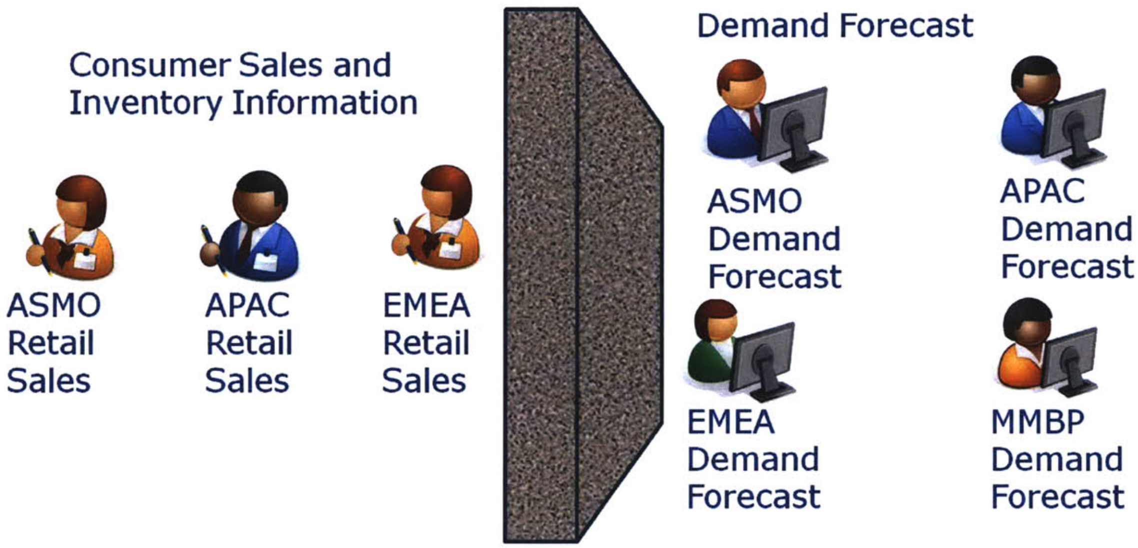

Figure 29: Divides in the Parties Controlling Demand Data

The first is the relationship between retail sales teams and demand forecast teams. Intel has retail sales teams that interface with retailers, Figure 29. Retail sales teams help retailers understand the Intel products; retailers get products; and monitor retail sales. During the data collection phase of this project it was determined that retail teams are monitoring and tracking retail sales and inventory. The sales teams collect information about retailers in their region but there is no transfer of this data from one region to another or to the demand forecast teams.

Similarly, there is a demand forecast team for each region. Demand forecast teams work with OEM sales teams to consolidate the OEM forecasts for their GEO. The OEM sales teams may talk to the retail sales teams when they judge demand, but any retail sales information is not conveyed to the demand forecast team for when they judge the OEM sales teams demand.

The databases being used or data being collected on OEM or retail sales is not being shared across GEOs. The APAC demand forecast team may only be responsible for judging

APAC OEMs' demand, but those OEMs are selling their products all over the world. For demand forecast teams to be able to judge demand they need to look ahead of the OEM demand to the retail demand. They need to be alerted to trends in demand ahead or changes in OEM ordering patterns.

The first step to improving the level of the current consumer sales data is to improve the transfer of information from the retail sales teams to demand forecast teams. Centralized demand information has been shown to reduce the bullwhip effect (Chen, Drezner, Ryan, & Simchi-Levi, 2000). This is a free step that could be easily completed without requiring the cooperation of retailers and OEMs.

The next step is to develop a relationship with retailers that results in consumer sales and inventory information being shared with Intel. Information sharing between organizations has been shown to reduce the bullwhip effect (Lee, So, & Tang, 2000). This step will only contribute to improving Intel's responsiveness if the information sharing across the groups is improved. It will be of no help if the retail teams get sales and inventory data from retailers and the information does not make it to the demand forecast teams. Retailer data would be most useful if it goes into making a MMBP global sales forecast and then gets translated back down to the GEO demand teams for predicting their billings.

5.4. Future Work

Retail data information sharing would also improve the granularity and timing of the data available to Intel. Current information comes in quarterly buckets and does not get released until a quarter after that quarter has closed. This means that Intel next quarter sales forecast is based on forecast of current quarter sales instead of actual data.

Getting retailers to share the shipping rate percentage along with having weekly sales information would also improve the error of the model discussed in this paper. As discussed previously in Chapter 4 this model has a fixed supply chain length of 14 weeks. This can be reduced when the percentage of computers shipped to retailers by air is increased. During the economic downturn retailers wanted to reduce inventory risk and increased the percentage of orders shipped by air. The tracking of this percentage along with data at the weekly level would allow the model to work with a varying average supply chain length.

6. Conclusion

Intel's current demand process is heavily reliant upon customers' forecasts. There is no incentive for customers not to overinflate their orders or to cancel their orders at the last possible moment. A new model based on Sterman's heuristic model that uses a top down approach can reduce the amount of forecast error that Intel is experiencing due to the bullwhip effect and phantom orders.

