Inventory Optimization in High Volume Aerospace Supply Chains

by

Brian Robert Masse

B.S. Mechanical Engineering, University of Notre Dame, 2005
M.S. Mechanical Engineering, University of Massachusetts Lowell, 2008

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

Master of Business Administration
and
Master of Science in Engineering Systems

In conjunction with the Leaders for Global Operations Program at the
Massachusetts Institute of Technology
June 2011

Signature of Author_______________

MIT Sloan School of Management
Department of Engineering Systems
May 6, 2011

Certified by_________________

Dr. David Simchi-Levi, Thesis Supervisor
Professor of Engineering Systems and Civil & Environmental Engineering

Certified by_________________

Dr. Roy E. Welsch, Thesis Supervisor
Professor of Statistics and Management Science

Accepted by

Dr. Nancy Leveson
Professor, Aeronautics and Astronautics and Engineering Systems Division
Chair, Engineering Systems Division Education Committee

Accepted by _______________

Debbie Berechman
Executive Director of MBA Program
MIT Sloan School of Management

This page has been intentionally left blank.

Inventory Optimization in High Volume Aerospace Supply Chains

by

Brian Robert Masse

Submitted to the MIT Sloan School of Management and the Department of Engineering Systems on May 6, 2011 in Partial Fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Engineering Systems

Abstract

The supply chains of aerospace products can be complex, involving thousands of components per product and hundreds of vendors spaced out over an increasingly global landscape. Managing all inputs necessary for these complex aerospace supply chains is a task that is critical to the success of any firm and requires extensive planning, close partnerships, and detailed analysis.

This thesis outlines a system for optimal safety stock management in high volume aerospace supply chains. Given such supply chain parameters as component inventory values, procurement and manufacturing lead times, demand distributions, and bills of material, the ideal safety stock locations and sizes which result in minimal overall inventory levels are calculated by a nonlinear optimization program. With this safety stock structure, aerospace firms can operate their supply chains with higher customer service rates and lower inventory levels.

A methodology is also developed to help aerospace companies improve their existing supply chains as efficiently as possible. Considering the limited time and resources available, a company may not be able to enhance all areas of its operations and determining where to improve with the greatest effect on customer service levels and inventory can be difficult. The framework developed provides general guidelines to ensure improvement resources are being deployed most efficiently. Finally, business environment and operations considerations are discussed to aid companies in the process of implementing supply chain improvements and instituting organizational change.

Thesis Supervisor: David Simchi-Levi
Professor of Engineering Systems and Civil & Environmental Engineering

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

This page has been intentionally left blank.

Acknowledgements

I would like to thank Pratt & Whitney for giving me the opportunity to work with them at their East Hartford, Connecticut location. Many P&W employees spent considerable amounts of time with me, teaching me their manufacturing processes and inventory systems and answering my questions when they no doubt had more critical tasks. I was given the opportunity to learn more from them than I will be able to return.

I would also like to thank the Leaders for Global Operations program at MIT for providing me with experiences and opportunities I would not have thought possible when starting this program two years ago. My views of operations and management have forever been changed, and I am a better person, personally and professionally, for it. I would also like to thank my LGO classmates, from whom I have learned more than from any other source during my time at MIT. Their experiences, talents, and support have made my experience that much more valuable.

Lastly, I would like to thank my parents for their continued support through all of my endeavors. Most of all, I would like to thank my wife Jaclyn for her constant encouragement and patience through my long nights at the office, weekends doing work, and inevitable frustrations. I could not have completed this experience without you.

This page has been intentionally left blank.

Table of Contents

List of Figures

1. Introduction

1.1. Overview

In many ways, the aerospace industry presents an exceptional set of challenges. Product development cycles are long, capital equipment costs are high, and cost pressures are severe. During the procurement process, customers can easily substitute one company for another because product differentiation is often small. In environments as competitive as these, successful companies need to implement world-class strategies in all areas of their business. Firms with excellent product offerings may still be uncompetitive due to substandard supply chain, aftermarket, manufacturing, or marketing services. Only by excelling in all of these areas can an aerospace firm be successful in the marketplace.

The supply chains of aerospace companies are generally expansive, consisting of several partners or suppliers, internal manufacturing, aftermarket support, and commodity producers. Deliverable products can consist of millions of individual parts, each of which must be individually sourced, controlled, and delivered to the point of use. Within this vast supply chain network, the margin of error can be very small. An individual missing rivet can delay the delivery of an entire aircraft, and a replacement engine blade can keep an in-service aircraft on the ground. With such severe consequences, aerospace supply chain operations are as critical to the success of the company as any other functional area. Only through close partnerships with suppliers, honest evaluation of internal manufacturing processes, and careful analysis of all network practices can an aerospace supply chain serve as the strategic asset necessary to the success of the firm.

1.2. Objective

Managing high volume supply chains is a difficult task for any aerospace company. As ever decreasing lead times are demanded by customers, many aerospace supply chains cannot operate under make-to-order systems or be buffered only at finished goods due to the enormous inventory levels that would be required. This thesis aims to provide a system for managing these supply chains by strategically locating safety stock to buffer against uncertain demand. Both safety stock levels and locations are optimized to provide the lowest overall inventory levels while maintaining a specified customer service rate. A framework for supply chain improvement is then developed which allows aerospace firms to target specific operations areas for further enhancement to ensure the firm's resources are used most efficiently. In these discussions, a sample Pratt & Whitney supply chain will be used as a basis for illustration, but the practices detailed can be applied to any number of products manufactured in high volume and with uncertain demand. Through these discussions, aerospace companies will be able to further use these practices to improve their supply chains and increase their competitive position in the marketplace.

All data in this thesis is fictionalized and provided solely for the purposes of illustration of the processes and frameworks described. To protect the proprietary property of Pratt & Whitney and United Technologies Corporation, all part numbers have been removed, all data has been altered, and the scales have been removed from a number of figures.

2. Background

2.1. Aerospace Industry Overview

The aerospace industry consists of several different markets and segments, with many having long design and development cycles and capital-intensive manufacturing. As a result, few competitors exist in each segment and barriers to entry are high. In the large commercial aircraft segment, Boeing and Airbus hold most of the market share, but are being challenged by such companies as Bombardier, Embraer, and Comac. Similarly, Lockheed Martin, Northrop Grumman, and Boeing control most of the U.S. military aircraft market. Each of these manufacturers uses several tiers of suppliers to contribute such critical subsystems as engines, avionics, body components, environmental systems, etc. Often, these suppliers are partners in new product development, sharing some of the design responsibility and financial risk associated with a new program.

2.2. Pratt & Whitney

Pratt & Whitney (P&W) is a division of United Technologies Corporation (UTC), which includes aerospace firms Sikorsky and Hamilton Sundstrand in addition to Carrier, Otis, and UTC Fire & Security. Pratt & Whitney, headquartered in East Hartford, Connecticut, develops and manufactures aircraft engines, gas turbines, and space propulsion systems. The company, founded in 1925, was a key manufacturer of aircraft engines during World War II, and operates heavily in both the commercial and military sides of the jet engine business. In the commercial space, it supplies engines to such aircraft as the Boeing 747, 767, and 777 and the Airbus A300, A320, A330, and A380. On the military side, Pratt & Whitney manufactures engines for the F-16 Fighting Falcon, F-22 Raptor, F-35 Lightning II, and C-17 Globemaster III.

Three main competitors exist in the aircraft engine market: Pratt & Whitney, GE Aviation, and Rolls-Royce. Competition between the three firms can be fierce on some programs, while they partner on others. For example, P&W and GE are joint partners on the GP7000 engine that powers the Airbus A380, but are direct competitors on a number of other aircraft. While airlines typically have an engine choice when purchasing a new aircraft, a particular engine may occasionally be the solitary option for a given plane. For example, Pratt & Whitney does not have an offering in the lucrative Boeing 737 jet engine market. However, they have recently won a series of contracts that could earn back significant market share in the narrow-body commercial jet segment. The Pratt & Whitney PurePower® engine uses an advanced gear system which allows the engine fan to rotate slower than the low pressure compressor and turbine. This separation of engine components results in significant decreases in fuel consumption, operating cost, environmental emissions, and noise over engines currently on the market. Using this advanced technology, P&W has won contracts to power the Mitsubishi Regional Jet, Bombardier CSeries, and Irkut MC-21, as well as a new engine option on the Airbus A320. With the emergence of this new technology, Pratt & Whitney is well poised to earn back a significant share of the narrow-body commercial aircraft market.

