

Managing Forecast Variability in a Build-to-Order Environment

by
Marshall Einhorn

Bachelor of Science in Mechanical Engineering, Tufts University (1998)

Submitted to the MIT Sloan School of Management and
Department of Mechanical Engineering
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 Manufacturing Program at the
Massachusetts Institute of Technology

June 2007

©2007 Massachusetts Institute of Technology. All rights reserved.


Signature of Author _____
MIT Sloan School of Management
Department of Mechanical Engineering
May 11, 2007

Certified by _____
David E. Hardt, Thesis Supervisor
Professor of Mechanical Engineering and Engineering Systems

Certified by _____
Roy E. Welsch, Thesis Supervisor
Professor of Statistics and Management Science

Accepted by _____
Lallit Anand, Chairman, Committee on Graduate Students
Department of Mechanical Engineering

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



BARKER

This page intentionally left blank

Managing Forecast Variability in a Build-to-Order Environment

by
Marshall Einhorn

Submitted to the MIT Sloan School of Management and Department of Mechanical Engineering on May 11, 2007 in Partial Fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Mechanical Engineering

ABSTRACT

In any production environment, managing demand variability is a delicate balancing act. Firms must constantly weigh potential obsolescence costs of unused inventory (should sales not materialize) against potential expedite costs or lost sales (should demand outpace available inventory). For build-to-order manufacturers such as Dell, the balancing act is even more challenging. While it offers a wide array of products, Dell does not hold its safety stock in the form of finished goods inventory. Instead, safety stock is held as parts inventory, sitting in supplier-owned supplier logistics centers. As a result, supplier stocking decisions may impact Dell's ability to respond to forecast variability. Other factors, such as globalization, product proliferation, and geo-manufacturing, all magnify the impact variability has on the forecasting process.

This thesis discusses two methods of dealing with demand variability. First, it examines the potential application of statistical modeling techniques to the part-level forecasting process. In particular, it looks at the use of time series models to forecast part-level demand. While the results did not merit a recommendation to utilize time series forecasts across the board (in lieu of the current process), certain supplemental applications of such forecasts would benefit Dell.

Second, it examines how hedging is currently utilized as a means to account for demand variability. While beneficial to Dell on the surface, a consistent hedge to the forecast is potentially detrimental to its vendor relationships. It has the direct impact of driving excess inventory onto the books of its vendors and it has the indirect impact of higher per part costs to Dell. It also exposes Dell to part shortages due to supplier decommits. To help counter these effects, the thesis identifies potential changes to the hedging process that Dell should consider.

Thesis Supervisor: David E. Hardt
Title: Professor of Mechanical Engineering and Engineering Systems

Thesis Supervisor: Roy E. Welsch
Title: Professor of Statistics and Management Science

This page intentionally left blank

I would like to thank everyone I worked with at Dell during the course of my internship. I had a tremendous experience working on this project and at Dell and it was primarily due to the great people I met. In particular, I'd like to thank my manager, Kevin Jones, for all his guidance and support during my six months in Austin. I learned a great deal from Kevin and owe my smooth transition into Dell to him. Jennifer Felch and Maggie Moore were also quite instrumental in this regard. I'd like to thank the rest of the Demand/Supply Tools and Processes team for their support and patience in answering my numerous questions throughout the internship: Sam Green, Jason Roberts, Joe Burnett, and JJ Lee. Ted Vanderlaan also taught me a lot about Dell's supply chain. I'd also like to thank Clark Ponthier and Annette Clayton for their support of the internship project and the LFM program as a whole.

I would like to thank my two faculty advisors, Roy Welsch and Dave Hardt, both for their visits to Austin (in the middle of the summer!) as well as their support from afar over the course of the internship.

I would like to acknowledge the Leaders for Manufacturing Program for its support of this work.

I would like to thank the administration of the LFM program and my LFM and Sloan classmates. My two years in Cambridge were very formative, primarily because of the relationships I built at MIT. I emerge from LFM knowing a whole lot more than when I came in.

I would like to thank my parents, Henry and Florrie, and my brothers, Ted and Jay, and their families, for their constant support and encouragement over the years.

Finally, I would like to thank my wife, Deborah, and our daughter, Elana. More than anyone else, Deb bore the burden of these last couple years and of this thesis in particular. I look forward to returning the favor on her dissertation!

This page intentionally left blank

Table of Contents

1 INTRODUCTION 9

1.1 PROBLEM STATEMENT 9

1.2 COMPANY OVERVIEW 9

1.3 DIRECT MODEL 10

1.4 SUPPLIER LOGISTICS CENTERS (SLCs) 11

1.5 VENDOR MANAGED INVENTORY 12

1.6 GLOBAL FOOTPRINT 12

1.7 PRODUCT PROLIFERATION 12

1.8 GEO-MANUFACTURING 13

1.9 THESIS FOCUS 14

2 STATISTICAL MODELING 17

2.1 SYSTEM SCHEDULING 17

2.2 PERIPHERAL SCHEDULING 18

2.3 INFORMATION MANAGEMENT SYSTEM 19

2.4 BENEFITS OF STATISTICS 19

2.4.1 IMPROVE FORECAST ACCURACY 20

2.4.2 EASE THE BURDEN 20

2.4.3 ALIGN WITH INDUSTRY BEST PRACTICE 21

2.5 APPROACH 22

2.5.1 GATHER DATA SETS 22

2.5.2 CREATE TIME SERIES MODELS 23

2.5.3 MAKE BACK-CAST PREDICTIONS 23

2.5.4 COMPARE RESULTS TO CURRENT PROCESS 24

2.6 DETAILS OF TIME SERIES MODELS 24

2.7 PROMINENT TIME SERIES MODELS 26

2.8 RESULTS 27

2.9 ANOTHER APPROACH 39

2.10 TAKEAWAYS 45

2.11 CONSIDERATIONS 47

3 SUPPLY HEDGING 49

3.1 CURRENT PROCESS 50

3.2 IMPACT OF THE HEDGE 52

3.3 IMPROVEMENT OPPORTUNITIES 54

3.4 IMPLEMENTATION CONSIDERATIONS 58

4 CONCLUSION 59

5 BIBLIOGRAPHY 61

Table of Figures

FIGURE 1 – REPLENISH-TO-PLAN VS. BUILD-TO-ORDER	11
FIGURE 2 – PRODUCT FAMILIES, 1998 TO 2006.....	13
FIGURE 3 – DEMAND PLANNING PROCESS	21
FIGURE 4 – TIME SERIES COMPONENTS	25
FIGURE 5 – 24X CD, 90-DAY FORECAST COMPARISON.....	29
FIGURE 6 – 24X CD, 60 DAY FORECAST COMPARISON	29
FIGURE 7 – 24X CD, 30 DAY FORECAST COMPARISON	30
FIGURE 8 – MODULAR BATTERY, 90 DAY FORECAST COMPARISON.....	31
FIGURE 9 – MODULAR BATTERY, 60 DAY FORECAST COMPARISON.....	31
FIGURE 10 – MODULAR BATTERY, 30 DAY FORECAST COMPARISON.....	32
FIGURE 11 – MODULAR BATTERY, MODEL VS. MASTER SCHEDULER COMPARISON.....	33
FIGURE 12 – 160G SATA HARD DRIVE, 90 DAY FORECAST COMPARISON	34
FIGURE 13 – 160G SATA HARD DRIVE, 60 DAY FORECAST COMPARISON	34
FIGURE 14 – 160G SATA HARD DRIVE, 30 DAY FORECAST COMPARISON	35
FIGURE 15 – 90 DAY FORECAST PERFORMANCE DISTRIBUTION, MODEL.....	36
FIGURE 16 – 90 DAY FORECAST PERFORMANCE DISTRIBUTION, MASTER SCHEDULER	36
FIGURE 17 – 60 DAY FORECAST PERFORMANCE DISTRIBUTION, MODEL.....	37
FIGURE 18 – 60 DAY FORECAST PERFORMANCE DISTRIBUTION, MASTER SCHEDULER	37
FIGURE 19 – 30 DAY FORECAST PERFORMANCE DISTRIBUTION, MODEL.....	38
FIGURE 20 – 30 DAY FORECAST PERFORMANCE DISTRIBUTION, MASTER SCHEDULER	38
FIGURE 21 – 90 DAY FORECAST COMPARISON, EUROPEAN FPMS, MODEL A	41
FIGURE 22 – 60 DAY FORECAST COMPARISON, EUROPEAN FPMS, MODEL A	41
FIGURE 23 – 30 DAY FORECAST COMPARISON, EUROPEAN FPMS, MODEL A	42
FIGURE 24 – 90 DAY FORECAST COMPARISON, EUROPEAN FPMS, MODEL B	43
FIGURE 25 – 60 DAY FORECAST COMPARISON, EUROPEAN FPMS, MODEL B	43
FIGURE 26 – 30 DAY FORECAST COMPARISON, EUROPEAN FPMS, MODEL B	44
FIGURE 27 – FORECAST ACCURACY METRICS SUMMARY, EUROPEAN FPMS	44
FIGURE 28 – FORECASTING LIFECYCLE VS. VOLUME.....	46
FIGURE 29 – FORECASTING VARIABILITY VS. VOLUME	47
FIGURE 30 – HIGH-LEVEL MAP OF THE FORECAST PROCESS	51
FIGURE 31 – COMPETING FORCES IN REGARDS TO HEDGING	52
FIGURE 32 – INVENTORY PULLS VS. HEDGED FORECAST	53
FIGURE 33 – SYSTEM DYNAMICS MODEL OF HEDGE	55
FIGURE 34 – SAFETY STOCK COMPARISON: HIGH VOLATILITY VS. LOW VOLATILITY.....	57

1 Introduction

1.1 Problem Statement

This thesis explores the impact of demand variability on a build-to-order production environment. Managing demand variability is critical to the success of any production operation. A supply chain that can accommodate sudden changes in customer demand is thus an essential ingredient to a firm's ability to deliver high customer satisfaction. This is particularly the case in the build-to-order world, where items are built only once an order is received. A supply chain ill-equipped to handle variations in customer demand would jeopardize the firm's ability to build its end product according to customer preferences.

An optimized approach to demand variability seeks to strike a balance between the potential costs of overstocking and lost revenue of stocking out. While overstocking threatens the product's 'freshness,' and incurs potentially significant inventory costs, stocking out weakens a company's ability to meet unexpected upsurges in demand. As this thesis will examine vis-à-vis Dell's supply chain, a firm must devise strategies for effectively dealing with these potential swings in demand.

Unlike a build-to-stock production environment, build-to-order operations do not build a set number of items each day. When a firm tailors its production schedule to meet customer orders, it places added pressure on the firm's ability to accurately forecast demand and manage any variability that may arise. With the continued trend toward the outsourcing of sub-assemblies, it puts added pressure on the parent firm to synch its entire supply chain around the build-to-order paradigm. While a firm may have been able to manufacture its own parts or sub-assemblies previously if demand spiked unexpectedly, this is no longer the case today.

In industries with rampant technology changes, accurate demand planning helps drive the company's bottom line through lower expedite costs and lower inventory holding costs. Managing forecast variability can also serve to drive an improved customer experience by stocking the shelves with the right product mix. A firm's ability to effectively handle variation in the demand planning process will determine its ability to respond to customer orders and demand trends. Ultimately, this ability to deal with demand variability will go a long way to determining the company's overall success.

1.2 Company Overview

The story of Dell Inc.'s rapid ascension to the upper echelons of the computer industry is well known. Founded in 1984 in Michael Dell's University of Texas dorm room, Dell was built on the premise that eliminating the middle man and dealing directly with the customer is better for all parties involved. Numerous books, articles, and publications discuss the dynamics of the personal computer industry in the 1980's and Dell's use of the Dell Direct Model to taking a

leading position in the industry. This introduction will focus on aspects of Dell's supply chain that directly impact the demand and supply forecasting process.

1.3 Direct Model

Dell's just-in-time, build-to-order manufacturing process is at the heart of its supply chain, as the Dell Direct Model enables consumers to configure systems to their specific needs. As a result, Dell builds its computers only once a customer order is received. Thus, Dell's facilities hold almost no finished goods inventory.

Still, Dell needs to build and ship new orders as quickly as possible to help foster a positive experience for the consumer. For individual customers, Dell competes both with other manufacturers who have followed its lead into direct sales and with traditional brick and mortar retailers. Since customers can walk into traditional stores and walk out with a computer a few minutes later, Dell needs to ship systems to the customer as quickly as possible. Although important for individual sales, rapid and on-time delivery of orders is also a key driver of customer satisfaction on the corporate side of the business (which makes up approximately 85% of revenues).

Given this juxtaposition – an environment that avoids finished goods inventory while emphasizing quick turnaround of customer orders – Dell's manufacturing system demands enough on-hand parts inventory to build incoming orders quickly. A traditional manufacturer might have a weekly or monthly order in advance from each of its customers, enabling them to forecast their parts needs more accurately and build systems on a more predictable schedule. Dell, on the other hand, must do its best to forecast expected demand and plan their supply needs according to the product mix laid out in that forecast.

Figure 1 demonstrates the dynamics of this environment. On the supply side of the equation, a replenish-to-plan approach is utilized to fill the supplier logistics center (SLC) with the parts needed to build the systems included in the sales forecast. In this sense, supply is "pushed" to the manufacturing sites based on the forecast. On the demand side of the equation, the build-to-order approach means supply is "pulled" as needed from the SLCs based on customer orders. At the dotted line, where "replenish-to-plan" meets "build-to-order" and where "push" meets "pull," lies the inherent complexity of demand and supply planning within Dell.