The new model is very dependent upon the sales forecast Intel creates using the data available. Intel needs to improve information sharing and coordination in the supply chain. The first step is improving the flow of sales information to regional demand forecast teams. The next step is to get retailer point of sale, inventory, and shipment method information. Intel can also improve supply chain coordination through improving its VMI policies to better meet customers' expectations. Taking these steps would greatly help Intel be aware and prepared during the next economic downturn.

Bibliography

Aviv, Y., & Federgruen, A. (1998). The Operational Benefits of Information Sharing and Vendor Managed Invenotry Programs. *Working Paper Washington University* .

Chen, F., Drezner, Z., Ryan, J. K., & Simchi-Levi, D. (2000). Quantifying the Bullwhip Effect in a Simple Supply Chain: The Impact of Forecasting, Lead Times, and Information. *Management Science , 46* (3), 436-443.

Chow, J. (2004). LFM Master's Thesis: Analysis of New Approaches to Improve the Customer Responsiveness of Intel's Microporcessor Supply Chain. MIT Sloan School of Management and MIT Department of Civil and Enviromental Engineering.

Curtis, D. (2002). What It Take to Implement and Maintain Successful Supplier-Managed Inventory Programs. *Hospital materiel Management Quarterly , 21* (4), 25-31.

Dedrick, J., & Kraemer, K. L. (2007, March 1). *Market Making in the PC Industry.* Retrieved January 15, 2010, from http://www.escholarship.org/uc/item/86979652

Deffree, S. (2009, January 16). *Intel Q4 Profit Plummets, No Q1 Sales Outlook Given, Continues Investing.* Retrieved January 15, 2010, from Electronic News: http://www.edn.com/article/CA6630205.html

Disney, S., & Towill, D. (2003). Vendor -managed Inventory and Bullwhip Reduction in a Two-level Supply Chain. *International Journal of Operations and Production Management , 23* (5-6), 625-651.

Dvorak, P. (2010). Clarity Is Missing Link in Supply Chain. *Wall Street Journal (Eastern Edition)* , A.1.

Fine, C. H. (1998). *Clockspeed : Winning Industry Control in the Age of Temporary Advantage.* Cambridge: Perseus Books.

Forrester, J. (1961). *Industrial Dynamics.* Cambridge, MA: MIT Press.

Fortune 500. (2009, May 4). *Fortune 500: Our Annual Ranking of America's Largest Corporations.* Retrieved December 28, 2009, from CNN Money: http://money.cnn.com/magazines/fortune/fortune500/2009/full_list/

Graves, S. (1999). A Single-Item Inventory Model for a Nonstationary Demand Process. *Manufacturing and Service Operations Management , 1* (1), 50-61.

Gupte, K. (2009). LFM Masters Thesis: Impact of Retail Sales and Outsourced Manufacturing on a Build-to-Order Supply Chain. MIT Sloan School of Management and MIT Department of Eletrical Engineering.

Intel. (2008). *Intel Corporate Website.* Retrieved December 28, 2009, from http://www.intc.com/intelAR2008/index.html

Lee, H., Padmanabhan, V., & Whang, S. (1997). The Bullwhip Effect in Supply Chains. *Sloan Management Review , 38* (3), 93-102.

Lee, H., So, K., & Tang, C. (2000). The Value of Information Sharing in a Two-Level Supply Chain. *Management Science*, *46* (5), 626-643.

Metters, R. (1997). Quantifying the Bullwhip Effect in Supply Chains. *Journal of Operations Management, 15*, 89-100.

National Bureau of Economic Reasearch. (2008). *Determiniation of the December 2007 Peak in Economic Activity*. Retrieved January 28, 2009, from http://www.nber.org/cycles/des2008.html

Pai, N. (2009, June). LFM Master's Thesis: Making Real-Time Precision Adjustments to World Wide Chip Production. MIT Sloan School of Management and MIT Department of Eletrical Engineering.

Rogoway, M. (2009, Janurary 7). *Intel Sales Plunged in Fourth Quarter.* Retrieved Janurary 15, 2010, from http://www.oregonlive.com/business/index.ssf/2009/01/intel_sales_plunged_in_fourth.html

Sahin, F., & Robinson, E. P. (2002). Flow Coordination and Information Sharing in Supply Chains: Review, Implications, and Directions for Future Research. *Decision Sciences*, *33* (4), 505-536.

Shah, A. (2010, January 25). *AMD takes processor market share from Intel, IDC says*. Retrieved Feburary 10, 2010, from Computer World: http://www.computerworld.com/s/article/9148638/AMD_takes_processor_market_share_from_Intel_IDC_says

Smaros, J., Lehtonen, J.-M., Appelqvist, P., & Holmstrom, J. (2003). The Impact of Increasing Demand Visibility on Production and Inventory Control Efficency. *Internationl Journal of Physical Distribution & Logistics Management*, *33* (4), 336-354.

Sterman, J. D. (1989, March). Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science* , 321-339.

Wright, D., & Xin, Y. (2008). Mitigating the bullwhip effect by ordering policies and forecasting methods. *International Journal of Production Economics*, *113* (2), 587-597.