2.3. Supply Chain Dynamics and Business Model

2.3.1. Aerospace Industry

The aerospace industry, in contrast to many consumer goods and other technology businesses, is characterized by very long product cycle times. New aircraft frequently require a decade or more to design and test, and may be in service for as many as 50 years. For example, the Boeing 747 first flew in 1969, and remains in production today, though it has undergone numerous design changes and upgrades. Similarly, the F-15 first flew in 1972 and is still serving the U.S. Air Force. With these product lifetimes, the aerospace business faces a set of challenges unlike those of many other industries. In such long design and development cycles, the technologies many new products use may become outdated by the time they enter the market. Advances in materials, processes, software, and other technologies that emerge once a design has been selected are challenging to implement, causing some products to become obsolete by the time they enter the market. The high development cost makes reactions to new breakthroughs difficult, and changes to proven designs can require extensive validation, removing the incentive to upgrade a product. For example, an aircraft which undergoes a design change to its wing to take advantage of new advances in aerodynamic simulation technology may have to perform the same costly validation procedures as a brand new aircraft introduction. As a result, the introduction of these technological advances may be delayed until the next new product release cycle.

With long product development cycles and service times, aerospace companies are often unable to respond directly to customer demand. Instead, they are forced to anticipate customer needs decades into the future. If firms only reacted to the immediate needs of its customers, large portions of the market that develop in the coming years may be missed and would be difficult to serve. As a result, aerospace companies are forced to determine where their industries are headed over the foreseeable lives of their products, and design solutions to meet their customers' needs before they develop. For example, a commercial aircraft manufacturer needs to analyze and plan for the needs of the industry over the next 20 to 40 years, even before airlines do. They need to examine such variables as the future of point-to-point versus hub-and-spoke travel, the role of cargo in commercial aviation, and the regions of the world likely to experience the greatest growth in airline traffic.

Furthermore, aerospace firms are forced to support their products for many years or decades after the initial sale to meet the needs of their customers. This support may come in the form of spare parts, repair, overhaul, or service, and can make up a large part of an aerospace firm's revenue. For new commercial aircraft purchases, airlines will typically spend an equivalent amount of money purchasing the aircraft as they will on parts and service over its life (Bernstein Global Wealth Management, 2003). On the engine side, GE estimates that for every dollar of value created by new engine purchases, $17 of net present value is created by the sales of parts and services to support it (Bernstein Global Wealth Management, 2003). In many aerospace segments, the Original Equipment Manufacturer (OEM) alone is able to capture the aftermarket value of its products due to industry regulation, customer loyalty, and barriers to entry. This large and mostly exclusive aftermarket creates a significant market opportunity for aerospace companies in terms of both revenue and profit.

The rapid technological change of the aerospace industry paired with the long product development and service lives force firms to support a wide array of applications. Aerospace companies are asked to support technologies that can be several decades old. Obsolescence of supplier parts can be a constant problem due to changing demands of other industries that share components. As a result, firms may have to bring more work in-house because few suppliers remain which still support it, adding cost to the business. Outdated capital equipment may have to be retained to process obsolete material. Select components with very infrequent demand still have to be produced to meet customer needs. Irregular production of these components incurs start/stop costs and can cause logistical problems.

However, even despite these issues, the aerospace aftermarket remains the most significant market opportunity for aerospace firms.

However, with these unique aerospace challenges come some considerable opportunities. The barriers to entry in the aerospace industry are extremely high, with the long development cycles and capital-intensive manufacturing requiring huge amounts of upfront cash in financing. As a result, relatively few firms can compete in each segment of the aerospace industry. The large commercial aircraft business is controlled by Boeing and Airbus, while the engine segment is covered by Pratt & Whitney, General Electric, and Rolls-Royce. Similar trends exist on the defense side of the business and for major subcontractors. With few competitors, an opportunity exists for each firm to enjoy large market shares, and work in closer coordination with its customers.

2.3.2. Pratt & Whitney

Like most of the aerospace industry, Pratt & Whitney relies heavily on aftermarket sales. Spare parts and service make up a large portion of Pratt & Whitney's annual revenue, and at a disproportionately higher profit margin than new engine sales. Due to its long history of successful engine programs, Pratt & Whitney has over 2900 commercial jet engines currently in service (Airline Business, 2010), including those operating within its joint venture programs. P&W engines power more than 30% of the world's passenger aircraft fleet, and such in-service military planes as the F-15 Eagle, F-16 Fighting Falcon, F-22 Raptor, and C-17 Globemaster III (Pratt & Whitney: An Overview, 2011). Some of these aircraft are still flying engines built several decades ago, which require additional spare parts and overhaul services than newer ones, and in turn generate further revenue for the company.

Due to its large revenue and profit contributions, Pratt & Whitney's aftermarket business is a critical area for the company and for UTC as a whole. In order to best serve this market, high customer service rates must be maintained through careful inventory and production management. Many of the aftermarket parts and services are life-limited, requiring replacements or overhauls after a certain number of flight hours or cycles, and can be forecast with some accuracy. Others, however, are elective or the result of failures. In these cases, an inability of Pratt & Whitney to immediately ship a component to a customer may result in that customer being unable to fly an airplane at all, resulting in an Airplane on Ground (AOG). These instances result in huge revenue losses for airlines due to their complete inability to use their asset, the cost of rebooking passengers, additional support needed, etc. As a result, it is critical that P&W maintain high service levels of its aftermarket components in order to meet customer need.

2.4. Problem Statement

Pratt & Whitney is under daily pressure to balance two conflicting priorities in its supply chain: inventory levels and customer service targets. P&W is held to tight expectations by UTC and its shareholders to operate with as little inventory as possible (by overall value). With minimal safety stock, cash is freed up to expand the business, move into new markets, and better utilize its current assets. If cash is tied up in inventory, P&W may lack the resources it needs to make strategic business moves such as bidding on new engine programs and moving into adjacent markets. Additionally, excess inventory can become obsolete if market demand or technology shifts. For example, if P&W holds a large number of a particular fan blade to sell for spare parts, and the demand for this particular product is displaced by other models, all of the cash that P&W put into the entire inventory of that blade is lost. Furthermore, P&W is better able to see the inefficiencies in its operations and identify improvement opportunities when operating with minimal inventory.

However, low inventory levels increase the risk that Pratt & Whitney will experience stockouts. On many of its components that support the high margin aftermarket business, stockouts are unacceptable. P&W's customers are generally unwilling to wait even a small portion of a component's manufacturing time for delivery after order placement, expecting prompt shipment in order to continue serving its passengers. If an airline experiences a failure in an engine part and is forced to wait for a replacement from P&W, the resulting AOG loses all revenues that aircraft would have generated. An event such as this can severely damage the image of P&W in the eyes of its customers, and should be avoided by any means necessary.

In accepting these market forces, the inherent challenge for an aerospace firm is ensuring very high customer service rates without excessive inventory levels. P&W must weigh the expectations from its different stakeholders in order to determine inventory levels of strategic items to hold and customer service rates to target. This thesis provides a tool to determine where and in what amounts inventory should be held in a supply chain. Using this tool, P&W can weigh the effects of varying customer service targets on necessary inventory sizes and locations. With this information, P&W can better plan its facilities, resources, and manpower in support of new engine programs and existing aftermarket activity. Specifically, this thesis will outline the methodology of the tool and give an overview of its application and use to P&W.

3. Inventory Management Literature Review

Locating safety stocks and managing inventory efficiently have been studied at length, with industry experiencing significant resulting benefits. Simpson (1958) develops a system for locating safety stocks in serial supply chains, similar to the successive machining operations of many aerospace components. This methodology can provide intermediate inventory sizes for parts in supply chains with flexible safety stock locations. It can also prove helpful to companies with complex machining operations or to size kanbans in large-scale assembly operations. Inderfurth and Minner (1998) present a system for safety stock optimization in multi-stage supply chains when demand follows a normal distribution. This system allows for very large supply chains (such as the entire BOM of a jet engine) to be analyzed, but is limited to customer deliverables with normal demand, which cannot always be proven. Graves and Willems (2000) give another safety stock optimization model which is capable of analyzing serial or branched supply chains. In doing so, they make no assumptions about demand distributions, allowing the model to use a number of peak demand approximations.

More recently, Miragliotta and Staudacher (2004) propose a method of reducing the effect of large orders unforeseen to the production organizations. To do so, they advocate the sharing of information between the sales team negotiating large orders and the production groups that will manufacture them once received. They perform a risk analysis weighing the benefits of starting production before demand is received against the costs of poor customer service rates and production spikes. Sitompul, Aghezzaf, Dullaert, and Landeghem (2008) propose a safety stock optimization system that takes capacity of each node into account, a strategy that could prove valuable to companies with assets operating near full utilization. Kanet, Gorman, and Stößlein (2010) outline a method in which safety stocks can be updated dynamically, allowing greater alignment with demand, and conduct a survey of U.S. industry to show the benefits of their approach.