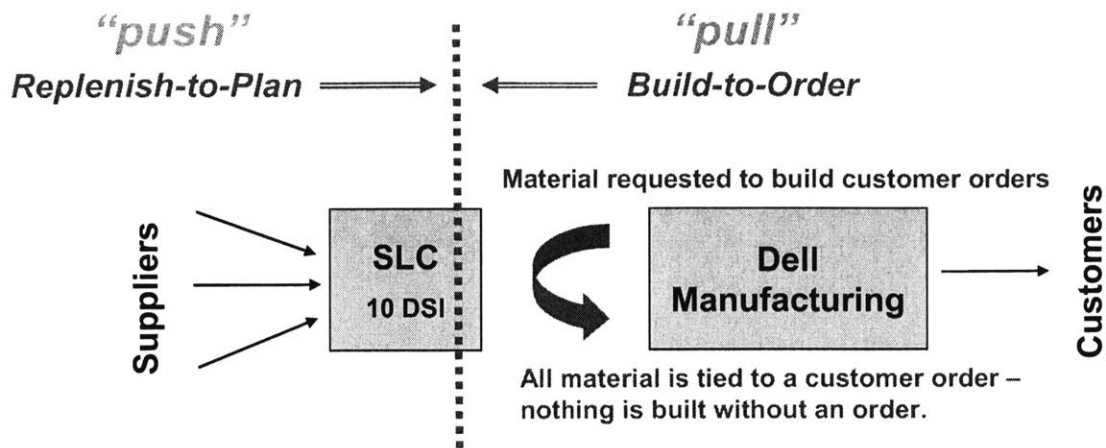


Figure 1 – Replenish-to-Plan vs. Build-to-Order¹

1.4 Supplier Logistics Centers (SLCs)

The SLC in Figure 1 represents a key cog in Dell’s supply chain. In “The World is Flat,” Thomas Friedman tracks the entire supply chain, including the unique role of Dell’s SLCs, that resulted in the delivery of his Dell Inspiron 600m notebook computer. In his words, “Surrounding every Dell factory in the world are these supplier logistics centers, owned by the different suppliers of Dell parts. These SLCs are like staging areas. If you are a Dell supplier anywhere in the world, your job is to keep your SLC full of your specific parts so they can constantly be trucked over to the Dell factory for just-in-time manufacturing.”²

Its use of SLCs enables Dell to carry virtually no parts inventory on its own financial balance sheet. In fact, with trucks arriving at each factory from its co-located SLC nearly every two hours, Dell typically has less than a day’s worth of parts inventory within the four walls of its production facilities. With no parts or finished goods to store, Dell is basically a warehouse-free operation.

The flip side of its lack of inventory and its utilization of SLCs is that Dell must rely heavily on the support of its suppliers. As one author puts it, “Dell has ripped away the psychological safety net that lots of inventory provides. Instead, it lives in a state of constant, self-imposed paranoia: It must meet demand, which is always in flux, with just the right amount of supply. If it fails, its manufacturing operations will crash within a matter of hours.”³

¹ Ponthier, Clark. “Dell Supply Chain Overview.” Presentation given to Pennsylvania State University, University Park, PA. 8 Sept. 2006. Wording adapted.

² Friedman, Thomas L. *The World is Flat: A Brief History of the Twenty-First Century*. New York: Farrar, Straus and Giroux, 2006. 516.

³ Breen, Bill. “Living in Dell Time.” *Fast Company* Nov. 2004. 27 Mar. 2007
<<http://www.fastcompany.com/magazine/88/dell.html>>.

1.5 Vendor Managed Inventory

To ensure that does not happen, suppliers are given the reigns for controlling the flow of supply into the SLCs. The vendor managed inventory (VMI) approach utilized by Dell is described as follows: “suppliers decide how much inventory to order and when to order while Dell sets target inventory levels and records suppliers’ deviations from the targets...and it uses a quarterly supplier scorecard to evaluate how well each supplier does in maintaining this target inventory in the [SLC].”⁴

To help determine how much supply to deliver to each SLC, suppliers are given Dell-generated demand forecasts on a weekly basis. Suppliers can weigh the figures they are given with the real-time demand they see in the form of the pulls from the SLCs. In this sense, the Dell supply chain provides a unique opportunity to each of its suppliers. Unlike their dealings with other manufacturers where they have to take the manufacturer at its word as far as customer demand, in Dell’s SLC-based manufacturing system, suppliers are given a direct signal of Dell customer demand.

1.6 Global Footprint

Much like the rest of its competition within the computer industry, many of Dell’s suppliers are now based in Asia. Early on, one of the primary benefits of Dell’s approach to manufacturing was its co-location with its supply base, as “Dell encouraged suppliers to locate warehouse and production facilities close to its assembly operations.”⁵ Dell was large enough that it could convince its suppliers to open new facilities near its own plants.

This was especially true with its production facilities in Austin, Texas where a significant number of electronics and technology companies had operations. Transportation lead time was minimal and as a result suppliers could react more quickly to a sudden need for a particular part. However, with the move of its supply base to Asia, transportation is now a much more significant portion of the overall lead time for a part. This is an important factor in the forecasting process, as it has become much more costly to expedite parts internationally.

1.7 Product Proliferation

As recently as 1998, Dell had 22 product families, primarily in its traditional core areas: desktops, portables (i.e. laptops), workstations, servers, and storage. By 2006, Dell had 180 product families. While some of these were in the services arena, many were in manufactured goods. Networking devices, projectors, flat-panel televisions, printers, and hand-held devices were all added to Dell’s product line in the last eight years. Appropriately, the “Computer” in “Dell Computer Corporation” was removed to reflect Dell’s move toward a full-blown

⁴ Kapuscinski, Roman, Rachel Q. Zhang, Paul Carbonneau, Robert Moore, and Bill Reeves. “Inventory Decisions in Dell’s Supply Chain.” *Interfaces* 34 (2004). 192.

⁵ Rivkin, Jan W., and Michael E. Porter. “Matching Dell.” Boston, MA: Harvard Business School, 1999. 8.

electronics provider. Figure 2 highlights the rapid expansion of Dell’s product line between 1998 and 2006.

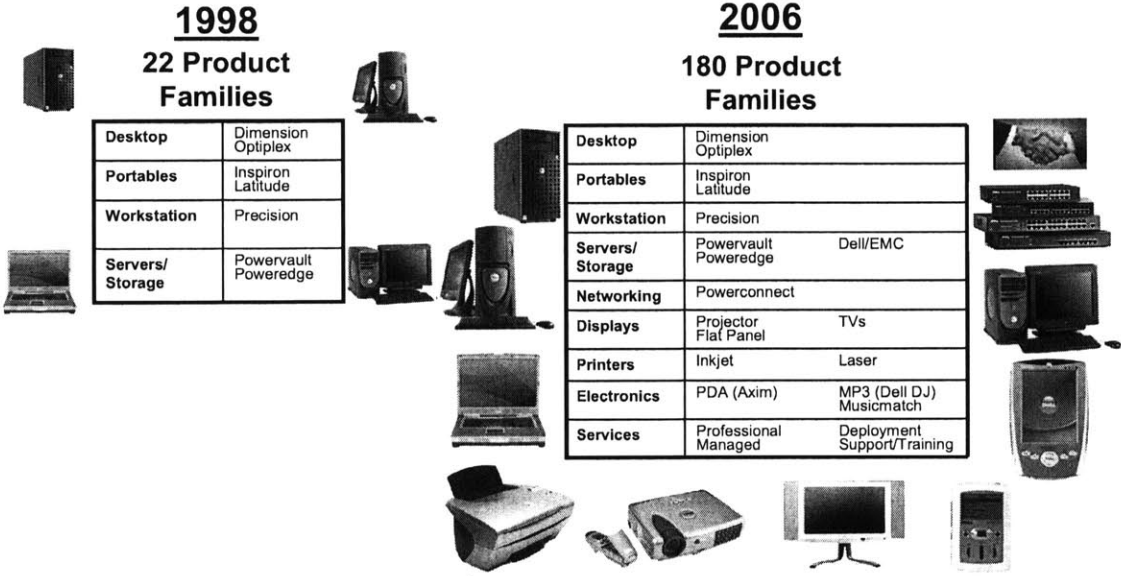


Figure 2 – Product Families, 1998 to 2006⁶

Dell’s product proliferation means more options for customers and more opportunities to increase revenue. The expanded offerings are meant to build upon its pre-existing supply chain strengths and in many ways that is the case. For example, as one of the world’s largest purchasers of glass for its monitor business, the move into televisions was a natural extension. It already possessed the supplier relationships and fulfillment infrastructure it would need for the television business. Yet, with each new product Dell adds to its portfolio, the company adds more complexity to the forecasting process. More products mean more systems and configurations to forecast, which mean more parts and materials to forecast, transport, and store in the SLCs.

1.8 Geo-Manufacturing

The final aspect of Dell’s operations that merits discussion in the context of the forecasting process is its move to geo-manufacturing in the last decade. In the past, each of Dell’s factories in the U.S. was dedicated to a specific product line. If an order came in for a desktop, it was assembled in one factory, while servers were assembled in another factory. Now, when a desktop order is placed, it is routed to one the three U.S. manufacturing facilities based on which facility is geographically closest to the shipping destination. While this results in significant savings in shipping and logistics costs, it creates more supply side costs. When demand is pooled and fulfilled at a single location, a manufacturer requires lower inventory to support the operation than it would if fulfilled out of multiple sites. Conversely, when demand is spread across multiple facilities, the amount of inventory in the entire system increases according to the

⁶ Ponthier.

number of facilities. This is known as the “square root law”, which states that the “total safety inventory required to provide a specified level of service increases by the square root of the number of locations in which it is held.”⁷

In her thesis “Multi-Site Inventory Balancing in an Extended Global Supply Chain”, Amy Reyner discusses the overall impact of globalization, product proliferation, and demand disaggregation on Dell’s supply chain:

Dell and many other industries are already seeing the effects of the challenges of increased lead times and demand disaggregation. A recent Annual State of Logistics Report (published by the Council of Supply Chain Management Professionals) stated that “As global supply chains have become longer and less predictable, companies have been carrying higher-than-ideal levels of inventory... prudent managers want to minimize inventory, but if they do that, they could be left with empty shelves.” Many speculations exist regarding Dell, specifically, as well. *Logistics* magazine states that Dell increased its days of inventory from three to four in FY 2005, and comments that “it may signal that the logistics icon is stretching its supply chain as it grows across product lines and geographic borders.” Goldman Sachs & Co. adds that Dell’s “increased size, larger international exposure, and much broader product line have reduced its nimbleness.” These analysts and many others will be watching carefully to learn how Dell and a number of other companies leap over such supply chain hurdles.⁸

1.9 Thesis Focus

This thesis examines two techniques that can be used to deal with variability when forecasting supply needs during the planning process: statistical modeling and hedging.

Section 2 looks at the use of statistical modeling techniques in the context of part-level demand. It discusses the potential benefits of a statistical approach to forecasting, attempts to create a viable model for doing so, and analyzes the model’s benefits and drawbacks. In addition to a quantitative approach to statistical modeling, Section 2 also examines more qualitatively how statistical modeling should be utilized in the supply forecasting process. With some schedulers forecasting the demand for hundreds of parts on a weekly basis, statistical modeling represents a potential way to ease the burden placed on schedulers. At the same time, statistical models cannot replace the tribal knowledge built up over years of forecasting.

The second technique discussed is the use of hedging as a means to account for forecast variability. With a globally dispersed supply chain forming the backbone of Dell’s operational capabilities, hedging is viewed as an effective way to protect against demand spikes and large orders. Section 3 provides an overview of the current hedge process and summarizes interviews

⁷ Anupindi, Ravi, Sunil Chopra, Sudhakar D. Deshmukh, Jan A. Van Mieghem, and Eitan Zemel. Managing Business Process Flows. Upper Saddle River, NJ: Prentice Hall, 1999. 146.

⁸ Reyner, Amy. Multi-Site Inventory Balancing in an Extended Global Supply Chain. Thesis. Massachusetts Institute of Technology, 2006.

with key personnel along the value chain regarding the current approach to hedging. Drawing upon these interviews, detailed recommendations for improvement opportunities are discussed.

This page intentionally left blank

2 Statistical Modeling

This section discusses the potential application of statistical modeling as a means of dealing with demand variability in a build-to-order environment. It first provides an overview of Dell's forecasting process, describing the role played by system master schedulers and peripheral master schedulers. To help frame the problem, this section discusses the organizational context of the project as well as the theory behind statistical forecasting. The section then looks at the use of time series models to forecast part-level demand. The predictions derived from these models are then compared to the actual demand encountered to determine the overall effectiveness of the statistical forecasts in comparison to the actual forecasts used. The section concludes with a discussion of the modeling results and the key takeaways from this analysis.

2.1 System Scheduling

The Demand/Supply organization within Dell America Operations is responsible for forecasting the demand for each part that goes into a Dell system. In all, this amounts to thousands of parts spread across multiple types of systems. Each week, master schedulers broken into four main forecasting teams look at the previous week's actual demand figures and the current master sales plan (MSP) to determine the forecasted demand for each system and part.

The master scheduling teams are grouped together based on the line of business they support:

- Desktops
- Portables
- Servers and storage
- Electronics and accessories

The part-level forecasting process begins with a forecast of system level demand. Each of the main lines of business has a master scheduler dedicated to forecasting demand for each of Dell's major platforms (e.g. Dimension, Latitude). These system level master schedulers base their projections on the MSPs that are received periodically throughout each quarter from the sales segments. At the aggregate level, the MSP sales projections have proven fairly accurate. Actual sales have traditionally run within a few percentage points of the number of system sales projected at the "Dimension" or "Inspiron" level, for example. The MSP is traditionally less accurate at lower levels, where it indicates the mix of each type of platform that will be sold (e.g. Dimension E521 vs. XPS710). It is inherently more difficult for Dell's sales segments to forecast the exact product mix that its customers will demand than the overall number of products.

To help narrow the gap that exists when forecasting the product mix, the Demand/Supply group maintains close ties to the product development and marketing organizations to help determine how upcoming product transitions might shift sales from one model to another. Each system master scheduler uses their knowledge of upcoming product transitions as well as potential large orders to produce a master production plan (MPP) for each platform in their domain. The weekly MPP breaks down the MSP into a weekly forecast for each system.

The forecasting process is described in more detail in 3.1.

2.2 Peripheral Scheduling

Based on the MPP, peripheral schedulers within each of the four scheduling teams forecast the part-level mix that they predict will accompany each system. This is typically accomplished by utilizing an attach rate approach to forecasting. The attach rate approach involves the use of percentages to determine the forecasted product mix among a system's different available features. In short, master schedulers attempt to predict what percentage of customers will buy a specific part when they buy a particular model number or group of models.

For example, consider a master scheduler who is responsible for forecasting the demand for hard drives that go into desktop systems. To forecast the number of 80 GB hard drives that will be sold, the master scheduler can say that three percent of Dimension E521 sales will include an 80 GB hard drive. This will then drive three percent of all forecasted Dimension E521's to the 80 GB hard drive, regardless of the actual number of E521's the system scheduler forecasts. The attach rate in this case is three.

In the above example, the master scheduler chose to attach the 80 GB hard drive to a low level of the product hierarchy (Dimension E521). The master scheduler could also have chosen to attach the hard drive to a higher level of the product hierarchy. For example, the master scheduler may choose to forecast that three percent of *all* forecasted Dimensions, regardless of the model number, would have an 80 GB hard drive. Even though the hard drive is attached to a different level of the product hierarchy, the attach rate concept remains the same. Only in this case, instead of determining the number of 80 GB hard drives based on the number of E521's forecasted, it is based on the total number of Dimensions forecasted.