4. Approach and Methodology

4.1. Current Inventory Analysis

Pratt & Whitney has a number of engine production and overhaul facilities across the globe. Engine components are produced at many locations, while assembly and test occur at a select few facilities. In addition to internally manufactured components, P&W procures large numbers of parts from its global supply base. Each of these parts requires their own manufacturing lead time (LT), and most have additional sub-tier components with their own procurement requirements. Within this complex supply chain, production scheduling and inventory analysis can become quite difficult. Further complicating this situation, P&W receives its orders in a number of ways. Its military orders, both for new engines and spare components, are typically received at full replenishment lead time and tend to change little over their duration. This structure allows P&W to produce or procure most of its military sales to order, giving it full visibility into its supply chain and the ability to keep military inventory to a minimum.

However, Pratt & Whitney's commercial orders are received in other ways and require different inventory strategies. Both new engines sales and spares orders are frequently placed short to full production lead time, which requires P&W to hold additional inventory to buffer against demand variation. Currently, P&W holds inventory in a number of manners depending on production volume and unit cost. High value items, such as engines and other large, critical subassemblies are built to order. Forecasting demand for these items would force P&W to carry significantly higher inventory values and bear large amounts of risk. If orders were not received in the forecasted manner for these high value parts, P&W would potentially be forced to carry millions of dollars in inventory across financial statement release dates, which would affect UTC as a whole. A similar strategy is used for items with inconsistent or rare demand. If P&W forecasted demand of these items, a portion of them would become obsolete before they are sold, resulting in significant losses.

High volume parts are handled differently, however. The components that Pratt & Whitney manufactures in the highest numbers, such as fan and turbine blades, typically serve the aftermarket in addition to new engine production. They are generally designed to a useful life less than that of the engine and must be replaced with some regularity, usually after a certain number of flight hours or takeoff and landing cycles. This forced replacement, combined with the large in-service base of P&W's engines, produces a high volume market on select P&W products. Due to volume and fairly regular demand, these parts require material buffering. Unlike P&W's military customers, commercial airlines do not place orders far in advance and require delivery far short to full replenishment lead time. To allow for this, P&W buffers with safety stock of only finished goods inventory. For example, to provide safety stock of a high volume fan blade, excess raw materials are purchased and undergo a series of machining operations to become the final deliverable to the customer. Safety stock (sometimes several months' worth of demand) is then held in finished goods form to account for any variation in demand that may occur over the full replenishment lead time of the component.

Other components that are low in value are bought to forecast and then held in inventory until needed. Components that fall into this category are typically not directly deliverable to P&W's customers, and include such low value items as fasteners, rivets, etc. Due to their low value but high impact, P&W has decided that taking a slight increase in inventory levels outweighs the cost of placing frequent orders and experiencing the occasional production stoppage due to a stockout.

4.2. Project Target Areas

High volume components, with their significant impact on aerospace firms, will be analyzed for inventory reductions. While production of these components is used to support both new engine production and the aftermarket, particular attention will be paid to the critical aftermarket segment, which drives the majority of demand. Within Pratt & Whitney, these high volume lines tend to be turbine and compressor blades and other similar products. Targeting these areas allows for safety stock analysis and measurable results. Their fairly steady demand allows for strategic inventory placement that would not be beneficial in low volume lines. Holding strategic safety stock in these situations with high demand variability would result in inventory levels disproportionately higher to the customer service benefits P&W would receive. Additionally, the analysis focuses on components that P&W manufactures internally, due to the increased control P&W has over its own safety stock inventory locations and sizes when compared to those of its suppliers. Within this framework, P&W can still hold stock of supplier goods within its own house, but only after it has received and owns the inventory. Asking first tier and lower suppliers to maintain specified inventory levels would cause P&W to experience significant additional cost and logistical complexity. As a result, this thesis will focus on the supply chains of high volume components manufactured by P&W.

4.3. Inventory Model

Pratt & Whitney's high volume products and their supply chains are an appropriate application for the strategic inventory placement model developed by Stephen C. Graves and Sean Willems at MIT (Graves & Willems, 2000). The methodology behind this model examines a given supply chain, or system of supply chains, and minimizes the overall inventory value carried given a specific customer service target. Due to the machining-heavy nature of P&W's high volume lines, their associated supply chains tend to involve only a few components but high labor content. These simple supply chains are an ideal fit for a spreadsheet-based tool which can be used by any member of the P&W community with little training.

4.3.1. Assumptions

The safety stock optimization model makes a number of assumptions to translate a complex supply chain into a set of mathematical equations. Each can greatly affect the results of the analysis, having large implications on the operations of a company, and must be carefully considered. First, the model uses a set of lead time data for various supply chain events (component procurement time, manufacturing lead time of a component, etc.). All analysis is then performed assuming each lead time is static and completely accurate. As a result, any lead times that are not true to reality can impact the performance of the supply chain and lead to excess inventory or stockouts. For example, if the delivery time of a supplier casting is estimated by a planner to be eight weeks, the model analyzes the supply chain using exactly that lead time as the delivery time for all those castings. However, if this eight week estimate is buffered to account for the longest realistic lead time or the longest past lead time experienced to save the planner any responsibility for stockouts, vast amounts of excess inventory will be held to account for this additional time. Internal manufacturing lead times can be influenced the same way. Therefore, it is important that all time data used in the model is completely accurate and not buffered to defer accountability. Lead time variability can, however, be factored into the model and will be considered later.

Customer demand, by its very nature, is uncertain, and very difficult for any mathematical model to evaluate. The model used here can incorporate a number of forecasts to predict demand, but all of them involve some prediction of the future which may or may not prove true. For stable product lines with constant or near-constant demand, the best method for predicting future performance is often based upon historical need. If demand appears to be normally distributed, its mean and standard deviation can be obtained over a given time period, and, along with a customer service target, be used to predict an upper bound of demand. For new programs or product lines with inconsistent demand, an estimate of the upper bound of demand can be made based on market research or similar predictions. As historical data becomes available and programs mature, these market estimates can be phased out in favor of historical distributions. No matter the method, however, the model requires some prediction of future demand upon which the analysis is based. Unforeseen spikes or drops in demand will not be able to be explained by the model, and will yield skewed results.

Furthermore, the model yields optimum safety stock locations and sizes, but does not analyze work-in-progress assemblies (WIP) or material flow. All analysis performed by the model is based upon the lead time data it is given, which must include all necessary transit times, order delay times, and production sequencing activities. Inconsistencies in the analysis may occur if certain parts are taken out of the work flow and quarantined due to a defect or other issue. Often, manual intervention is used to resequence work flow, giving priority to more visible or critical programs. This practice can cause the lower priority components to back up and their manufacturing times to greatly increase, which in turn can cause safety stocks to dry up and potential stockouts to occur. The factors affecting work flow and production sequencing cannot be directly analyzed by the model, which assumes that they are all built into the given lead times.

4.3.2. Methodology

In order to analyze a complex supply chain, the model breaks the entire procurement and manufacturing process into a series of "nodes", each of which represents a distinct step or activity. These nodes can vary greatly in scope, and are specific to the level and complexity of the supply chain being analyzed. At Pratt & Whitney, many of the nodes used represent large operations, such as the whole procurement process of a supplier part or the entire in-house manufacturing time of a blade. Essentially, each node is a part number and a new node is created when a component or assembly changes part numbers. This distinction was chosen because the greatest sources of inventory control exist at these interfaces. Physical boundaries (between manufacturing departments or facilities) and virtual distinctions (in inventory tracking systems, etc) occur between these different part numbers. At P&W, inventory can only be controlled once supplier components are received in-house because of the lack of ownership of inventory at suppliers' facilities. However, other manufacturing processes may require different node granularity. For some simpler components or assemblies with irregular value-add profiles, the correct nodes may be individual manufacturing operations (grind, deburr, assemble, etc) rather than part number interfaces to provide the lowest overall inventory value.

These nodes are then arranged into a tree system, showing the bill of material (BOM) of the entire product. The final deliverable (at the end of the tree) is assigned a maximum service time to the customer, which represents the time frame in which the customer expects delivery from the time the order is placed. Working backward in the node tree from this final deliverable, each successive node is assigned a guaranteed service time to its downstream nodes and an inbound service time (which accounts for the lead times of its successive upstream nodes, allowing the operation to collect all components necessary for production to commence). Safety stocks are then located and assigned to cover the peak demand likely to occur over the lead time of that given node. This process is repeated until all nodes in the model are covered and assigned safety stocks. Several service time initial conditions may have to be tested to achieve the global optimum, with all service times set to zero or to the node lead times often showing success. Once reached, however, this optimum will yield safety stocks sufficient to achieve the desired customer service rate with the lowest overall inventory value.