The attach rate used is based on many factors. The master schedulers will work with the following organizations to help determine the attach rates for their parts:

- Dell's internal strategy organization (the COC) – How will upcoming promotions impact the attach rate for specific parts?
- Product engineering – How will upcoming system and part transitions effect the attach rates?
- Procurement – Are there any supply issues that might impact part availability?

So, while a part may have a three percent attach rate one week, it is entirely possible that it will have a six percent attach rate the next.

A peripheral master scheduler is typically responsible for less than five commodities (groups of like parts), yet individual master schedulers can have over 500 system-part combinations to forecast each week. For example, the master scheduler responsible for desktop graphics cards – considered a single commodity – may have 30 different graphics cards under their domain. True to Dell's core mission of allowing the customer to configure their system according to their own

needs, each of those 30 graphics cards may be available in 20 different desktop platforms. As a result, the master scheduler may have 600 system-part combinations to forecast each week.

This is an example of the Dell Direct Model having a significant impact on the forecasting process. A basic principle of forecasting is that the fewer individual items being forecasted, the more accurate the resulting forecast. If each of those 30 graphics card were available in ten different platforms as opposed to 20, it would result in a more accurate forecast. More specifically, the forecast for the *sum* of all graphics cards will be more accurate than the forecast for each of the graphics card types individually. Yet, obtaining the correct forecast for the *individual* types of graphics cards is also a key outcome of the forecasting process. In Dell's case, by forecasting an attach rate to lower levels of the product hierarchy, the sum of the forecasts against lower levels typically produce a more accurate summed forecast than a forecast based on attach rates against the higher product level.

2.3 Information Management System

The information management system that manages the Demand/Supply forecasting process is critical to the success of the overall process. It provides an interface for the master schedulers to both view the projected system demand and input the forecasted part-level demand. Every week, it produces a demand forecast for every part used in each of Dell's three global regions. This forecast is in turn used by the procurement team to drive the flow of supplies into the SLCs. It is also used by production control to give the production team a sense of near term forecasted demand and the impact on work shifts in the factories.

A project team within the Demand/Supply Tools and Processes team is driving changes into the current information system, and the thesis project was housed within this team. A benefit being sought of the new system design is that it creates data analysis and statistical modeling opportunities that were previously unavailable to master schedulers within the forecasting system.

2.4 Benefits of Statistics

Before discussing the feasibility of adding statistical modeling to the redesigned system and the details of which modeling techniques to make available, it is important to understand why it should be added. According to Douglas Montgomery, "Statistics is a collection of techniques useful for making decisions about a process or population based on an analysis of the information contained in a sample from that population."⁹ There are three main reasons a forecasting organization might consider adding a statistical component to the forecasting process.

⁹ Montgomery, Douglas C. Introduction to Statistical Quality Control. 4th ed. New York, NY: John Wiley & Sons, Inc., 2001. 37.

2.4.1 Improve Forecast Accuracy

A forecast is only as good as its eventual accuracy. At Dell, forecast accuracy is determined by the following formula:

$$\text{Forecast Accuracy} = \text{Actual Demand} / \text{Forecasted Demand}$$

Equation 1 – Forecast Accuracy

A forecast that predicts demand of 100,000 units in a particular month sets the expectation for the entire supply chain. Everyone from third tier suppliers to final assembly factories will plan their activities accordingly. If demand actually comes in at 125,000 units, there are 25,000 orders that the supply chain has to scramble to fill. The forecast accuracy in this case is 125%. While there may be enough buffer inventory in the SLC to meet this unforeseen demand, it leaves the manufacturer at risk until that buffer is refilled. Likewise, there may be enough capacity in the factory to handle the additional orders, but production control may have to schedule overtime instead of a standard shift to handle the extra demand. A more accurate forecast may have enabled production control to avoid costly overtime.

Or, consider the scenario where forecasted demand is 100,000 units and actual demand comes in at 75,000 units. In this case, forecast accuracy is 75%. Material has been put into the supply chain to meet 25,000 units of demand that did not materialize. Schedules have been set on the final production lines to meet the full 100,000 units. Some companies would still build to the full 100,000 units, hoping to sell the excess at a later date. For Dell, carrying finished goods inventory is not an option. The supply associated with those excess units will remain in the supply hub until orders come through that need it. And in an industry where “many components lose .5 to 2.0 percent of their value per week”¹⁰, having 25,000 extra units represents a significant liability.

Because the forecast drives so many aspects of the process, both internally and externally, forecast accuracy is one of the key metrics that is used to determine the effectiveness of the master scheduling team. If statistical modeling can improve upon the current forecast accuracy performance, it would be a simple decision to implement it into the forecasting process.

2.4.2 Ease the Burden

As discussed above, the process of breaking down the master sales plan and converting it into a part-level production plan is time intensive. Peripheral master schedulers have hundreds of system-part combinations to forecast each week and thus hundreds of attach rates to enter into the information system each week. Decreasing the amount of time master schedulers spend entering figures into the system would give them more time to focus on other aspects of their jobs (e.g. collaborating with counterparts in other parts of the value chain, implementing process improvements).

¹⁰ Kapuscinski, 191.

There are two main opportunities for statistical modeling to decrease the amount of time spent on the forecasting process. One, the system could utilize past data to automatically suggest forecast figures. Master schedulers could then determine which system-part combinations to focus their attention. Two, the inclusion of statistical modeling in the tool will lessen the need for individual models to be stored on master schedulers' computers. With more information and models within the system, it will give master schedulers access both to their own models and others' models as well. This will result in less time spent creating their own models or asking for others' and more time spent analyzing their commodities and collaborating with others.

2.4.3 Align with Industry Best Practice

Other companies are utilizing statistical modeling to help drive their demand planning organizations. In their book "Demand Management Best Practices: Process, Principles and Collaboration", Crum and Palmatier identify the ideal process that an organization can utilize in creating a demand plan:

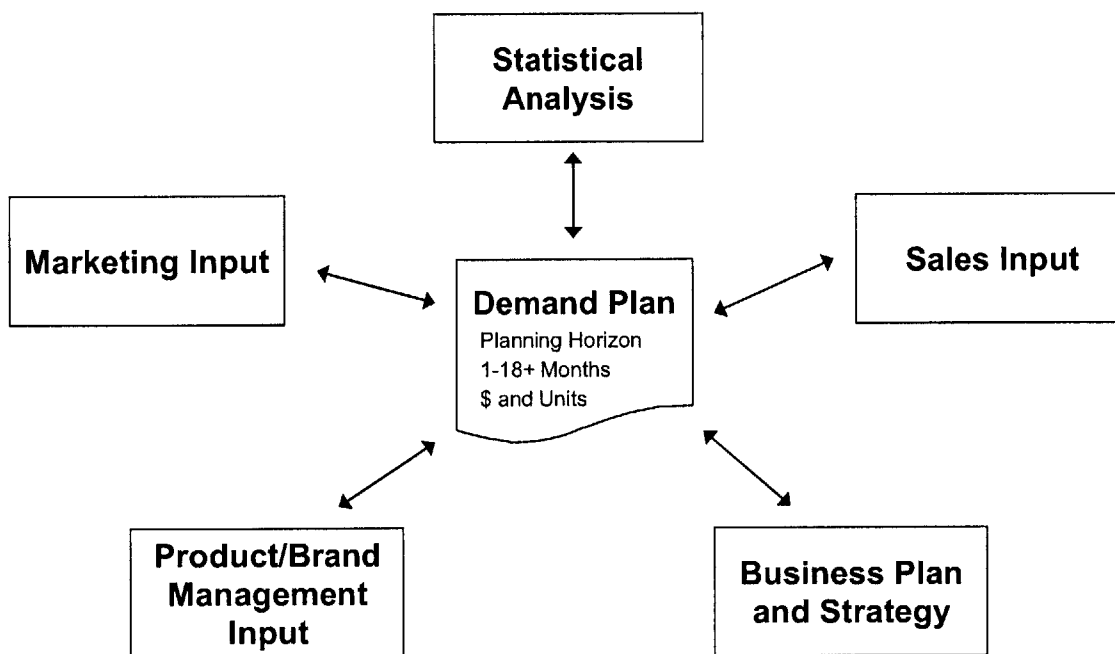


Figure 3 – Demand Planning Process ¹¹

Of the five boxes feeding into the demand plan in the above diagram, statistical analysis represents the biggest opportunity for the Demand/Supply group to level itself with demand planning best practice.

¹¹ Crum, Colleen, with George E. Palmatier. Demand Management Best Practices: Process, Principles and Collaboration. Boca Raton, FL: J. Ross, Inc., 2003. 29.

2.5 Approach

In order to determine the potential impact of statistical modeling on overall forecast performance, the study focused on the forecast accuracy metric. If the forecast accuracy of the statistical models put together could consistently perform as well as or outperform the current approach to forecasting, it would satisfy the first criteria discussed above.

In order to make this determination, the following process was used:

1. Gather Data Sets
2. Create Time Series Models
3. Make Back-cast Predictions
4. Compare Results to Current Process

2.5.1 Gather Data Sets

At the start of the study, a significant amount of time was spent determining which data sets were appropriate for the modeling effort. The following characteristics were considered.

- At least 12 months of data – As discussed above, Dell’s sales are cyclical in nature. Some of these cycles are three months long, driven by the quarterly pressure placed on any publicly traded company. Other cycles are annual in nature (e.g. holiday shopping season, federal buying season, back-to-school shopping season). In order to capture these annual demand patterns, the study chose to focus on part data that went back at least 12 months.
- One-to-one relationship with a system – For a part that is available in multiple systems, it is difficult to determine from a single part-level volume which system is driving the demand. If demand for a part was 10,000 units one month and it attached to eight different platforms, it would be difficult to pinpoint the drivers behind those 10,000 units. Since a potential direction of the study involved drawing correlations between parts and specific systems, the study attempted to focus on those parts that only attached to one system.
- As true a demand signal as possible – In an ideal world, the part-level data would represent customer demand as opposed to customer sales. However, capturing a true demand signal is impossible in this environment. Whether a customer buys a system off www.dell.com or through the customer call center, there is no way to determine the system and components they truly wanted to buy when they first started the sales process. Pricing, promotions, stockouts, and up-selling all represent potential ways in which the customer’s true demand signal gets diluted during the sales process. Initially, the study hoped to focus on those parts that were least exposed to these factors.

2.5.2 Create Time Series Models

The analysis focused on a particular type of statistical modeling technique known as time series modeling. A time series model is a projection derived from historical data that is broken out into regular intervals (e.g. days, weeks, months). Because the projections are based on regularly spaced time buckets, the model predicts the future in those same time buckets. Examples of time series models include:

- Daily rainfall predictions based on past daily rainfall measurements
- Monthly shopping mall revenue based on past monthly revenue
- Annual automobile sales projections based on past sales performance

The advantages to time series forecasting stem primarily from the fact that time series models rely on nothing other than the specific historical data itself. Whereas one could make a city's rainfall predictions based on other factors such as macro indicators of global weather patterns or the amount of snowfall the previous year in nearby regions, a time series model relies solely on previous rainfall data. In this sense, the time series approach lacks any bias to other factors. It also makes it an effective way to forecast a large number of items. Unlike regression analysis, there is no need to gather other sets of data to feed into the model.¹²

Within the realm of time series forecasting, there are numerous models that emphasize different aspects of the historical data to predict future values. A more detailed discussion of these models is included below.

2.5.3 Make Back-cast Predictions

By examining historical demand patterns we can make “back-cast” predictions using time series forecasting models. The predictions derived from these models can then be compared to the actual demand encountered to determine the overall effectiveness of the statistical forecast.

To understand the back-casting approach, consider the case of a part that has 36 months of historical data. To evaluate the effectiveness of the model's predictions, we decide to compare the most recent 12 months of data to the model's predictions for those 12 months. However, the predictions that are most relevant to the study are those made 30 days, 60 days, and 90 days in advance. These are the predictions that drive the 30, 60, and 90 forecast accuracy metrics that are used to determine the effectiveness of the actual forecast. Therefore, if we want to create a full 12 months of 30, 60, and 90 day forecasts, we need to create a series of rolling models based on a data set that grows by one month with each new prediction. The following list outlines the steps needed to create 30, 60, and 90 forecasts for months 25 through 36 in the 36 month data set.

1. Remove months 22 through 36 of data and create a model based off the first 21 months of data. Assign the prediction for month 25 as the 90 day forecast for month 25.

¹² Crum, Chapters 4 and 8.

2. Add month 22's sales data to the data set and create a model based off the first 22 months of data. Assign the prediction for month 25 as the 60 day forecast for month 25. Assign the prediction for month 26 as the 90 day forecast for month 26.
3. Add month 23's sales data to the data set and create a model based off the first 23 months of data. Assign the prediction for month 25 as the 30 day forecast for month 25. Assign the prediction for month 26 as the 60 day forecast for month 26. Assign the prediction for month 26 as the 90 day forecast for month 26. There is now a full set of 30, 60, and 90 forecasts for month 25.
4. Repeat this process for each successive month of data.

With different time series models from which to choose, it is necessary to select one that closely predicts a particular part's demand pattern. For the data sets discussed below, this comparison was easy to make simply by viewing the forecast charts of the various models, as many of the models were clearly an inappropriate fit. In a handful of cases, two models closely tracked the actual demand pattern exhibited based on the shape of their forecast charts. In those cases, data was pulled from both models and compared to the actual demand as described in section 2.5.4 to determine the closest fit.

2.5.4 Compare Results to Current Process

The final step in the process is to determine the effectiveness of the model. There are two aspects of this determination. First, the model's predictions should be compared to actual sales in a similar manner using the same forecast performance metrics utilized by the Demand/Supply group. With the model's 30, 60, and 90 day forecasts in hand this is simply a matter of using the formula identified in Equation 1 above. This will give a sense of how the model performed relative to actual demand.

The model's performance metrics should also be compared to the performance metrics of the actual forecast that was used. This will give a sense of how the model performed relative to the current forecasting methodology.

2.6 Details of Time Series Models

As described by Hanke and Wichern in "Business Forecasting", there are four components to any time series model:¹³

- Trend (T)
- Seasonal (S)
- Cyclical (C)
- Irregular (I)

These components are represented graphically in Figure 4:

¹³ Hanke, John E., and Dean W. Wichern. Business Forecasting. 8th ed. Upper Saddle River, NJ: Pearson Prentice Hall, 2005. 158-159.