This optimization function is then minimized over the feasible space to obtain the ideal safety stock locations and sizes. In mathematical terms, the optimization function is given below (Graves & Willems, 2000).

$$\sum_{j=1}^{N} V_j \left\{ D_j\left(SI_j + T_j - S_j\right) - \left(SI_j + T_j - S_j\right)\mu_j \right\}$$

Equation 1: Inventory Model Optimization Function

Where: V_j = unit inventory value at node j
SI_j = inbound service time to node j
S_j = guaranteed service time of node j

D_j = peak demand per time period experienced by node j
μ_j = average demand per time period experienced by node j
T_j = lead time of node j
s_j = maximum service time for demand node j

Subject to the constraints: $S_j - SI_j \leq T_j$

$SI_j - S_i \geq 0$

$S_j \leq s_j$ for all demand nodes j

S_j, $SI_j \geq 0$ and integer

The optimization program minimizes the inventory value held by every node in the supply chain, from j=1 to N. In this way, the safety stock held is actually the difference between the peak and mean demand over the net replenishment time of each node. The given constraints restrict all solutions to possible conditions in the supply chain. For example, the first constraint ensures that net replenishment time of each node is greater than zero, and the second forces each node to wait until all previous nodes have delivered materials before it can commence production. Using this set of mathematical equations, the ideal safety stock locations and sizes can be obtained.

This model formulation is a nonlinear optimization problem, which can be solved by a number of commercially available software programs. At P&W, all calculations were entered into a spreadsheet, and the optimization problem was solved using an add-on program. The spreadsheet format (see Appendix 1) was used at P&W to maximize usefulness of the product because it is available to all members of the organization, increasing the likelihood that the tool will become widely used. If the tool is built in a more powerful but sparsely available specialized software package, its adoption would be slowed and its traction in the organization diminished. In the relatively simple supply chains of

P&W's spare parts business, the spreadsheet add-on platform is sufficient to properly solve the nonlinear optimization problem and common enough to gain support within the company.

5. Analysis and Results

5.1. Application

Perhaps the best way to illustrate the usefulness and effectiveness of the safety stock optimization tool is to review an example of its application. An example of a P&W high volume product is shown in the sample supply chain Figure 1.

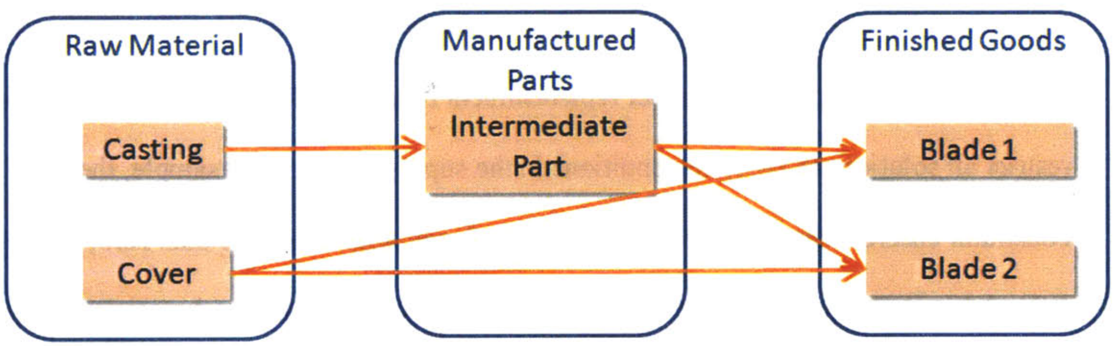

Figure 1: Sample P&W High Volume Supply Chain

This supply chain is typical of many of the high volume parts at P&W. This product has two different blades as end deliverables to customers. Material inputs are few (only two in this case) and much of the inventory value is added through machining and labor. The casting is machined into an intermediate part with its own unique part number, which is essentially a break in the manufacturing process because the other supplier component, a cover, is not added until the final blade. Each of the parts in Figure 1 represents a node in the optimization model where safety stock could be held. The raw materials are bought from suppliers and would be held as safety stock once received at a P&W facility. All further operations are performed within P&W, allowing easier inventory control.

The static lead times for each component in this supply chain and their associated inventory value at each point in the manufacturing process are shown in Table 1.

Part	Inventory Value	Quoted Lead Time (weeks)
Blade 1	$400.00	4.0
Blade 2	$425.00	4.5
Intermediate Part	$250.00	2.0
Cover	$2.00	1.0
Casting	$75.00	8.0

Table 1: Component Inventory Values and Lead Times

As this data shows, the inventory of the final deliverable product to the customers, the blades, are more than five times the inventory value of the raw materials. This significant difference creates an opportunity to hold inventory in locations of lesser value. P&W values its inventory in two ways. For purchased parts (such as the casting and cover), it is valued at the purchase price from the vendor. For parts manufactured in-house, inventory is valued based on the labor or machine time added to the part plus the cost of any supplier components. Labor and machine hours both have associated standard hourly rates based on labor grade and type of machine. P&W's manufacturing IT system tracks the real-time machine and labor content added to a given part and calculates the inventory value.

Furthermore, as shown in Table 1, the majority of the full replenishment lead time of the entire supply chain is in the procurement of the casting. From the time P&W places an order, two months elapse before they actually receive the casting. In comparison, the critical path of the entire subsequent supply chain is only 6.5 weeks. By having the bulk of the lead time built into a lower value area of the supply chain, safety stock may be able to be held there at a lower rate to cover more expensive areas of the supply chain. The cover, at just $2 and with a one week lead time, has little impact on the supply chain in terms of both price and schedule.

The model runs assuming all demand and production occurs within a specified time period, with the length of time having significant impact on the usefulness of the results. In this case, a time period of one week is used to yield the most meaningful results. If longer time periods are used, the variation in demand may diminish, but this length of time is too long to accurately model supply chains where some components have lead times as short as one week. As with most mature products, these two blades have relatively stable demand, with the standard deviation of weekly demand rarely varying more than 50% week to week. While this may seem volatile, demand is rarely zero in any one week and spikes are infrequent. For the model, two years' worth of demand data was used to give an accurate representation of the production needs for this project.

The safety stock optimization model requires some assumptions about demand, but these can be given in a number of ways. In short, it requires a "peak" demand that the end items are likely to see any given time period. If demand exceeds this limit, insufficient safety stock may exist, and stockouts may occur. However, this "peak" demand can be determined in a number of ways, and it even may be appropriate to estimate the peak demand when poor forecasts or little history exists. Executives may conservatively set these estimates, or they may be set to align with production capacity. Where sufficient demand history exists and is likely to align with future orders, analysis of historical demand may yield better results. A common way to set the peak demand based on historical data is to specify a service level based on a normal distribution. This way, a company can determine that it wants its safety stocks to buffer for demand with 90%, 95%, 99%, etc certainty.

To check if demand data follows a normal distribution, several steps must be taken. If demand of a particular product is normally distributed, a histogram of it should follow a bell-shape curve (Vining & Kowalski, 2006). If demand of a particular product from a time period is represented by y, the data must be arranged in ascending order such that $y_1 \leq y_2 \leq y_3 \cdots \leq y_n$. The cumulative probability point, $P(i)$ is then

$$P(i) = \frac{i - 0.5}{n}$$

Equation 2: Cumulative Probability Point Calculation

where i is the rank of a particular demand value. The cumulative probability points from Equation 2 are then plotted on y-axis, with each demand figure on the x-axis. A sample normal probability plot from a Pratt & Whitney product line is shown in Figure 2.

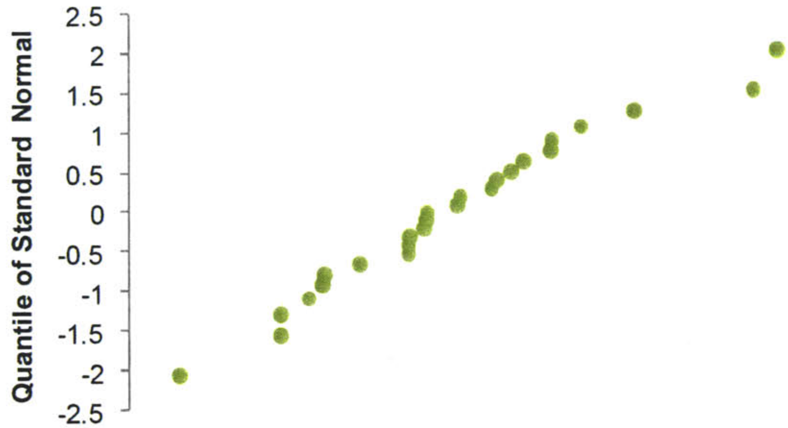

Figure 2: Normal Probability Plot of a Sample Product

If the data graphed in the normal probability plot is roughly linear, then the demand can be assumed to follow a normal distribution. A common test for linearity is to place a "fat pencil" on the page and ensure that all data points are covered by the pencil. This simple test accounts for any slight imperfections in the data, as no demand is perfectly normal. The data shown in Figure 2 does appear to follow a normal distribution. All data points are roughly linear and pass the "fat pencil" test. If multiple data points fell outside this linear trend, or if the graph was blatantly nonlinear (such as parabolic), the normal distribution could not be used to analyze the demand of this product.

Once demand data is confirmed as following a normal distribution, the demand limit can be estimated by

$$Demand\ Limit = \mu + z\sigma$$

Equation 3: Demand Limit Calculation

where μ is the mean demand per time period, z is the standard normal factor corresponding to a customer service rate, and σ is the standard deviation of demand per time period. With this demand limit, the model can now be used to determine the optimum safety stock locations and sizes.

Another important consideration in the model is the service time for the demand nodes. This factor represents the amount of time a company has to ship a product once it receives an order. For example, if an order is received on day one, but the customer does not require shipment for another two weeks, the service time of the demand node would be 14 days. This service time to the customer is typically determined by customer expectations and market trends. In some industries with highly customized and low volume products, long customer service times may be acceptable. Aircraft, for example, are typically built to order for a specific customer and are unique to other planes made by that company. In this case, the airlines expect a long service time for their product, allowing the aircraft manufacturer the freedom to not hold safety stocks of finished aircraft. However, the business environment of Pratt & Whitney's aftermarket does not allow for this type of long customer service times. Typically, parts sold as spares need to ship immediately because they prevent a customer from utilizing their asset. For this reason, the customer service times in all of the P&W supply chains studied were set to zero to model a ship-from-stock environment. As a result, safety stock may have to be held at the finished goods stage of every product line.

5.2. Model Results

The safety stock optimization model yielded a number of interesting insights into Pratt & Whitney's high volume spare components supply chains. With the parameters in

Section 5.1 entered into the spreadsheet-based optimization program in Appendix 1, a detailed safety stock map is obtained. A summary of the results is shown in Figure 3.

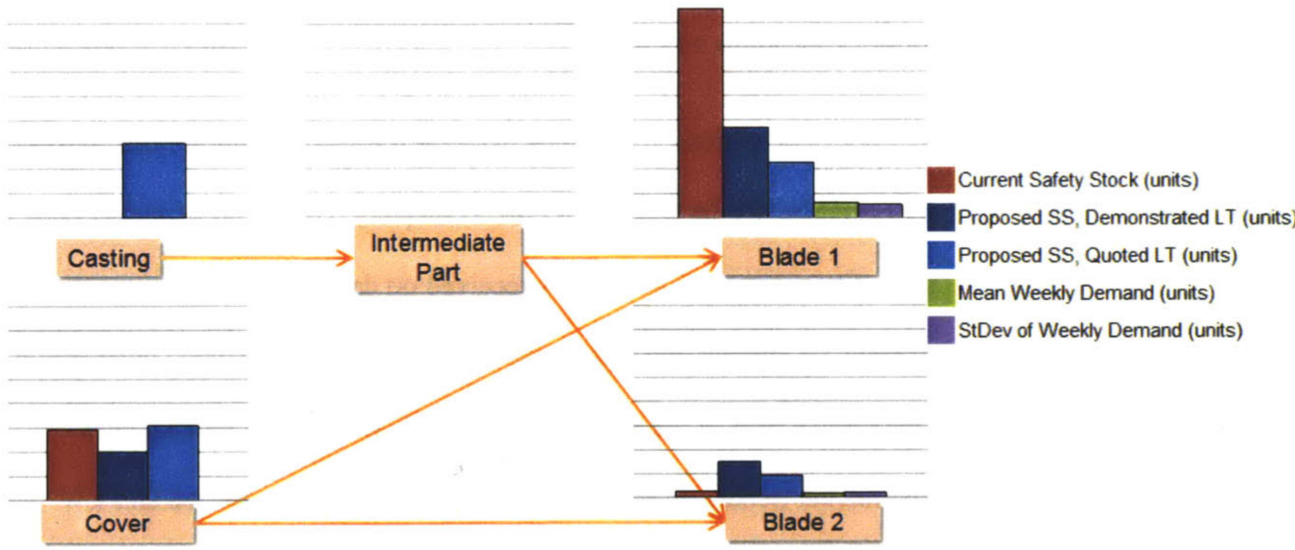

Figure 3: Example Safety Stock Optimization Results

As shown, P&W previously held large amounts of safety stock at finished goods. While this practice may have provided very high service levels, it forced the company to carry high inventory levels. As a reference to the safety stock targets, mean weekly demand and the associated standard deviation of weekly demand are shown in Figure 3. For the Blade 1 product, for example, P&W is holding many, many weeks' worth of demand in safety stock. These high levels are necessary to buffer against the variation demand across the full replenishment lead time of the entire supply chain. However, the unit value of the blades at this point in the supply chain is \$400, compared to \$75 at the casting level. As a result, the location choice of this safety stock forces P&W to carry much more inventory (by value) than is necessary. At the time of analysis, Blade 2 was still gaining market share, and not experiencing the same demand levels as Blade 1.

Little historical demand data was available, so the initial safety stock level was manually set low. As also shown in Figure 3, safety stock of the cover was carried due to its extremely low inventory value and short lead time. As with some parts, it is simply easier to carry higher inventory levels of these parts than to order them just-in-time because of their low inventory impact and ordering and logistics cost.

In this model, two different parameters for the lead time of each component are used to obtain a better frame of reference for the safety stock results. First, the model was run with the lead times for all components set to their "quoted" lead time, which is the static measure direct from P&W's Material Requirements Planning (MRP) system. These numbers are used to plan material requirements and order dates, but are updated infrequently and do not take into account quality or supplier issues. For purchased parts, the quoted lead times account for all time from the order placement date until the shipment reaches P&W and is ready for use. For parts manufactured in-house, the quoted lead time includes any necessary manufacturing and transportation times.

However, upon further analysis, the quoted lead times of various manufacturing steps did not always match the real process times. In reality, each of these lead times varied occasionally and disrupted the supply chain. The organization requested a conservative lead time value to protect against stockouts. Upon further inspection, the historical lead times of each component approximately followed a normal distribution, so a "demonstrated lead time" figure was created with a 95% certainty level. These lead times are buffered to allow for the natural supply chain variation, and tend to be longer than quoted lead times of a particular component. Actual historical lead time data was not available for purchased parts, so only nodes in the supply chain that are performed internal to P&W are able to use these demonstrated lead times.

As Figure 3 also shows, safety stock has been lowered in the later stages and moved earlier in the supply chain. In this case, fewer units of finished blades, specifically Blade 1, are necessary to buffer the supply chain against demand variability by providing additional safety stocks of castings. Figure 3 shows safety stocks in units, where the real value of this proposed safety stock strategy may not be apparent. However, by moving the inventory from Blade 1 finished goods, at a value of $425 per unit, to castings, at $75 per unit, the real cost savings becomes apparent. As shown, the safety stocks that were calculated from the model using demonstrated lead times tend to be longer than those using quoted lead times, but both still show significant savings over the current state. Specifically, an inventory reduction of over $1.2M is possible in this supply chain alone by moving the inventory as shown.

In the particular supply chain analysis shown in Figure 3, a few observations can be made which appear counterintuitive. First, in the case of the cover component, the model recommends lower safety stock when using demonstrated lead times than when using quoted lead times, when the opposite is usually true. Pratt & Whitney was more willing to hold inventory of the covers than other items because of its low inventory value. Second, Blade 2 shows higher recommended safety stock levels than P&W is currently holding. As mentioned earlier, this product is still ramping up production and gaining traction in the marketplace, so little historical demand data was available. As such, initial safety stock levels were set manually based on sales estimates, when additional levels were actually necessary to account for demand variability.