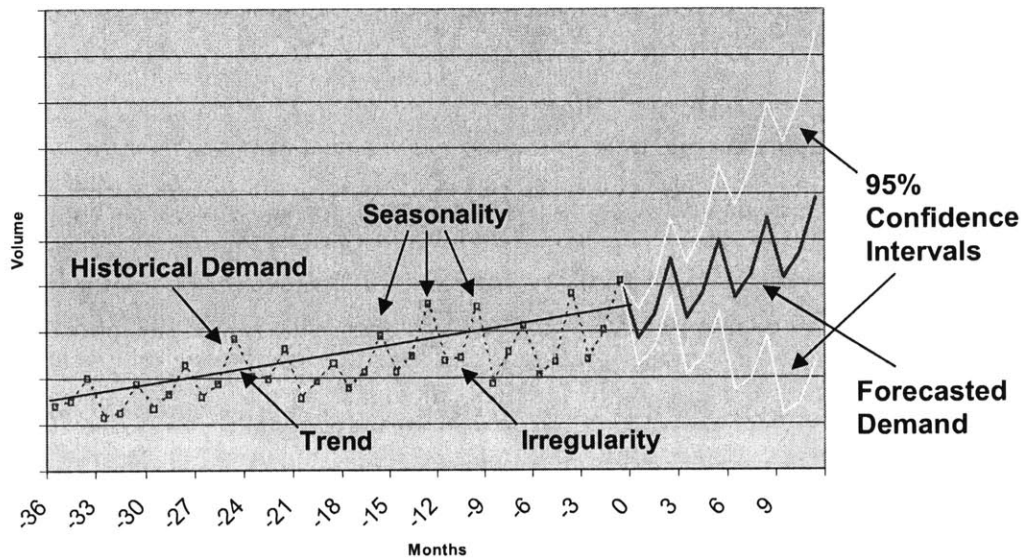


Figure 4 – Time Series Components

The *trend* component represents the underlying growth or decline exhibited by the data set. A gradual change over time would be captured in the trend variable. Hanke and Wichern reference consistent population changes and inflation as two examples of the trend component. In the world of personal computers, the general shift toward portables and away from desktops is an example of trend. Another example is the move away from cathode ray tube monitors and toward flat panel monitors.

The *seasonal* component represents repeatable patterns in the data set. These are typically driven by the weather or “calendar-related events such as school vacations and national holidays”.¹⁴ In the case of most publicly traded companies, sales will follow a hockey stick pattern within a quarter, with sales spiking toward the end of each quarter. Likewise, other sales patterns contribute to the seasonal component of Dell’s time series models. The buying season between Thanksgiving and Christmas, the United States federal government buying season in September, and the back-to-school buying season in August are all examples of the seasonal nature of the personal computer industry.

The *cyclical* component represents fluctuations that are greater than one year in length. These fluctuations are often difficult to extract from a time series and are typically included in the trend portion of the decomposition. An example from the personal computer world is the dot com boom and bust of the late 1990’s. In recent years, the emergence of new technologies and startup companies under the umbrella of “Web 2.0” could signal the continuation of this same cyclical component.

¹⁴ Hanke, 160.

The final piece of the time series puzzle is the *irregular* component. It represents changes in the data that are unpredictable and random.

With the components of a time series now identified, we can turn our attention to the relationship between the components in predicting the value of the data set. Hanke and Wichern describe it as follows. “To study the components of a time series, the analyst must consider how the components relate to the original series. This task is accomplished by specifying a *model* (mathematical relationship) that expresses the time-series variable Y in terms of the components T, C, S, and I.”¹⁵

Hanke and Wichern go on to identify the difference between an additive components model and a multiplicative model. An additive model considers Y as a sum of the above components while a multiplicative model considers it as a product of the components.

- Time series additive decomposition: $Y_t = T_t + S_t + I_t$
- Time series multiplicative decomposition: $Y_t = T_t \times S_t \times I_t$

Hanke and Wichern summarize the reasons for decomposing time series models:

One reason for decomposing a time series is to isolate and examine the components of the series. After the analyst is able to look at the trend, seasonal, cyclical, and irregular components of a series one at a time, insights into the pattern in the original data values may be gained. Also, once the components have been isolated, they may be recombined or synthesized to produce forecasts of future values of the time series.¹⁶

2.7 Prominent Time Series Models

The following discussion of specific models is meant to give a sense of the different techniques that have been developed to make time series forecasts.

Exponential Smoothing involves computing a moving average of the observed value, “where the current and immediately preceding (‘younger’) observations are assigned greater weight than the respective older observations. Simple exponential smoothing accomplishes exactly such weighting, where exponentially smaller weights are assigned to older observations.”¹⁷ In order to give more weight to the most recent observation, the following equation is used:

$$S_t = \alpha * X_t + (1-\alpha) * S_{t-1}$$

Equation 2 – Simple Exponential Smoothing

¹⁵ Hanke, 160.

¹⁶ Hanke, 172.

¹⁷ “Time Series Analysis.” StatSoft, Inc. <<http://www.statsoft.com/textbook/sttimser.html#exponential>>.

S_t is the smoothed prediction for the upcoming period, α is the smoothing constant between 0 and 1, X_t is the most recent observation, and S_{t-1} is the smoothed prediction for the previous period. Clearly, the value selected for α is key. If one wants the most recent observation to strongly influence the forecast for the next period, α should be close to one. If one wants the previous prediction to be more influential, α should be closer to zero.

Numerous variations of the exponential smoothing method of time series forecasting exist. *Holt's linear exponential smoothing* (a.k.a. double exponential smoothing) "allows for evolving local linear trends in a time series"¹⁸, while the *Winters method* takes Holt's method one step further by adding a seasonality estimate to the forecast.

Hanke and Wichern summarize the benefits of exponential smoothing methods, "Exponential smoothing is a popular technique for short-run forecasting. Its major advantages are low cost and simplicity. When forecasts are needed for inventory systems containing thousands of items, smoothing methods are often the only acceptable approach."¹⁹

Autoregressive integrated moving average (ARIMA) represent another category of time series forecasting methods. Also known as *Box-Jenkins*, this model "uses an iterative approach of identifying a possible model from a general class of models."²⁰

Seasonal ARIMA models are a subset of ARIMA models that specifically target data sets that are particularly seasonal in nature. Seasonal ARIMA models are typically listed with a notation of ARIMA (p, d, q) (P, D, Q)_S where:

- p = number of regular autoregressive terms
- d = number of regular differences
- q = number of regular moving average terms
- P = number of seasonal autoregressive terms
- D = number of seasonal differences
- Q = number of seasonal moving average terms
- S = length of seasonal period

For a detailed discussion of the details of the ARIMA and seasonal ARIMA models, consult a forecasting textbook, such as Hanke and Wichern.

2.8 Results

This section will discuss various part-level back-casting models created during the course of the internship and attempt to draw conclusions based on the results.

The following example is of a low-end optical drive (24X CD drive). The part is used solely in desktop systems sold to business customers. A Seasonal ARIMA (1,0,0;1,0,1)₃ model was used

¹⁸ Hanke, 121.

¹⁹ Hanke, 129.

²⁰ Hanke, 381.

to create the time series predictions. In this case, the model had mixed results. Using the forecast accuracy formula identified in Equation 1, we can compare the performance of the model to the performance of the actual forecast that was created by the master scheduler.

In making this comparison, it is important to understand the method used to compare the two forecast accuracy figures. The goal of any forecast is to result in 100% accuracy. If forecast A results in 75% accuracy and forecast B results in 80% accuracy, it can be said that forecast B was 5 percentage points more accurate than forecast A.

The answer is less straightforward for situations where forecast A is 75% accurate and forecast B is 120% accurate. In this case, forecast B is still 5 percentage points more accurate than forecast A because B is 5 percentage points closer to 100% than A. Yet, there are other considerations when comparing two forecasts that fall on either side of the 100% mark. In this example, forecast B under-forecasted the actual demand. This could lead to poor customer experience in the form of stockouts and delivery delays and increased supply chain costs in the form of expedites. Similarly, forecast A's over-forecast could lead to additional costs as well. With excess supply in the supply chain, inventory holding costs increase in addition to exposure to obsolescence costs.

For the purposes of this analysis, we will compare under-forecasts and over-forecasts on the same scale – relative to 100% accuracy. Other factors, such as supply chain costs and customer experience, which are impacted by an under-forecast or over-forecasted will be discussed in Section 3.

For the 24X CD drive, the model predicted August sales 6.79 percentage points more accurately than the master scheduler at 60 days (78.71% accuracy vs. 71.92%) and 6.66 points more accurately at 30 days (81.69% vs. 75.03%). For September, the accuracy improvements were negligible, with 0.69, 1.74, and 0.77 percentage point improvements for the 90, 60, and 30 day forecasts, respectively.

The model significantly underperforms the master scheduler forecast in October and November, with 90-60-30 performance declines of 16.02, 13.59, and 16.74 percentage points in October and 90-60 declines of 40.5 and 6.78 percentage points in November. For the two months, only the 30 day November forecast was an improvement in the model, but just a 1.78 percentage point improvement at that. For these two months, the model picked up the seasonal nature of the time series and correctly predicted a peak in October followed by a dip in November. It even picked up on the general trend downward for the part (see the decreasing quarterly peaks in January, April, and July). But it did not weight this decline enough when predicting the values for October and November.

Figure 5, Figure 6, and Figure 7 compare the actual demand to the 90, 60, and 30 day forecasts of the master scheduler and the model. In each figure, there is one-month gap in the master scheduler forecast due to missing data in the information system. This gap did not have an impact on the model forecast and the model forecast was not compared to the master scheduler forecast for this particular month.




Figure 5 – 24X CD, 90-Day Forecast Comparison

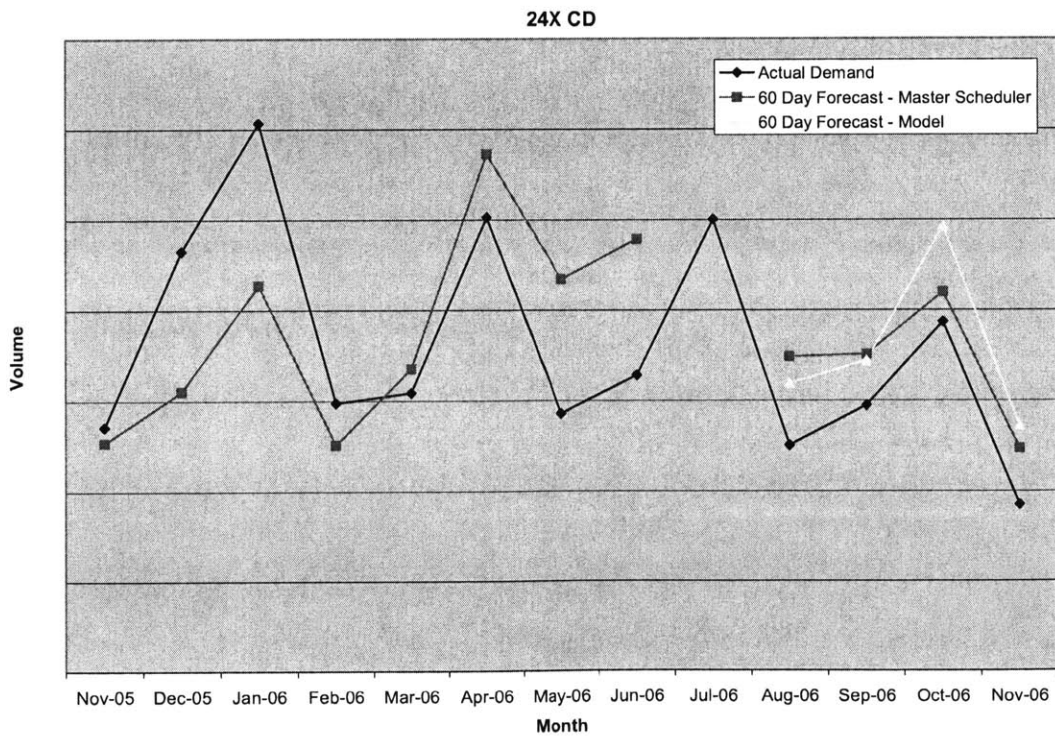


Figure 6 – 24X CD, 60 Day Forecast Comparison




Figure 7 – 24X CD, 30 Day Forecast Comparison

The next example is of a modular laptop battery. While every laptop automatically comes with one battery, a modular battery refers to a second battery that can be purchased and used in the modular bay of the device. Unlike most of the components that make up Dell’s systems, there have been relatively few technological changes to the modular battery. As such, historical sales data dates back to 2003.

The following three charts (Figure 8, Figure 9, and Figure 10) compare actual modular battery sales to the 90, 60, and 30 day forecasts of the master scheduler and the model. There is a general upward trend in modular battery sales between March 2003 and July 2005. Aside from a large spike in July 2006, sales appear to level off after July 2005.

A Seasonal ARIMA (1,0,1;1,0,1)₃ model was used to make the predictions.




Figure 8 – Modular Battery, 90 Day Forecast Comparison

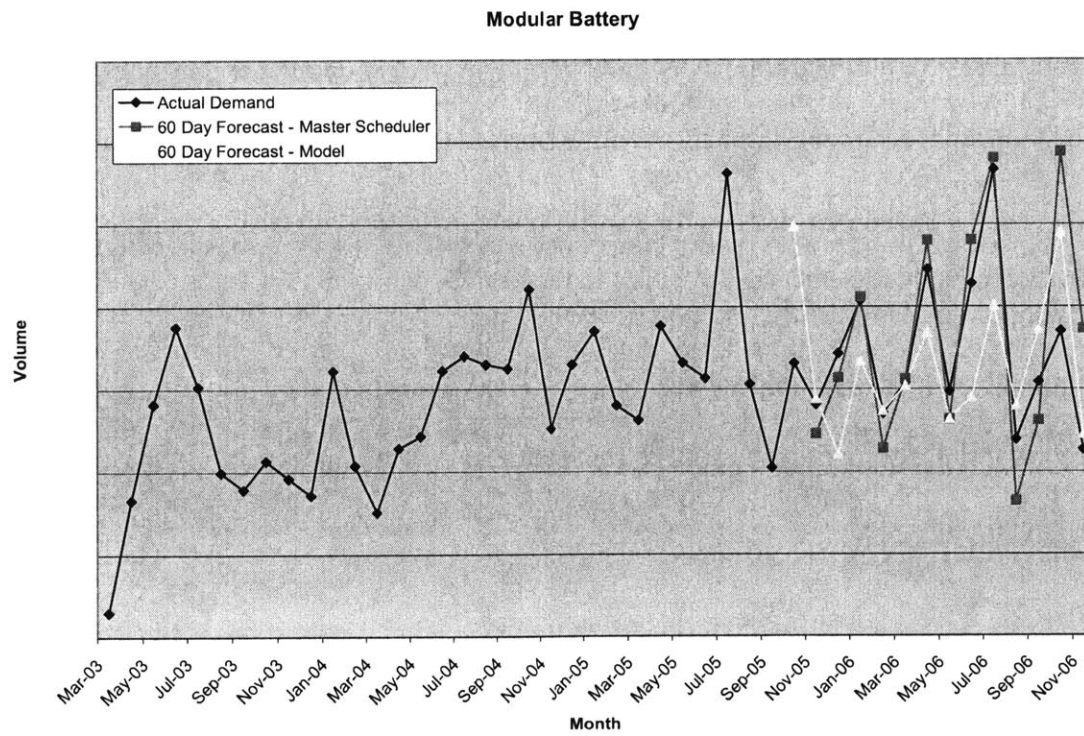


Figure 9 – Modular Battery, 60 Day Forecast Comparison




Figure 10 – Modular Battery, 30 Day Forecast Comparison

The model results were mixed for the modular battery. Figure 11 shows the percentage point differences between the model forecast and the master scheduler forecast. Values that are not highlighted indicate months where the model outperformed the master scheduler. Values that are highlighted indicate months where the master scheduler outperformed the model. There were 18 forecast periods where the model outperformed the master scheduler and 21 where the master scheduler outperformed the model. The average improvement was 14.53 percentage points, the average decline was 19.27 percentage points, and the overall average change was a decline of 8.27 percentage points.

In particular, the model had difficulty predicting the sales volumes for the December 2005 holiday buying season. It also had difficulty predicting the sales spikes in June 2006 and July 2006. All three months were missed by a significant amount and all three were missed by an under-forecast.

Month-Year	Model vs. Master Scheduler		
	90 Day	60 Day	30 Day
November-05	-10.33%	12.22%	-3.09%
December-05	-11.99%	-45.64%	-51.57%
January-06	-8.78%	-20.91%	-1.10%
February-06	-3.61%	17.53%	-7.90%
March-06	-16.84%	0.32%	-3.05%
April-06	-9.91%	-13.55%	-25.44%
May-06	9.30%	0.32%	40.30%
June-06	-36.24%	-36.57%	-23.18%
July-06	-54.41%	-38.75%	-40.77%
August-06	26.46%	30.76%	15.39%
September-06	7.49%	0.84%	15.38%
October-06	4.10%	12.55%	41.10%
November-06	3.61%	31.56%	18.70%

Figure 11 – Modular Battery, Model vs. Master Scheduler Comparison

The following example is of the 160G SATA hard drive, a high volume hard drive that customers could add to a desktop system. Demand for the 160G SATA hard drive followed an upward trend from November 2004, peaking in December 2005, before tapering off through April 2006. A Seasonal ARIMA (1,0,1;1,0,1)₃ model was used to predict sales figures for January through April of 2006.

The two 90 day forecasts produced by the model (Figure 12) do not perform as well as the master scheduler forecasts for the same time periods. In particular, the model missed significantly for the 90 day forecast of April 2006, over-forecasting sales at 49.46% accuracy versus 68.62% for the master scheduler. However, for the 60 and 30 day forecasts (Figure 13 and Figure 14) the model improved five out of seven opportunities.

An interesting month to examine is April 2006. As mentioned above, the 90 day forecast of the model underperformed the master scheduler forecast. However, the 60 and 30 day forecasts both improved upon the master scheduler forecast. If we take the actual sales to be X, the 90 day forecasts for model and the master scheduler were 1.45X and 2.02X, respectively. From there, the two forecasts go in opposite directions. The master scheduler forecast assumes that demand will continue to rise, making 60 and 30 day predictions of 2.51X and 2.95X. This is a logical assumption, given that April is the end of Dell's first quarter and the previous end of quarter numbers were strong (in January). However, the model picks up on a downward trend in the sales figures and adjusts its 60 and 30 day predictions to 1.77X and 1.27X. The model ultimately ends up at being 78.86% accurate at 30 days, which is a significant improvement from the model's 90 day starting point of 49.46%.




Figure 12 – 160G SATA Hard Drive, 90 Day Forecast Comparison




Figure 13 – 160G SATA Hard Drive, 60 Day Forecast Comparison




Figure 14 – 160G SATA Hard Drive, 30 Day Forecast Comparison

The three parts discussed thus far are typical of the complete set of parts that were analyzed. In general, the models exhibited pockets of improvement in concert with pockets of degradation. Of the 160 forecasts made, 66 resulted in improvements to the master scheduler forecast while 94 resulted in declines relative to the master scheduler forecast. The following charts show the distributions of 90 day forecast accuracy metrics (Figure 15 and Figure 16), 60 day forecast accuracy metrics (Figure 17 and Figure 18), and 30 day forecast metrics (Figure 19 and Figure 20) for the 11 parts included in this analysis for both the model predictions and the master scheduler forecasts.

90 Day Performance Distribution - Model




Figure 15 – 90 Day Forecast Performance Distribution, Model

90 Day Performance Distribution - Master Scheduler




Figure 16 – 90 Day Forecast Performance Distribution, Master Scheduler

60 Day Performance Distribution - Model

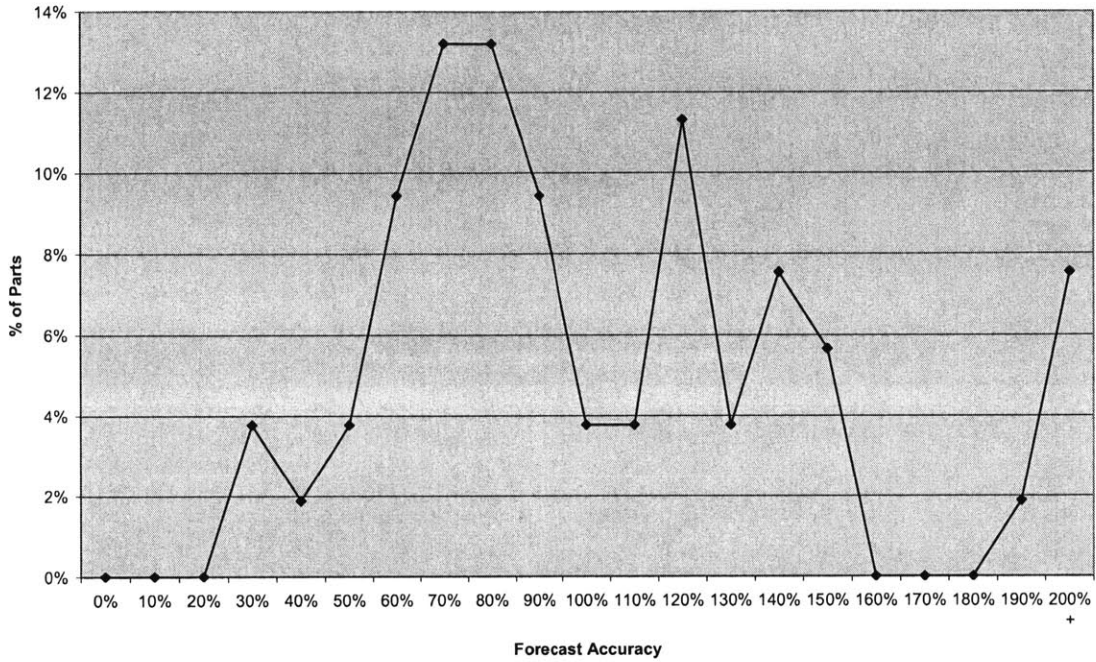


Figure 17 – 60 Day Forecast Performance Distribution, Model

60 Day Performance Distribution - Master Scheduler




Figure 18 – 60 Day Forecast Performance Distribution, Master Scheduler

30 Day Performance Distribution - Model




Figure 19 – 30 Day Forecast Performance Distribution, Model

30 Day Performance Distribution - Master Scheduler




Figure 20 – 30 Day Forecast Performance Distribution, Master Scheduler

The above analysis compared the standard attach rate method used by master schedulers to a statistical approach using a pure aggregate part demand (not as an attach percentage). As discussed above, the master schedulers take an attach rate approach whereby they assign an attach rate percentage for each part to the systems they attach to. Multiplying these attach rates by the number of predicted system sales and adding up like parts, the information system can determine the aggregate number of parts that are forecasted to be sold in a given time period.

The analysis attempted to determine if a time series approach would produce more effective results. In other words:

- Can a part's aggregate demand data be used to predict the aggregate demand for the same part in the future?
- Can time series modeling be used to replace attach rate forecasting in certain situations?

Based on the time series forecasts produced as part of this research, the answer is that it cannot replace attach rate forecasting. The subtle changes caused by promotions, demand shaping, pricing, supply disruptions, and the introduction of new products cannot be replicated in a time series approach to forecasting at the part level.

A more fitting use of this time series approach is as a tool for master schedulers to create a quick snapshot of the potential part-level demand forecast. Currently, master schedulers utilize various reports (such as the "trend report") to obtain a statistical view of future demand. The trend report takes the monthly sales to date and extrapolates it out to the entire month. A more robust approach could involve the use of time series models as a baseline against which master schedulers can compare their forecasts. It would be an effective way for master schedulers to view aggregated part-level data for all their parts against a statistical forecast at the same aggregate level. If their forecast falls outside a certain boundary, the tool could be designed to highlight the part to the scheduler. They may choose to keep their existing forecast, but at least they will have a quick litmus test for comparison's sake.

2.9 Another Approach

Another aspect of the analysis involved the creation of time series models at a higher level in the product hierarchy, using flat panel monitors (FPM) for the European market as a case study. Dell's European FPMs are currently forecasted at a higher level than the parts discussed so far. In essence, European FPMs are treated as their own line of business, as opposed to a peripheral part that is attached to systems. As its own line of business, the business creates a top line forecast for all FPM sales. This single sales estimate is then broken down into specific estimates for each available model (e.g. 17-inch, 19-inch) using a product mix that assigns a percentage to each of the models. The product mix evolves each month based on past monitor trends, master scheduler analysis of upcoming product information (e.g. promotions, new products), and the overall volume forecast for desktop systems. In essence, they effectively create an attach rate at the desktop level, then split it out to specific monitor models based on this analysis.

To determine the effectiveness of time series modeling in regard to European FPMs, the back-casting approach discussed above was utilized. Using historical aggregate demand dating back to February 2005, two different time series models were developed to predict FPM sales for December 2006 to September 2007.

Model A was developed using a Seasonal ARIMA (1,0,1;1,0,1)₃ model and model B was developed using a Seasonal ARIMA (1,0,1;1,0,1)₁₂ model. Note that the second model was created using a 12 month seasonality period, compared to a three month seasonality period for the first model. This will put a greater emphasis on the annual seasonality exhibited versus the quarterly seasonality. Both models exhibited statistically accurate results, as indicated by the mean absolute percent error (MAPE). MAPE is calculated according to the following formula:

$$MAPE = \frac{1}{n} \sum_{t=1}^n \left| \frac{P_n - A_n}{P_n} \right|$$

Equation 3 – Mean Absolute Percent Error

In the above equation, n is the number of data points in the time series, P_n is the predicted value, and A_n is the actual value. For model A, the 90, 60, and 30 day MAPE was .185, .129, and .112, respectively. For model B, the MAPE was .096, .068, and .052.

The following charts (Figure 21, Figure 22, and Figure 23) show the 90, 60, and 30 day forecast comparisons for model A and the master scheduler. Aside from the initial two months in which the model underperformed the master scheduler, model A outperforms the master scheduler forecast for all but three forecasts (each less than a 4% differential).




Figure 21 – 90 Day Forecast Comparison, European FPMs, Model A




Figure 22 – 60 Day Forecast Comparison, European FPMs, Model A




Figure 23 – 30 Day Forecast Comparison, European FPMs, Model A

The following charts (Figure 24, Figure 25, and Figure 26) show the 90, 60, and 30 day forecast comparisons for model B and the master scheduler. Model B proves to be even more effective than model A in predicting FPM sales. Here, the initial learning period is not as disruptive and over the course of the final nine months of the forecast, there are no 90, 60, or 30 day forecasts that do not outperform the master scheduler forecast. Over the entire period, the average 90, 60, and 30 day forecast accuracy is improved by 11.76 percentage points, 12.59 percentage points, and 11.48 percentage points.




Figure 24 – 90 Day Forecast Comparison, European FPMs, Model B




Figure 25 – 60 Day Forecast Comparison, European FPMs, Model B

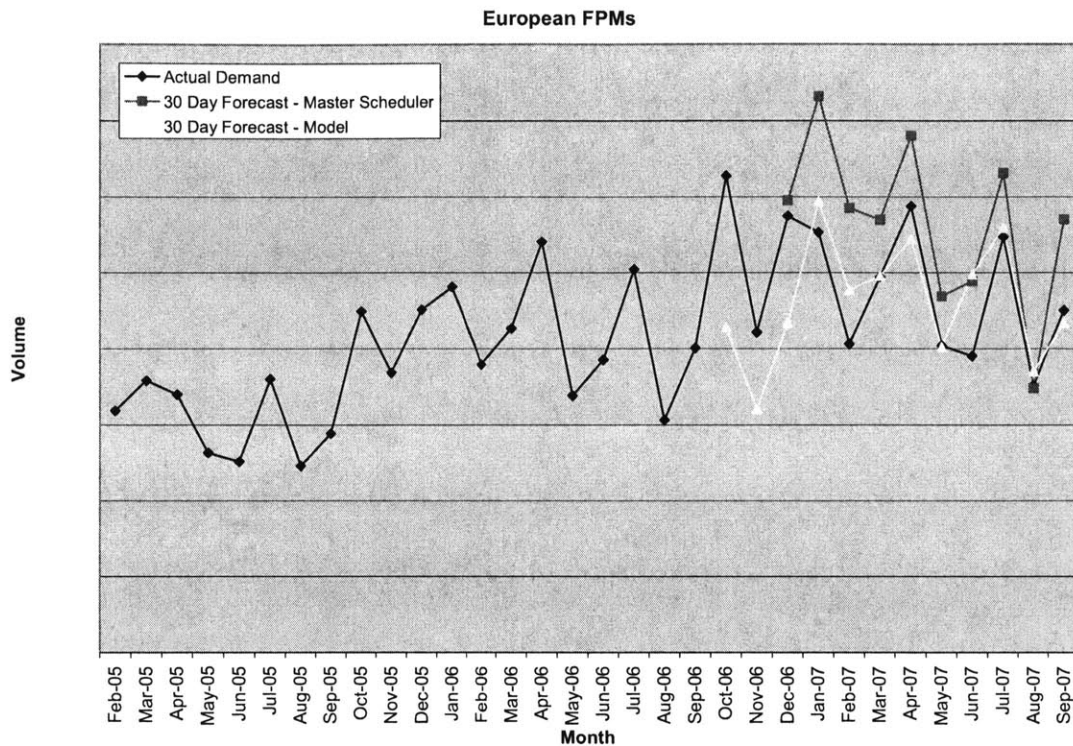


Figure 26 – 30 Day Forecast Comparison, European FPMs, Model B

The following table summarizes the forecast accuracy metrics for the master scheduler, model A, and model B.