5.3. Sensitivity Analysis

Once an organization optimizes its safety stock levels and locations, continuous improvement efforts are often initiated, but choosing in which area to concentrate resources can be difficult. The improvement options on any given supply chain are almost endless. An organization can work with its suppliers to lower the unit cost of parts, reduce lead time, or even decrease lead time variability. Additionally, a firm can look inward to try to reduce its own manufacturing time, variability in its manufacturing time, or service time to its customers. It can even examine its own inventory valuation system to see if the labor or machine rates applied to components are excessive, or add value incrementally (at part number interchanges, for example) rather than the strict hourly value-add time that Pratt & Whitney uses. Some companies also attempt to alter their demand patterns. For example, demand can be smoothed out to reduce variation, but this practice damages customer service rates, and cannot be considered for P&W's high volume aftermarket business. However, some industries with long customer service times, especially those that use a make-to-order system, such as new aircraft production, can effectively smooth out demand and carry little safety stock, but at the cost of customer service times.

While improving all areas of a supply chain would be ideal, most companies do not have the required time or resources. With so many areas to focus on, organizations often have a difficult time determining where to start. It can be challenging to decide if efforts should begin with their supplier base, or with internal manufacturing processes. Once these decisions are made, choices must then be made between lead time, variability, and cost reductions. Furthermore, some companies have very little leverage over their suppliers, either because they are a small percentage of their overall business or buy sole-sourced parts. In these cases, improving the vendor area of the supply chain may be difficult, so firms can only concentrate on the internal parts of their supply chain. By eliminating improvements areas that are not feasible, it may become clear to many firms where to concentrate improvement efforts.

However, even when improvement projects are not limited to internal processes, the choice of where to start can be difficult. For the supply chain outlined in Figure 1, a number of options are possible. Internal Pratt & Whitney manufacturing processes could be examined to shorten their manufacturing time, allowing less safety stock to be held to cover the uncertainty in demand. Efforts could be focused on bringing the cost of inventory down by reductions in labor rates, energy input, machine costs, chemical costs, facility charges, supervision, etc. Additionally, the variability in manufacturing time can be a large driver of demand, but can be reduced by looking at defect rates, which cause large delays in processing time, and unplanned machine downtime, which interrupts the flow of material and causes fluctuations in lead time. Figure 4 shows the effects of each of these improvement efforts.

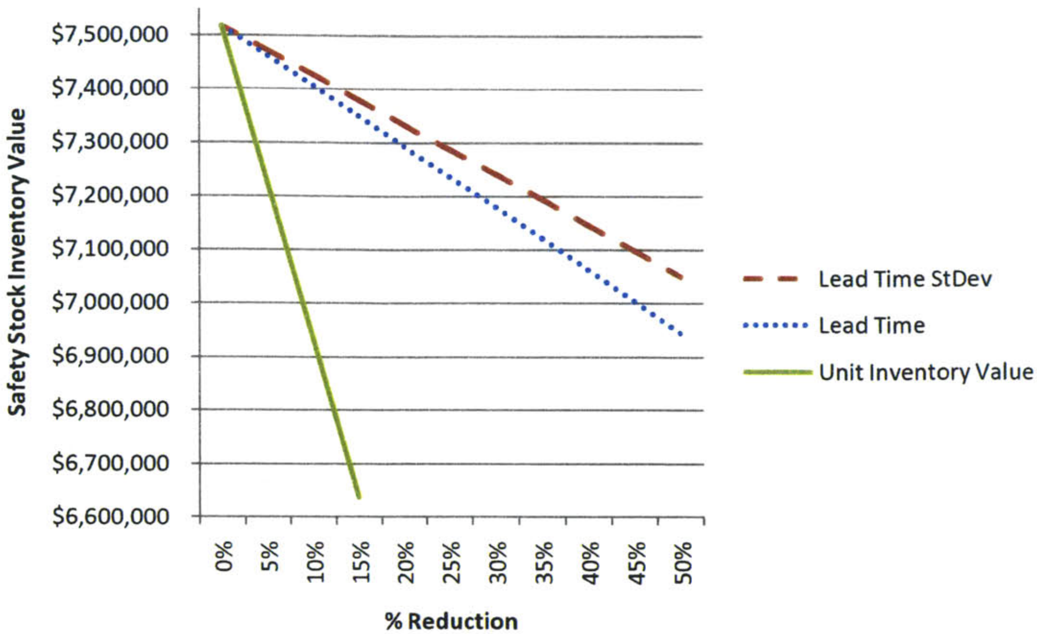

Figure 4: Internal Improvement Effort Analysis for Blades

As shown, each improvement effort does result in overall safety stock inventory reductions, but with varying success. The safety stock values shown on the vertical axis are from a nominal case. To obtain this data, the safety stock model was run successive times to yield the overall supply chain safety stock value, while reducing the lead times, lead time standard deviation, and unit inventory value of Blade 1 and 2 by 5% reduction increments. The improvement efforts in each of these three areas resulted in a linear relationship between improvement and safety stock value, with unit inventory value being the most effective by far. Lead time improvements had a greater impact than decreases in its standard deviation, but only a 10% reduction in unit inventory value had a greater effect on the overall inventory value than a 50% reduction in the other factors. This significant difference is intuitive, however, because of the large number of blades held in safety stock (approximately 2900 and 1980 units of Blades 1 and 2, respectively) and their higher inventory values. By reducing the high inventory value of these finished goods by 10%, cost is saved on each one of these units in safety stock, and at the highest cost location in the supply chain. Meanwhile, reducing the lead times and LT standard deviations allow the supply chain to hold fewer units in safety stock because it reduced the interval over which demand variation must be covered, but these units are still held at the higher inventory level. As a result, if improvement efforts can only focus on one area of the supply chain, inventory cost reductions should be addressed.

In some cases, companies may have better control over their suppliers or little room for improvement in their own manufacturing processes due to past advances in productivity. In these cases, it may be easier for a company to work with its suppliers on their processes in order to make the relationship more valuable and cost effective for both parties. This relationship is more likely to be effective when a close partnering history exists between the two firms and their future success is dependent on one other. In such a case, Figure 5 shows the effectiveness analysis of each improvement method at the supplier level.

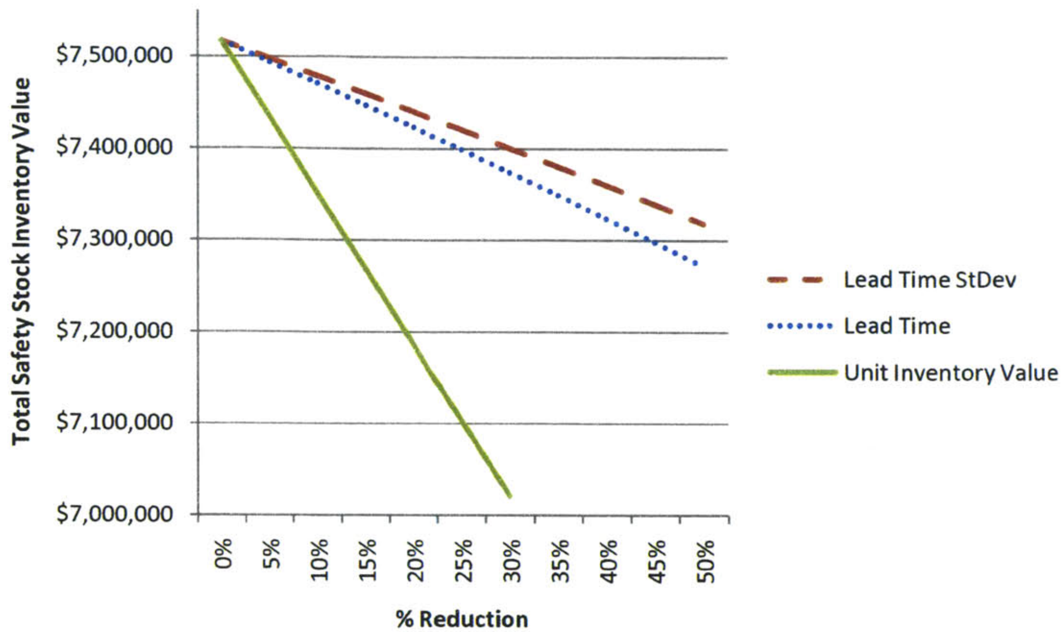

Figure 5: Supplier Improvement Effort Analysis

As shown, the effects of each improvement strategy at the supplier level are similar to those internal to P&W in Figure 4. Cost reduction efforts hold significantly more value than the other two, but by a lower margin. Here, a 15% reduction in casting unit cost results in a lower overall inventory value than a 50% reduction in either casting lead time or LT standard deviation. This result, again, is intuitive because of the lower inventory value of the castings when compared to the finished blades. Additionally, the casting lead time, at eight weeks, is significantly longer than that of the blades, causing reductions in the casting lead time or its variation to result in greater inventory savings than similar blade efforts. This increased effectiveness is due to the additional safety stock necessary to cover the longer lead time and demand variation of the castings. However, despite the additional value of the lead time reductions, improvements in the unit inventory cost are still the most effective choice.