	Master Scheduler			Model A			Model B		
	90 Day	60 Day	30 Day	90 Day	60 Day	30 Day	90 Day	60 Day	30 Day
Dec-06	83%	90%	97%	159%	146%	133%	126%	120%	106%
Jan-07	72%	90%	76%	138%	94%	93%	111%	99%	104%
Feb-07	74%	67%	69%	104%	100%	85%	86%	90%	86%
Mar-07	82%	80%	87%	113%	86%	100%	103%	99%	99%
Apr-07	82%	80%	86%	88%	108%	108%	94%	93%	96%
May-07	81%	86%	86%	101%	100%	101%	92%	97%	96%
Jun-07	75%	77%	80%	80%	81%	78%	83%	82%	84%
Jul-07	82%	88%	87%	103%	97%	98%	97%	99%	99%
Aug-07	85%	85%	101%	85%	87%	95%	94%	95%	100%
Sep-07	70%	79%	79%	93%	107%	104%	98%	103%	101%

Figure 27 – Forecast Accuracy Metrics Summary, European FPMs

An interesting takeaway of this comparison relates to the consistent bias towards over-forecasting exhibited by the master scheduler. Clearly, factors other than forecast accuracy play an important role when measuring the results of forecasting process. While the statistical model was generally more accurate than the master scheduler, there were no instances where the master

scheduler was under-forecast. On the other hand, the model under-forecast the majority of the time at the 90-day level. The master scheduler's consistent track record of over-forecasting results from the fact that they know there are high stock-out costs when it comes to monitors – they have a long lead-time and they are bulky, making them costly to expedite by air.

2.10 Takeaways

The following takeaways emerged from the statistical modeling discussed in this section.

1. Historical Data – Due to the short lifecycles associated with most parts, model generation was a challenge. While creating time series models based on shorter data sets is possible, the models generated from shorter data sets were less accurate because the trend and seasonality of the data was more difficult to pick up. This is less of a problem in longer data sets (at least one year in duration), as the time series analysis is able to capture the overall trends and seasonality patterns.

The relationship between product lifecycles and product volumes in regard to statistical forecasting can be considered in one of the four quadrants of Figure 28:

- Low Volume-Short Lifecycle
- Low Volume-Long Lifecycle
- High Volume-Short Lifecycle
- High Volume-Long Lifecycle

High volume parts with long lifecycles are the only parts for which time series forecasting is an exact fit. This was the case with the analysis conducted as part of the research effort. The modular battery discussed above had a significantly longer product lifecycle than the majority of the parts available in current systems. It also had enough volume to endure any one-time volatility due to promotions and market trends. As Figure 28 highlights, this is more difficult at lower volumes. While long lifecycle, low volume parts can still benefit from time series models, they need to be augmented by input from the marketing and sales teams. At lower volumes, the forecasts are much more susceptible to slight forecast misses due to pricing, promotions, or supply availability.

For short lifecycle parts, time series modeling is much more difficult to implement. Time series profiles of similar parts may offer insights into a newer part. In the world of Dell, where part lifecycles are often 12 months or less, a part profiling approach may offer the optimal approach. The parts themselves do not have enough history of their own to make time series forecasting a viable option. However, once a part reaches the end of its life, its historical demand profile can be used to build a demand forecast for like parts in the future. Each time a similar part reaches its end of life, the part's data can be added to the next iteration of the model, thus improving the overall accuracy of the model.

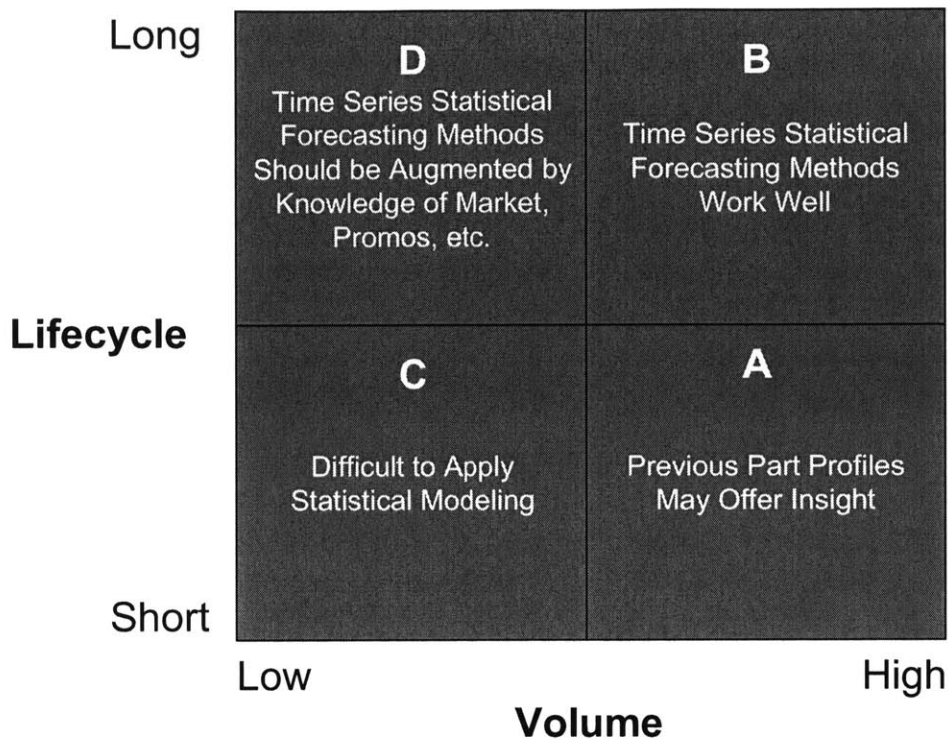


Figure 28 – Forecasting Lifecycle vs. Volume²¹

- Spikes and Dips – Unexpected changes in the demand data had two effects. One, they caused forecast accuracy degradation in the actual month where the demand shift occurred. The models based their predictions solely on historical data with no input regarding promotions or supply continuity issues. Two, once these spikes and dips became part of the historical data set, the model would factor them into future predictions of trend and seasonality. In this regard, promotions and demand shaping that are meant to drive demand in a particular month can actually have the effect of influencing the time series forecast for future months.

One way to minimize the impact of these spikes and dips (i.e. outliers) is to systematically remove them from the data set before building the model. In their place, the model could utilize dummy data points that serve to smooth out the peaks and valleys. A key to making this approach successful is determining whether a data point truly represents an outlier brought on by other factors (e.g. promotions, supply constraints) or whether it is simply a reflection of true demand. It would be difficult to build this determination into the model automatically. Instead, a master scheduler with detailed knowledge of a commodity would need to systematically classify potential outliers.

Crum and Palmatier discuss the relationship between demand variability and product volumes, as identified in Figure 29. Low variability, high volume parts offer the best combination for time series forecasting. Low variability, low volume parts are also a good fit

²¹ Crum, 139.

for time series forecasting, but require the same reinforcement from marketing and sales as discussed above for low volume parts with long lifecycles.

Highly variable parts offer a greater challenge for time series forecasting. With little in the way of recognizable trend and seasonality patterns, it is more difficult for the time series model to make accurate predictions.

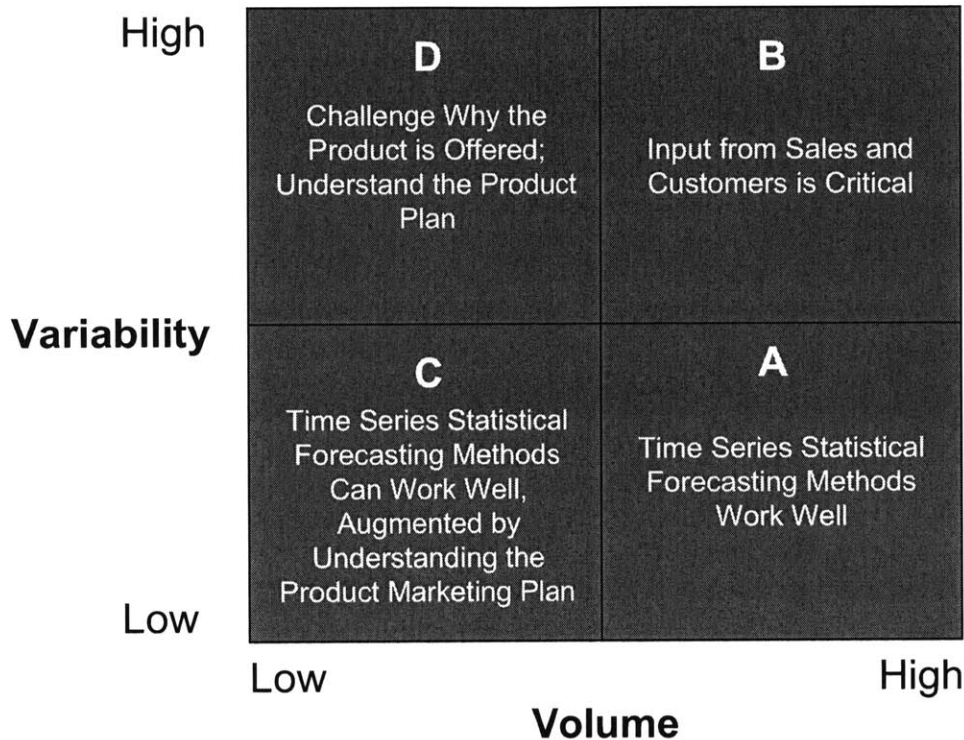


Figure 29 – Forecasting Variability vs. Volume²²

- Part Aggregation – Not surprisingly, the farther up the product hierarchy predictions were made, the more accurate the results. This was most evident from the European FPM example. To apply this lesson to the forecasting process, master schedulers should consider applying time series models to entire categories of parts. They can then make sure the sum of their individual part forecasts falls in close proximity to the overall time series forecast.

2.11 Considerations

One of the overall goals of the statistical study was to compare the effectiveness of attach rate forecasting to attribute based forecasting. The attach rate approach utilized by the master scheduling team has some significant advantages. It enables the master schedulers to forecast part-level demand as a function of system-level demand. Master schedulers do not need to know what system-level demand will be, simply what percentage of a system’s demand will contribute

²² Crum, 139.

to a specific part's demand. Attach rate forecasting also allows master schedulers to attach parts at any level of the product hierarchy.

The attribute-based approach, of which time series modeling is one potential application, takes system level demand out of the equation. Instead, it utilizes attributes of the part itself to make predictions. In the case of time series modeling, the attribute used is the historical sales data of the part. Other applications could involve cost, price, promotion, or lifecycle information.

Based on the results of this analysis, attribute-based forecasting in this form should not replace attach rate forecasting. Subtle changes in demand patterns based on system-level demand are not captured in attribute-based forecasting. As Hanke and Wichern put it, "A judicious mixture of quantitative techniques with common sense is always necessary if forecasts are to be accurate, understood, and used by the firm's decision makers."²³

Instead of replacing attach rate forecasting, attribute-based forecasting should be used as a complement to it. Attribute-based forecasting has the advantage of requiring very little in the way of manual input. An information technology system can be set up to create an attribute-based forecast on a recurring basis. This forecast can then be used as a litmus test to compare attach rate forecasts against. It can also be used to create the forecast itself for certain types of parts that have longer lifecycles and fewer promotions.

An important consideration in implementing any type of statistical forecasting into the current process is the managerial aspect of doing so. As part of the analysis described above, numerous interviews with current master schedulers were conducted. Many master schedulers were enthusiastic about the potential uses for statistics in the forecasting process. In fact, many are already using statistics in one form or another to support their current forecasting process. Others are less familiar with the potential application of statistics in the forecasting process.

Any organization that hopes to successfully implement a change to a pre-existing process needs to consider the human aspect of the change. In this case, part-level forecasts are currently the domain of the master schedulers. Any change to the part-level forecasting process thus needs to involve master schedulers in the planning stages. There is a sizable body of tribal knowledge within the Demand/Supply organization that has grown over years of forecasting that cannot be replaced by statistical modeling.

One path is to create a master scheduler statistics working group. This would create a dual effect. First, it would join together the current pockets of statistics that currently reside within the master scheduler community. Second, it would create a team of stakeholders to drive changes into the existing process. If statistical analysis is viewed as a mandate that is coming down from above, it has less of a chance of actually being implemented. If it is driven into the process from the grassroots level, it is more likely to be utilized.

²³ Hanke, 485.

3 Supply Hedging

To help account for forecast variability, the Demand/Supply team has the option of adding a hedge to the forecast of a particular system or part. Much like a financial hedge that is designed to protect an investor from a loss, a supply hedge is designed to reduce the company's risk exposure to a forecast miss.

Typically, a supply hedge will involve a master scheduler adding volume to a particular forecast. Occasionally, a supply hedge will involve a master scheduler subtracting volume from a particular forecast. For example, consider the case of a part where supply is expected to outpace demand and result in an excess supply. A master scheduler may negatively hedge the forecast for another part in that same family if they anticipate promotions to shift demand toward the excess supply of the original part.

However, a more typical scenario involves a supply hedge that is designed to *increase* the available supply for a particular system or part. At the system level, positive hedges are most commonly associated with two situations: new product introductions and seasonality. When new products are identified, they are tracked on the product roadmap corresponding to one of the four master scheduler teams. Months or years of planning go into the identification, design, development, testing, and rollout of new products. Nearly every organization within Dell is involved during this process: engineering, marketing, procurement, manufacturing, support, and sales. Advanced sales, strategic marketing, and profiles of previous system introductions can all be used to hone in on a specific demand forecast.

Yet, despite the significant planning that takes place surrounding the creation of new products, there is an inherent level of uncertainty associated with the introduction of these products into the marketplace:

- How will customers react?
- How will suppliers react?
- How will the competition react?

The answer to these and other questions will all impact the demand levels associated with the new products. To help account for the inherent uncertainty surrounding the introduction of new products, master schedulers will often add a hedge to the sales forecasts they receive for a new product. For example, if the master sales plan (MSP) indicates the company will sell 100,000 units of a new product in the first month, a master schedule may decide to drive a 5% hedge into the product just in case demand outpaces the sales forecast. The resulting forecast of 105,000 units would be provided to the procurement team and ultimately to the vendors responsible for the new product.

Hedges are also common during peak periods. A main focus of the customer experience is product availability, and this is particularly important during peak periods of demand. For individual consumers, sales typically peak during the back-to-school season and the holiday shopping season. For government customers, sales peak during state and federal buying seasons.

During these periods of heightened demand, it is crucial that no customers are given an extended lead time on their product due to insufficient availability. Even if a customer decides to wait for the order, it could lead them to use a competitor the next time they have a purchase to make. To help protect against this potential outcome, master schedulers can add a hedge to the sales forecasts to increase the overall availability of supply.

On the surface, the hedge has significant benefits for Dell. It increases the available supply should demand increase beyond forecasted levels, creating a safety net to help protect against unforeseen demand spikes. At the same time, if sales do not outpace the forecast, the excess supply inventory remains outside of Dell's factories in the SLCs. However, a closer look reveals a more complex paradigm for both Dell and its suppliers.

3.1 Current Process

In order to more fully understand the hedge, interviews were conducted with various stakeholders throughout the organization. The goal was to determine the impact the hedge has both internally on Dell operations as well as externally on vendors throughout the supply chain. Over the course of the analysis, 41 interviews were conducted with various internal stakeholders in the forecasting process. They fell into the following groups:

- Demand/Supply Master Schedulers – 11
- Demand/Supply Managers – 8
- Demand/Supply Analysts – 2
- Procurement – 8
- Buyers – 2
- COC – 5
- Production Control – 3
- Manufacturing – 2

An important takeaway from these interviews was the fact that adjustments are made to the forecast at multiple points in the forecasting process. While only one of these adjustments is in fact called a “hedge”, the adjustments all have the same effect of modifying the original forecast figure. In “Inventory Decisions in Dell’s Supply Chain”, Kapuscinski, et al describe this phenomena at Dell. “While innocently acting with the company’s best interest in mind, each group used its own judgment and biases to modify the forecasts of demand. Such iterative hedges and adjustments can erode the quality of the original information.”²⁴

The high-level process map in Figure 30 identifies the basic flow to the Dell forecasting process that emerged from these interviews.

²⁴ Kapuscinski, 199.

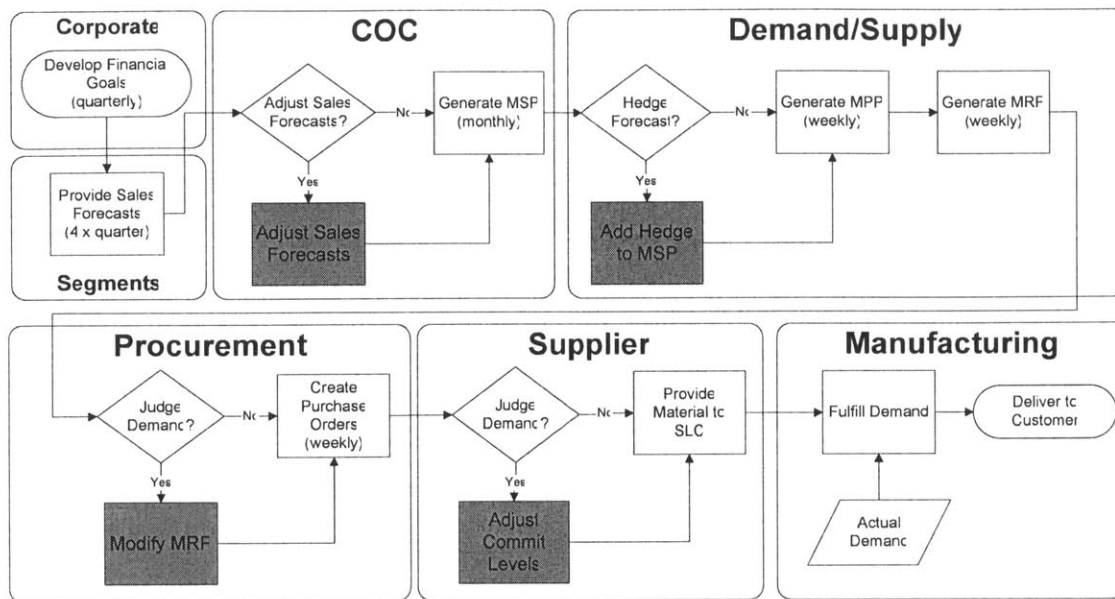


Figure 30 – High-Level Map of the Forecast Process

The process begins with the financial targets that are generated at the corporate level on a quarterly basis. The corporate financial targets represent the high-level revenue goals for the company as a whole. They are not broken down into a mix of specific systems. This begins to take place within the different sales segments responsible for each Dell customer segment. Three to four times per quarter, the sales segments produce sales forecasts identifying how many of each system they will sell to meet the corporate financial targets. The individual sales forecasts from each segment are then passed onto the COC to consolidate across the entire company. At this point, the COC will determine if any adjustments are needed to the consolidated sales forecasts before generating the monthly master sales plan.

When the MSP is passed on to Demand/Supply, the system master schedulers have the opportunity to add a hedge to the forecast. This will involve input from the COC to determine the level of confidence surrounding individual system projections in the MSP. If a new product is being introduced, if a seasonal sales period is approaching, or if a promotion is being planned, the master scheduler and their COC counterpart will often add a hedge to the forecast. During more stable periods, they will often decide to add little or no hedge to the forecast.

The part-level MRP is produced on a weekly basis and is passed onto the procurement organization, which is responsible for generating purchase orders for each part and supplier combination. Typically, procurement will take the forecast as is and place supplier orders based on it. However, procurement does have the opportunity to provide input on the forecast and will occasionally suggest changes to the forecast levels to Demand/Supply. This can be an adjustment to either increase or decrease the forecasted supply level. While infrequent, this judgment is typically made when suppliers have expressed concern over high supply levels in the SLC for a particular part. In turn, procurement may decide that current supply levels provide enough of a cushion such that the forecast can be adjusted downward. On the occasions where procurement

wants to adjust the forecast, Demand/Supply and procurement must agree to the adjusted level before officially doing so.

The suppliers themselves can also make a similar judgment. Even though they are contractually required to supply the agreed upon number of parts identified in each of their purchase orders, suppliers will sometimes decide to send fewer parts. The dynamics surrounding this decision get to the heart of the complexity of the hedge question for a build-to-order manufacturing system.

3.2 Impact of the Hedge

It is important to understand there are competing forces when it comes to the impact of a hedged forecast. Figure 31 is adapted from a presentation given by Kevin T. Jones, the Manager of the Tools and Processes team within the Demand/Supply organization. It identifies the different perspectives represented by Dell and its suppliers regarding the hedge.

<u>Dell Perspective</u>		<u>Supplier Perspective</u>
<ul style="list-style-type: none">• Except for EOL phase, there is virtually <i>no downside</i> to a hedged forecast	▶	<ul style="list-style-type: none">• Too much inventory means absorbing <i>carrying charges</i>, potential obsolescence
<ul style="list-style-type: none">• Hedged forecast will generally keep liability for <i>expedites</i> in supplier's court	▶	<ul style="list-style-type: none">• Too little inventory may require expediting (question of <i>who pays</i>)
<ul style="list-style-type: none">• Significant downside for <i>under-forecasting</i> when this leads to a supply issue	▶	<ul style="list-style-type: none">• Dell forecast numbers are historically <i>overstated</i>

Figure 31 – Competing Forces in Regards to Hedging²⁵

1. Other than a part's end-of-life phase, there is little downside for Dell to adding a hedge to a forecast. This is because the inventory remains on the supplier's financial books until it enters the Dell factory to be added to a customer order. On the other hand, the supplier is required to pay holding costs for that inventory throughout the process – from the production of the part in its own factory to the transport of the part to the supplier logistics center to the time it spends in the center awaiting a customer order. This will, of course, hit Dell's books eventually. An over-stocked supply chain means parts lose their value as they wait to be included in systems. Historically, though, the focus has been on avoiding out-of-stock situations.

²⁵ Jones, Kevin T. "COS FY07 Re-Engineering: Hedging." 3 Jan. 2006.

2. Adding hedge to a forecast will generally keep the exposure to part expediting costs with the supplier. Supplier contracts require them to provide the supply levels identified by Dell. If they do not, they are technically on the hook for any expedite costs that are required to rush supply to Dell factories to fulfill customer demand. If demand outpaces the forecast provided by Dell, it is Dell's responsibility to pay for the costs to expedite materials. Thus, it is in Dell's best interest to upwardly hedge its forecast to reduce its exposure to expedite costs.

3. With its focus on the customer experience, it is crucial for Dell to avoid supply issues caused by an under-forecast. Culturally, there is significant emphasis placed on maintaining the continuity of supply (COS). Other than exposure to end-of-life obsolescence costs, there is very little contractual downside for Dell to carrying excess inventory. With an emphasis on COS and little in the way of contractual downside, this naturally leads to a pattern of buffered forecast figures to ensure that demand can be met by actual supply levels. Not surprisingly, overall average forecast accuracy is below 100%, indicating a bias to overstating the true forecasted demand. Because of the just-in-time, build-to-order nature of Dell's supply chain, its suppliers know this. Suppliers see the demand patterns in near real-time because they see their parts pulled from the supplier logistics centers in response to customer orders. They can then compare the total number of parts pulled from the SLC for a given week or month and compare it to the total number of parts forecasted by Dell. If the numbers differ, it is likely that the forecasted figures are greater than the actual pulls. This pattern of conflicting signals is identified in Figure 32 below, from a presentation by Ted Vanderlaan.

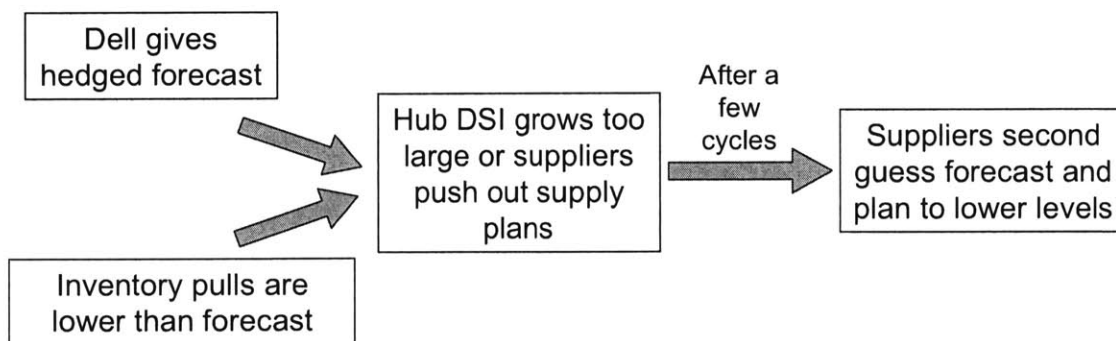


Figure 32 – Inventory Pulls vs. Hedged Forecast²⁶

In this model, the supplier repeatedly sees inventory pulls that are lower than forecasted pulls. As sales come in lower than forecasted, the level of days of sales inventory (DSI) in the supply hub steadily increases. If this pattern repeats for a number of months, the suppliers will no longer trust the forecasted figures and may decide to lower their incoming supply, regardless of the forecast.

Throughout the course of the interviews, there were numerous anecdotes of this taking place. One such case involved supplier “decommit” on a large capacity hard drive that customers

²⁶ Jones.

could add to Dell's high-end desktop systems. As is typical at the high-end of the system spectrum, volumes were relatively low on this particular hard drive. Yet, the emphasis on the customer experience for customers of high-end products is critical. Customers of high-end products help shape overall public opinion, as they often share their experiences via word-of-mouth and on-line channels.

As discussed in Section 2, low volume products are particularly difficult to forecast accurately. Add the high-end nature of the product on top of that and there was a recipe for over-forecasting from the start. After three months of seeing the over-forecasting firsthand, the supplier decided to adjust its supply commitments. Once this supply adjustment made it into Dell's tracking system, it immediately raised a red flag to the master scheduler in charge of the particular hard drive. The situation was eventually resolved through a series of communications between the master scheduler, the COC, the buyer, and the supplier. The case study points to the competing forces and conflicting signals that can arise from a hedge in a build-to-order supply chain.

Noticeably, the signaling issues that can arise are dramatically different from other, non-build-to-order manufacturing companies. Jones considers the difference most distinct in the short term. For most companies, the short term procurement process is fairly stable. While forecasts may fluctuate outside of a certain number of days (typically the lead time for a given part or system), within that lead time there is little that changes in the way of en route deliveries. Once a purchase order is in place, if there are too many parts or systems on their way, they will sit in the factory or warehouse of the retailer. Dell's short term window, on the other hand, behaves a little differently. Dell typically over-forecasts just in case demand hits one of its spikes. Yet, it only pulls what it needs from the parts inventory. So unless one of the demand spikes actually takes place, the pulls are going to be less than what it has asked its suppliers for in its hedged forecast. As Jones put it, this "leads to a credibility issue and non-value-add second-guessing."²⁷

3.3 Improvement Opportunities

If Dell wants to improve the signaling issue that has developed with its suppliers, it needs to modify its approach to supply hedging. Changes to its hedging philosophy will also enable Dell to reduce its exposure to obsolescence cost and improve its overall forecast accuracy. Of course, any changes to the forecast hedging approach need to be weighed against the perceived downsides. Currently, the quantifiable downsides to the hedge are the exposure to obsolescence costs and the cost declines associated with over-stocking. If inventory that Dell requested is never sold, it will ultimately be transferred off of the supplier's balance sheet and onto Dell's. Just as it tracks its costs in expediting materials to fulfill customer orders, Dell also vigorously tracks its obsolescence costs due to customer demand that never materializes. Driving E&O ("Excess and Obsolete) costs as low as possible is a constant goal of the Demand/Supply organization.

²⁷ Jones, Kevin T. E-mail to Marshall Einhorn. 30 Jan. 2007.