Considering that inventory unit cost is a significantly better choice on which to focus improvement efforts, Figure 6 compares the inventory savings resulting from unit cost reductions in the blades and castings.

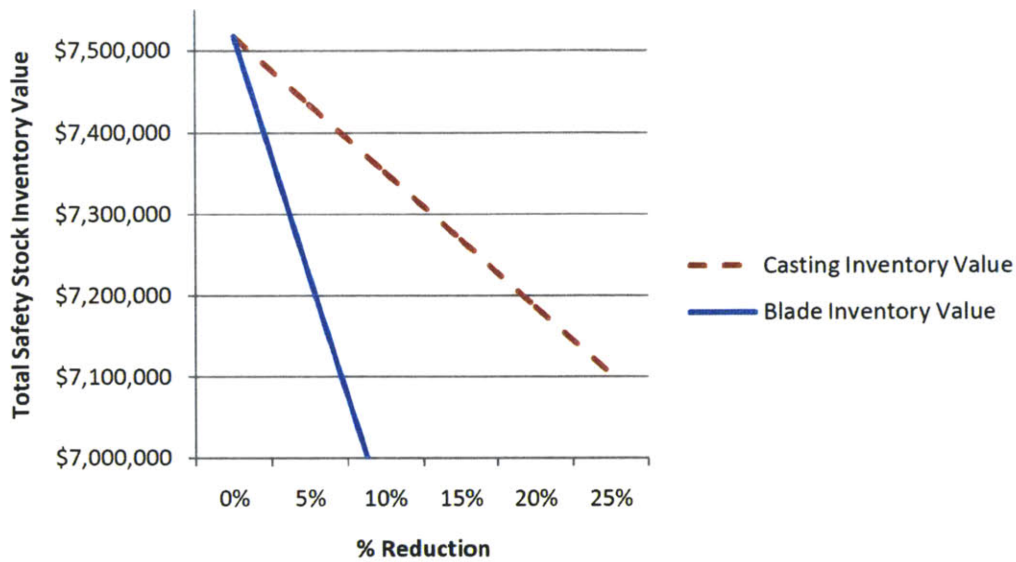

Figure 6: Casting and Blade Cost Reduction Comparison

As shown, reducing the unit inventory value of the blade yields greater returns than that of the casting, primarily due to its higher initial cost. This is true of all supply chains that follow a similar trend of increasing inventory value as parts flow closer to finished goods, and is not specific to this one P&W supply chain. As a result, any improvement efforts should be focused as close to the end of the supply chain as possible because of the high per-unit inventory values associated with this location. In some supply chains, cost reductions at finished goods may not be possible, so efforts should be undertaken moving back in the supply chain one node at a time from customer delivery.

However, when possible, the most effective means of reducing safety stock is through demand smoothing. Demand variability is the one supply chain measure that truly drives the need for safety stock. Without demand variability, there would be no need to carry safety stock because the demand experienced in every time period would be exactly the same as the last, with no chance of sudden spikes or drops in demand. As a result, production could be completely level-loaded, making it more economical in a number of ways. Labor charges would fall because there would be no need for overtime. Machine capacity could be sized to align with the static demand, eliminating the need to purchase additional capacity to be able to respond to demand spikes. However, demand smoothing is not possible in many industries because it drives up customer service times. When not possible, even slight decreases in demand variability can yield drastic declines in necessary safety stock levels. Figure 7 shows the same effects of reducing unit inventory cost of the casting and blades as Figure 6, with the additional factor of reducing demand standard deviation.

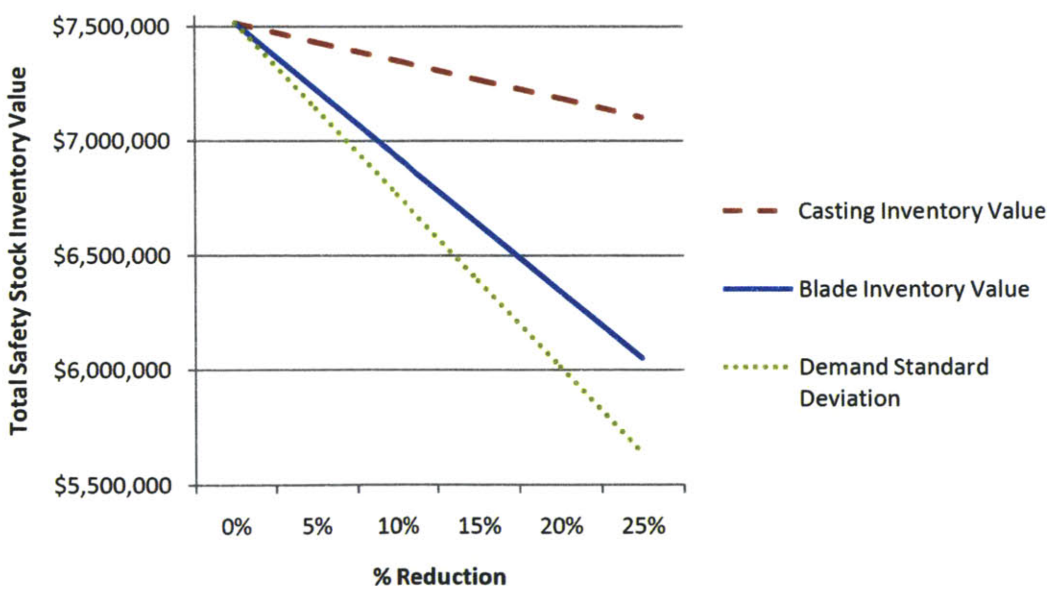

Figure 7: Unit Inventory Cost and Demand Smoothing Analysis

As shown, reducing demand variability is the single most effective measure that can be taken to reduce the necessary safety stock carried by a supply chain, even more so than reducing unit inventory values at the most expensive point. While this technique may not be feasible in such markets as the aftermarket aerospace industry in which Pratt & Whitney operates, other companies may find this to be a viable strategy. If the benefits are sufficient, production could be kept at a level rate in anticipation of future demand. Even slow production adjustments such as specifying one production rate for a given time period can be very beneficial. The automotive industry works under this production strategy. Automobiles are produced at a specified rate, which can be changed in the long run, but not in immediate response to customer demand. Lower safety stock is then possible, but customers may have to wait longer for their automobiles if none are available that match their needs. However, when demand cannot be adjusted or controlled, improving the supply chain at its costliest point, finished goods, is the ideal place to begin improvement initiatives.

5.4. Implementation

5.4.1. Challenges

Once inventory recommendations are made, implementing them can be another matter altogether. Significant amounts of discipline are typically necessary for an organization to bring in additional material and even more to reserve it as safety stock. In any large organization supply chain, the most attention is typically paid at the latest stages. The business is driven by customer demand, causing the organization and its personnel to focus on finished goods. In the example of a jet engine blade, a customer order triggers an employee to look at the number of completed blades. If these are not available, the next upstream node is then examined for availability or potential issues. By this method, large amounts of attention are paid to the last few stages in a supply chain. However, by ignoring the earlier stages, organizations take a shortsighted view of the process. In fact, the earliest stages of a supply chain, typically consisting of suppliers and their sub-tier suppliers, are just as important to the overall success of a product line as the final few process steps before a component is shipped to the customer.

Frequently, as demand peaks, organizations will signal the supply chain to produce additional components, even if sufficient safety stock exists to accommodate the surge. This is known as the "bullwhip effect" (Wangphanich, Kara, & Kayis, 2010), resulting in huge amounts of excess and wasted inventory. To counteract this, the further upstream nodes have to be disciplined to only produce to order, and the downstream nodes cannot request more than their customer demand. Upstream nodes frequently feel as though they need to produce as long as they have material to keep the "more important" downstream nodes running. However, this practice is actually contrary to the purpose of safety stock. The upstream nodes have to show restraint to hold large amounts of raw material inventory and not process it, even if workers or machines are idled. Producing excess material, even to keep other nodes running, only increases the inventory an organization is forced to carry and the likelihood it will become obsolete. As a result, in order for safety stock to be held effectively, manufacturing groups, especially upstream ones, must remained disciplined enough to not process any material until demand is received.

5.4.2. Business Environment Considerations

The model outlined here assumes an ordering system that may be contrary to that currently used by many aerospace companies. Its analysis is based on a supply chain in which customer demand is received at the end node, and then necessary components are ordered from upstream nodes in the same time period. In this system, no advance production planning or time fences are needed and all production reacts only and immediately to demand. The lead time of each manufacturing or procurement stage, however, is factored into the safety stock that each node carries. By reacting immediately to customer demand, the need for complicated planning systems is eliminated, and all necessary buffers are built into the optimized safety stock.

The MRP systems that most aerospace companies use, however, can be in conflict with the simple ordering strategy employed by the inventory optimization model discussed here. These MRP systems are rigid computer structures that plan production and procurement activities using the final component requirement date then determining all other necessary material requirement dates by their lead times. In doing so, MRP systems require long planning times and react poorly to shifts in customer demand. In fact, they work best when demand is placed far in advance and components do not have to be shipped from stock. Changes can be made to most MRP systems to allow for better response to customer demand, but many others still struggle. If a company's IT systems are not aligned with its supply chain strategy, neither are likely to succeed. For a safety stock optimization strategy such as the one outlined here to become widely used in an organization, the production and procurement planning system used must be flexible enough to allow for immediate reactions to customer demand and not burden an organization with long planning time fences and rigid production schedules.

5.4.3. Possible Areas for Further Projects

While this project addresses the needs of Pratt & Whitney's high volume supply chains, significant opportunity exists for further projects. Although the high volume aftermarket makes up a considerable portion of P&W's revenue, many other critical production areas operate at far lower rates. The safety stock optimization model discussed here is of little use in these lower volume areas due to their uncertain demand. Additional projects could address these lower rate areas, such as engine production and overhaul. These programs make up a large portion of the inventory value that P&W carries because of the longer lead time of each assembly and higher per-unit inventory values. Future projects could address these lower volume areas and seek to minimize the necessary inventory value in make-to-order systems.

Vendor management is another source of supply chain uncertainty that is not addressed in this thesis. In the safety stock optimization model discussed, supplier lead times are held as completely accurate with no allowance for variability or interruptions. In reality, P&W demand may exceed supplier capacity or the vendors themselves may have internal issues that cause shipment delays. By examining supplier stability, capacity, and past performance, tools could be developed which provide a more robust view of the supply chain and allow for variation. Future projects could work more closely with the supply base to identify improvement opportunities and reduce cost for themselves as well as P&W.

6. Recommendations and Conclusions

The dynamics of the aerospace industry present a number of challenges unlike those faced by many other industries. Lead times of components are very long, and suppliers with the necessary capabilities can be few. Capital equipment and tooling costs are substantial, creating high barriers to entry for new competitors and putting significant pressure on existing firms. Many components, such as the jet engine blade discussed here, need to be shipped immediately when sold because the availability of much larger assets depend on them. These aftermarket parts are typically sold on large margins, and make up significant portions of a firm's profits. The aerospace industry as a whole also tends to be cyclical, further adding to the challenges the few firms in each segment experience.

However, despite these challenges of the aerospace industry, rewards for competing firms can be high. In the high volume aftermarket business, high customer service rates and low inventory levels can be crucial to a firm's success. The safety stock optimization model discussed here addresses both of these needs. It provides the ability to specify varying customer service rates (90, 95%, 99%, etc), the ability to ship in any time window, and can use demand limits in a variety of forms. By using basic manufacturing and procurement parameters, such as lead time and inventory unit cost, safety stock levels and locations across an entire supply chain are optimized to provide the lowest overall inventory value. By reducing these inventory levels without sacrificing customer service rates, firms are able to remain competitive and better respond to customer demand.

As important as the safety stock optimization tool itself is the ability to pinpoint the most efficient locations in a supply chain in which to make improvements. By simply trying to improve all areas of a product's supply chain, firms may spread their resources too thin, generating benefits in some areas while experiencing little to no returns in others. However, by using the safety stock optimization model to analyze the ideal areas of a supply chain to target, improvement efforts become their most effective. As shown, improvement initiatives tend to be most efficient when directed at areas of high inventory value, such as finished goods, and at reducing unit inventory cost. Reducing raw material cost, lead time, or LT variability do not have as significant an effect on overall safety stock inventory as finished goods improvements. By instituting an optimized safety stock structure and then choosing specific areas of the supply chain upon which to improve, an aerospace firm can remain competitive in its aftermarket industry and better serve the needs of its customers.

Appendix 1: Safety Stock Optimization Model Interface

Inputs

Confidence Interval that No Stage will Stock Out	95%

	Blade 1	Blade 2
Mean Weekly Demand, μ_j	500.00	600.00
Standard Deviation of Weekly Demand	200.00	250.00
Maximum Service Time for Demand Node (weeks), s_j	0	0
Correlation of Demand (0=Independent, 1=Completely Correlated)	0	

Legend
General Inputs
Part-Specific Inputs
Safety Stock Values
Decision Variables
Used in Calculations
Objective Function

Calculations

	Part Number	Average Lead Time (weeks)	StDev of Lead Time (weeks)	Inventory Value / Unit	Effective Lead Time, T_i	**Safety Stock, $E[I_j]$**	Safety Stock Value	Inbound Service Time (weeks), SI_j	Guaranteed Service Time (weeks), S_j or S_{ij}	Net Replenishment Time (weeks), $SI_j + T_j - S_j$	Demand Over Net Replenishment Time, $D_j(SI_j + T_j - S_j)$	$\mu_j(SI_j + T_j - S_j)$
Level 3	Casting	8.0	0	$75.00	8.00	1489.48	$111,711	0	0	8.00	10289.48	8800
Level 2	Intermediate Part	2.0	0	$250.00	2.00	0.00	$0	0	2	0.00	0.00	0
	Cover	1.0	0	$2.00	1.00	0.00	$0	0	1	0.00	0.00	0
Level 1	Blade 1	4.0	0	$400.00	4.00	805.81	$322,324	2	0	6.00	3805.81	3000
	Blade 2	4.5	0	$425.00	4.50	1048.39	$445,567	2	0	6.50	4948.39	3900

Optimization Function

$879,602

Constraints

		Constraint 1			Constraint 2			Constraint 3			Constraint 4
	Part Number	$S_j - SI_j$	≤	T_j	$SI_j - S_i$	≥	0	S_j	≤	s_j	$S_j, SI_j \geq 0$
Level 3	Casting	0	≤	8.00							
Level 2	Intermediate Part	2	≤	2.00	0	≥	0				
	Cover	1	≤	1.00							
Level 1	Blade 1	-2	≤	4.00	0	≥	0	0	≤	0	
	Blade 2	-2	≤	4.50	0	≥	0	0	≤	0	

Bibliography

Bernstein Global Wealth Management. Aftermarkets: Staging a Comeback in 2004. (2003). *Black Book - U.S. Multi-Industry: Commercial Aerospace in a Holding Pattern*, 72-87.

Graves, S. C., & Willems, S. P. (2000). Optimizing Strategic Safety Stock Placement in Supply Chains. *Manufacturing & Service Operations Management, 2*(1), 68.

Inderfurth, K., & Minner, S. (1998). Safety stocks in multi-stage inventory systems under different service measures. *European Journal of Operational Research, 106*(1), 57-73.

Jet Engine Market Statistics. (2010). *Airline Business, 26*(4), 54-54.

Kanet, J. J., Gorman, M. F., & Stosslein, M. (2010). Dynamic planned safety stocks in supply networks. *International Journal of Production Research, 48*(22), 6859-6880.

Miragliotta, G., & Staudacher, A. P. (2004). Exploiting information sharing, stock management and capacity oversizing in the management of lumpy demand. *International Journal of Production Research, 42*(13), 2533-2554.

Pratt & Whitney: An Overview. (2011). Retrieved January 2011, from Pratt & Whitney corporate website: http://www.pw.utc.com/About+Us/An+Overview.

Simpson Jr., K. F. (1958). In-process inventories. *Operations Research, 6*(6), 863.

Sitompul, C., Aghezzaf, E., Dullaert, W., & Landeghem, H. V. (2008). Safety stock placement problem in capacitated supply chains. *International Journal of Production Research, 46*(17), 4709-4727.

Vining, G., & Kowalski, S. M. (2006). *Statistical Methods for Engineers.* Belmont, California: Thomson Higher Education.

Wangphanich, P., Kara, S., & Kayis, B. (2010). Analysis of the bullwhip effect in multi-product, multi-stage supply chain systems - a simulation approach. *International Journal of Production Research, 48*(15), 4501-4517.