While this is the main quantifiable downside of excess hedging, there are other potential pitfalls that are more difficult to put a number on. One is the potential disruption to Dell's continuity of supply. If the pattern of over-forecasting described above leads to a supply issue because of a supplier decommit, it may mean a lost sale due to unavailable supply. Another is the idea that the pattern of over-forecasting leads to higher per part prices from suppliers. This was widely recognized during the course of the interviews as a disadvantage of the standard approach to hedging. While it is difficult to quantify, the system dynamics that lead to that conclusion are identified in the following figure.

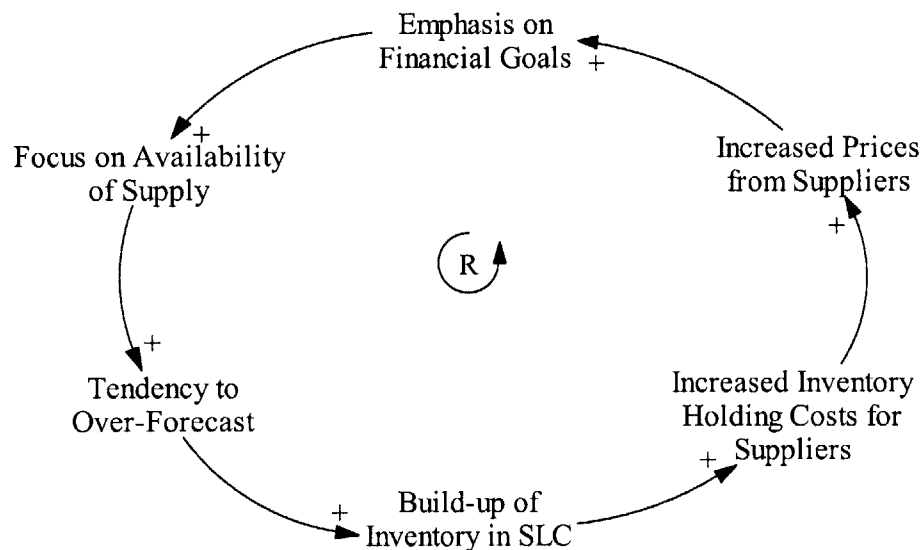


Figure 33 – System Dynamics Model of Hedge

The overall dynamics of the situation are as follows. As with any publicly traded company, it begins with an emphasis on the financial goals of the organization. With this as a backdrop, it places a focus on the availability of supply to meet customer demand. Positive customer experience is a key aspect of attaining financial success and having supply available to meet customer orders is critical to making that happen. Due to the nature of the contracts between Dell and its suppliers, this leads to a tendency of Dell to over-forecast to ensure supply is available. In turn, this leads to a build-up of inventory in the supplier logistics centers. Increased inventory leads to increased inventory holding costs, which hit the supplier financial figures. With higher costs, suppliers have to charge higher prices to Dell to meet their own internal financial goals. These higher prices in turn increase the price of doing business for Dell. Ultimately, this places an even greater emphasis on meeting financial goals and not losing any potential sales due to supply issues.

Based on the common themes from the interviews, the following series of recommendations for optimizing Dell's hedge strategy emerged.

1. Forecasts are currently spread evenly across each week within a month. However, there are rarely months when demand consistently hits the same figure week after week. As a result, there is logic built into Dell's IT systems that handles how forecast misses are tracked from

an inventory perspective. Currently, this logic is setup to take the delta between actual sales and forecasted sales for a particular part and add that difference into the forecast figure for the following week. While this helps the forecast track the hockey stick nature of weekly demand patterns (demand in week 1 is typically less than week 2 which is typically less than week 3 and so on), it also rolls over the missed hedge percentage from week to week. As a result, buyers will sometimes remove this miss when providing updated figures to suppliers, even though master schedulers assume the miss is included.

Rolling over the forecast miss is designed to portray the hockey stick demand pattern. To obtain the same results, a project is currently underway that would allow master schedulers to enter weekly forecasts instead of monthly forecasts, thus eliminating the need to spread forecasts evenly across each week. With the ability to mirror the ramp up of demand from week to week, there would be no need to roll over the forecast miss from week to week and any missed hedge would not be rolled over from week to week as well.

2. Many were in agreement that driving the hedge as low as possible was the desired outcome. The hedge effectively shields Dell from the downside of not matching supply with demand, but it forces substantial waste into the process. Even though most of this waste currently resides in the supplier's court, a holistic view of the supply chain would demand a new approach to the hedge. One approach to reducing the impact the hedge has directly on the supplier's books and indirectly on Dell's is to use a true forecast signal. Instead of a hedge, each forecast could include a measure of forecast error such as standard deviation or confidence intervals. The supplier contracts could be setup to provide incentives if suppliers provide inventory to handle an agreed upon level of forecast variance (e.g. two standard deviations above the forecasted mean). It could accomplish the desired result of the hedge – enough supply to deal with potential demand spikes – in a way that puts both parties on the same page.

This focus on a truer signal of forecasted demand echoes a key tenet of the Dell Direct Model. Michael Dell makes reference to the role information sharing plays in its supplier relationships, "The key is in providing your supplier with all the information they need to make an informed decision. A lot of that has to do with sharing your strategies and goals openly with them."²⁸ The move away from a fully hedged forecast is simply an extension of this principle.

3. A key tool in dealing with demand variability is the use of safety stock, as described in the book "Managing Business Process Flows":

... because customer-demand forecasts are usually inaccurate, planning process output to meet only forecasted demand may result in stockouts, delayed deliveries, lost sales, and customer dissatisfaction. Thus, many producers maintain safety inventories of outputs to absorb excess demand and to ensure product availability and customer service despite forecast errors.²⁹

²⁸ Dell, Michael, with Catherine Fredman. *Direct From Dell: Strategies That Revolutionized an Industry*. New York, NY: HarperCollins, 1999. 187

²⁹ Anupindi, 111.

In Dell's case, maintaining safety inventories of *outputs* is not an option. However, with a production time that is measured in hours instead of days for most of its products, Dell can accomplish the same effect by carrying a safety inventory of parts.

The equation to determine the appropriate level of safety stock (I_s) to hold for an individual part is given in Equation 4, where z is the desired level of service, σ is the standard deviation of the demand for the given part, and L is the lead time.

$$I_s = z * \sigma * \sqrt{L}$$

Equation 4 – Safety Stock

Based on this view, safety stock will increase (decrease) with an increase (decrease) in:

- Demand variability or forecast error
- Delivery lead time for the same level of service
- Delivery lead time variability for the same level of service³⁰

Currently, Dell's safety stock policy requires each of its suppliers to generally maintain the same number of days of inventory for each part. Based on this number, individual part inventory levels are tracked on a weekly basis and deemed safe or at risk of a stockout based on the recent run rate of the part. Instead of applying a standard requirement to the number of days of inventory, a more appropriate approach would involve tailoring the safety stock policy to each specific part or type of part.

For example, if we assume that each part has the same lead time and the same desired service level, then the variability (standard deviation) of the part's demand becomes the driving factor in determining safety stock levels. For parts that have volatile demand, and thus a higher standard deviation, the safety stock equation demands a higher level of safety inventory. For parts that have more stable demand, the equation requires lower safety inventory. The following figure highlights two different parts that exhibit the same average demand but different levels of variability. As a result the level of safety stock that is required is different.




Figure 34 – Safety Stock Comparison: High Volatility vs. Low Volatility

³⁰ Novak, Sharon. "Class #11." Course 15.761: Operations Management. Massachusetts Institute of Technology, Cambridge, MA. July 2005.

Once other factors enter into the equation, such as segmented levels of service based on the part type (e.g. higher levels of service for higher end systems) or different lead times (e.g. monitors coming from Mexico vs. Asia), it becomes even more appropriate to differentiate safety stock requirements as well. Within Dell, this concept has already taken shape. In particular, Ted Vanderlaan has put together an intriguing analysis that points to a safety stock policy based on demand volatility and part cost.³¹

3.4 Implementation Considerations

Any changes to the hedge should be implemented with participation from all parts of the supply chain. The system that is currently in place has been developed through both formal (e.g. contracts, meetings) and informal (e.g. relationships, oral agreements) channels over many years. Making changes to the forecasting process without proper buy-in from the stakeholders involved in those channels is not advisable. To help drive changes into the process, Dell should consider the creation of a task force involving representatives from each part of the organization that is impacted by the forecast. This task force should also include representatives from one or more suppliers.

The task force should start small. Starting with wholesale changes to an established process such as the hedge or the safety stock policy will be fruitless. Instead, the changes should initially focus on short term wins, such as the modification of the forecast miss rollover logic. After aligning this process across the supply chain, the team can then move onto bigger ticket items.

³¹ Vanderlaan, Ted. "DSI Target Setting." 27 Nov. 2006.

4 Conclusion

Dealing with demand variability is a delicate balancing act for any manufacturer or retailer. On the one hand, there are potential obsolescence costs of unused inventory should sales not materialize. On the other hand, there are potential expedite costs, or worse, lost sales, should sales outpace available inventory. As both a manufacturer and a retailer, Dell is in a unique position; the company strives to deliver an extensive array of products and features while optimizing the amount of inventory required to support those product lines throughout its supply chain.

Dealing with demand variability is further complicated by features of the Dell Direct Model. With its build-to-order manufacturing system, Dell cannot hold its safety stock in the form of finished goods inventory. Instead, safety stock must be held as parts inventory, sitting in supplier-owned SLCs. Dell is ultimately responsible for the inventory in the SLCs but it remains off its books until an order is started. With vendors in charge of stocking the SLCs based on Dell's forecasts, it creates a potential disconnect – stocking a full set of safety stock may be in Dell's best interests but it may not be in the supplier's best interests. As a result, supplier stocking decisions may impact Dell's ability to respond to forecast variability. Likewise, globalization, product proliferation, and geo-manufacturing all magnify the impact variability has on the forecasting process.

Two methods of dealing with demand variability were discussed in this thesis. Section 2 looked at the potential application of statistical modeling techniques to the part-level forecasting process. In particular, it looked at the use of time series models to forecast part-level demand. A series of back-cast forecasts were created to examine the effectiveness of the selected models. While the results did not merit a recommendation of utilizing time series forecasts across the board in lieu of the current process, there are certain applications of time series forecasts that Dell should consider.

One, time series modeling can be an effective means to quickly provide a straw forecast for one or many parts. Adding time series models to the upcoming system upgrade would give master schedulers the ability to quickly assess their forecasts against a statistical litmus test. Two, for parts where historical data is plentiful, demand is relatively smooth, or parts are aggregated with one another, it may be appropriate to use time series forecasts. This would allow master schedulers to focus on parts that are more volatile and parts that are newly introduced or soon-to-be discontinued.

Section 3 examined how hedging is currently utilized as a means to account for demand variability. While beneficial to Dell on the surface, its track record of consistently adding a hedge to the forecast is potentially detrimental to its vendor relationships. It has the direct impact of driving excess inventory onto the books of its vendors and it has the indirect impact of higher per part costs to Dell. It also exposes Dell to part shortages due to supplier decommits. To help counter these effects, Dell should consider three changes, as described below.

First, the company should consider modifying its information systems to allow for different forecasts from week to week. This will help avoid the information confusion that results from the forecast miss rollover each week. Second, Dell should work with its suppliers as it drives hedges lower. To help maintain the same level of service, it could share a forecast error metric with its suppliers and incentivize them to meet a pre-determined level of variability. Third, Dell should implement a non-uniform safety inventory policy based on the specific features (e.g. variability, cost, lead time) of individual parts or categories of parts.

Moving forward, the Demand/Supply team should continue to analyze the use of regression models in forecasting part-level demand. While time series techniques draw upon the historical patterns in the sales data itself, statistical modeling techniques that attempt to draw relationships between part-level demand and other factors (such as part lifecycle stage, price, promotions, and demand for other parts) represent a potential implementation path for statistical modeling. If Dell can navigate past complications surrounding data gathering for these other factors, it could prove more fruitful than the time series approach described in Section 2. Likewise, demand profiling is another approach that merits consideration. One of the main hurdles faced by the time series approach involved the short lifecycles of most Dell parts. To help circumvent that issue, like parts can be combined to produce lifecycle profiles to use as the basis for future part forecasts.

5 Bibliography

- Anupindi, Ravi, Sunil Chopra, Sudhakar D. Deshmukh, Jan A. Van Mieghem, and Eitan Zemel. Managing Business Process Flows. Upper Saddle River, NJ: Prentice Hall, 1999.
- Breen, Bill. "Living in Dell Time." Fast Company Nov. 2004. 27 Mar. 2007
<<http://www.fastcompany.com/magazine/88/dell.html>>.
- Crum, Colleen, with George E. Palmatier. Demand Management Best Practices: Process, Principles and Collaboration. Boca Raton, FL: J. Ross, Inc., 2003.
- Dell, Michael, with Catherine Fredman. Direct From Dell: Strategies That Revolutionized an Industry. New York, NY: HarperCollins, 1999.
- Friedman, Thomas L. The World is Flat: a Brief History of the Twenty-First Century. New York: Farrar, Straus and Giroux, 2006.
- Hanke, John E., and Dean W. Wichern. Business Forecasting, 8th ed. Upper Saddle River, NJ: Pearson Prentice Hall, 2005.
- Jones, Kevin T. "COS FY07 Re-Engineering: Hedging." 3 Jan. 2006.
- Jones, Kevin T. E-mail to Marshall Einhorn. 30 Jan. 2007.
- Kapuscinski, Roman, Rachel Q. Zhang, Paul Carbonneau, Robert Moore, and Bill Reeves. "Inventory Decisions in Dell's Supply Chain." Interfaces 34 (2004): 191-205.
- Montgomery, Douglas C. Introduction to Statistical Quality Control. 4th ed. New York, NY: John Wiley & Sons, Inc., 2001.
- Novak, Sharon. "Class #11." Course 15.761: Operations Management. Massachusetts Institute of Technology, Cambridge, MA. July 2005.
- Ponthier, Clark. "Dell Supply Chain Overview." Presentation given to Pennsylvania State University, University Park, PA. 8 Sept. 2006.
- Reyner, Amy. Multi-Site Inventory Balancing in an Extended Global Supply Chain. Thesis. Massachusetts Institute of Technology, 2006.
- Rivkin, Jan W., and Michael E. Porter. "Matching Dell." Boston, MA: Harvard Business School, 1999.
- "Time Series Analysis." StatSoft, Inc. 27 Mar. 2007
<<http://www.statsoft.com/textbook/sttimser.html#exponential>>.
- Vanderlaan, Ted. "DSI Target Setting." 27 Nov. 2006.